

^zVirtuozzo

Virtuozzo Storage 2.3

User's Guide

December 12, 2017

Virtuozzo International GmbH

Vordergasse 59

8200 Schaffhausen

Switzerland

Tel: + 41 52 632 0411

Fax: + 41 52 672 2010

<https://virtuozzo.com>

Copyright ©2001-2017 Virtuozzo International GmbH. All rights reserved.

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

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

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

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

Contents

- 1. About Virtuozzo Storage 1**
 - 1.1 Supported Storage Types 1

- 2. Accessing Virtuozzo Storage S3 Buckets 2**
 - 2.1 Managing Buckets via Virtuozzo Storage User Panel 2
 - 2.1.1 Logging in to Virtuozzo Storage User Panel 3
 - 2.1.2 Adding, Deleting, and Listing Buckets 4
 - 2.1.2.1 Listing Bucket Contents in a Browser 4
 - 2.1.3 Creating, Deleting, and Listing Folders 5
 - 2.1.4 Uploading and Downloading Files 5
 - 2.1.5 Obtaining and Validating File Certificates 6
 - 2.2 Accessing S3 Storage with CyberDuck 7
 - 2.2.1 Managing Bucket Versions 9
 - 2.3 Mounting S3 Storage with Mountain Duck 10
 - 2.3.1 Creating Buckets on Mounted S3 Storage 13
 - 2.4 Configuring Backup Exec to Keep Backups in S3 Storage 13
 - 2.5 Bucket and Key Naming Policies 20

- 3. Accessing Virtuozzo Storage iSCSI Targets 22**
 - 3.1 Accessing Virtuozzo Storage iSCSI Targets from CentOS 6 22
 - 3.2 Accessing Virtuozzo Storage iSCSI Targets from Microsoft Windows Server 2012 R2 23
 - 3.3 Accessing Virtuozzo Storage iSCSI Targets from VMware ESXi 30
 - 3.4 Accessing Virtuozzo Storage iSCSI Targets from Citrix XenServer 6.2 31
 - 3.5 Accessing Virtuozzo Storage iSCSI Targets from Microsoft Hyper-V 31

- 4. Accessing Virtuozzo Storage NFS Shares 33**
 - 4.1 Mounting NFS Exports on Linux 33

4.2 Mounting NFS Exports on MacOS 34

CHAPTER 1

About Virtuozzo Storage

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

Virtuozzo Storage offers integration with the Acronis Notary service to leverage blockchain notarization and ensure the immutability of data saved in object storage clusters. The notary service time-stamps the data and stores data fingerprints in a public blockchain ledger, not owned by Acronis, ensuring data immutability. Notarization is critical in situations when the integrity of the original data must be maintained.

Blockchain notarization is commonly used to maintain authenticity of such data as property and medical records, chain-of-evidence for court documents, police video or security camera footage, intellectual property bound documents, long-term archiving that could be subject to IT audits, and 'consortium' data storage, where multiple entities or individuals need to securely store and exchange massive amounts of data and information.

1.1 Supported Storage Types

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

- S3 object storage for storing an unlimited number of files.
- iSCSI block storage for virtualization, databases, and other needs.
- NFS shares.

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

CHAPTER 2

Accessing Virtuozzo Storage S3 Buckets

To access Virtuozzo Storage S3, 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 Storage S3 allows you to access your data in several ways:

- via Virtuozzo Storage user panel;
- via a third-party S3 application like Cyberduck, Mountain Duck, Backup Exec, etc.

This sections describes how to use the aforementioned applications to access data in S3.

2.1 Managing Buckets via Virtuozzo Storage User Panel

This section describes how to manage buckets and their contents using Virtuozzo Storage user panel.

2.1. Managing Buckets via Virtuozzo Storage User Panel

2.1.1 Logging in to Virtuozzo Storage User Panel

To log in to the Virtuozzo Storage 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/`.

Log in

ENDPOINT

Use secure transfer (SSL)

ACCESS KEY ID


SECRET ACCESS KEY

LOG IN

2. On the Virtuozzo Storage login screen, enter your credentials in the corresponding fields and 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 click **Log out**.



2.1.2 Adding, Deleting, and Listing Buckets

On the **Buckets** screen:

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

The screenshot shows a dialog box titled "Add bucket" with a close button (X) on the left. Inside the dialog, there is a label "Bucket name" above a text input field containing the text "bucket1". 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, see *Bucket and Key Naming Policies* on page 20.

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

2.1.2.1 Listing 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,

2.1. Managing Buckets via Virtuozzo Storage User Panel


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**.

2.1.3 Creating, Deleting, and Listing Folders

On the bucket contents screen:

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



×

New folder

Folder name

folder1

Add Cancel

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

2.1.4 Uploading and Downloading Files

On the bucket or folder contents screen:

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

Buckets > bucket1 👤

Type	Name	Size	Last modified	
📁	folder1/			📁 New folder
📄	file1.txt	445 bytes	Jun 29 15:03	⬇️ Download file
				⬆️ Upload
				⚙️ Get Certificate
				✅ Validate
				✖️ Delete

Uploading file 1 of 1

db1.db 100 MB

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

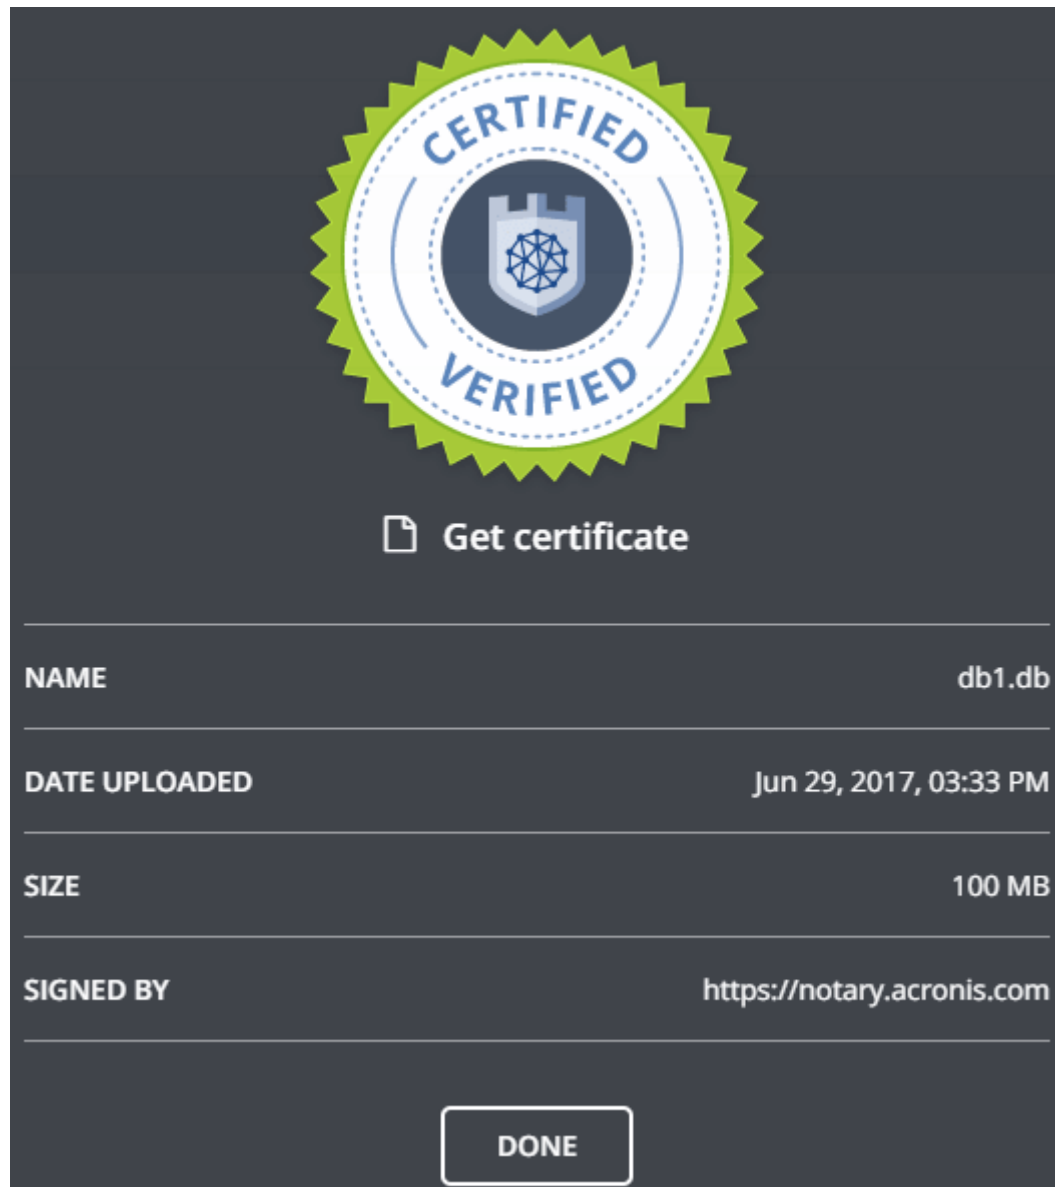
2.1.5 Obtaining and Validating File Certificates

Virtuozzo Storage 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.

On the bucket's or folder's contents screen:

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

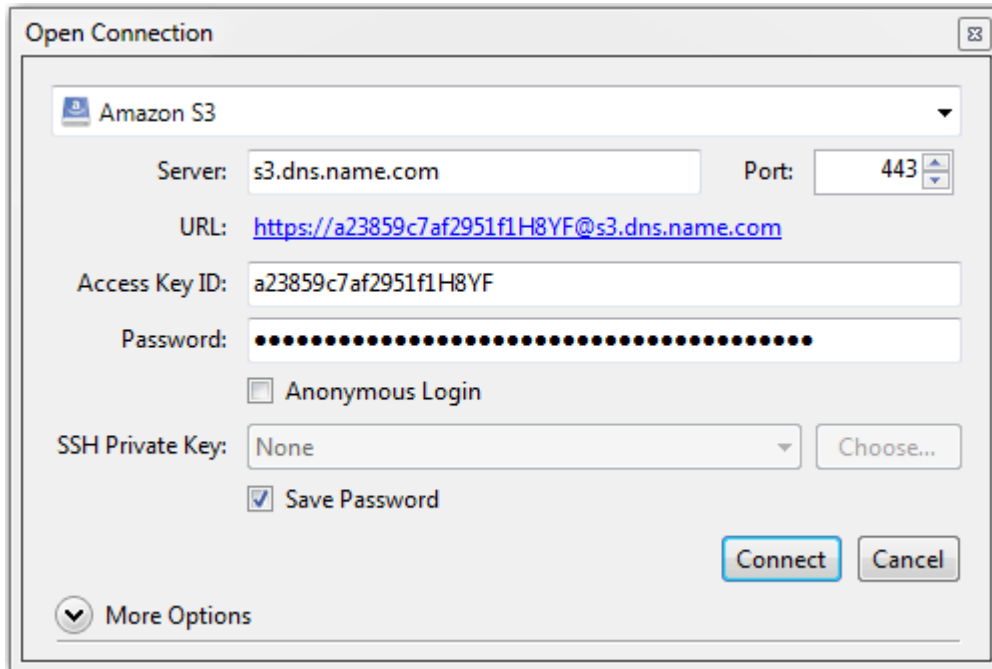


2.2 Accessing S3 Storage with CyberDuck

To access Virtuozzo Storage 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** field, 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 from <https://trac.cyberduck.io/wiki/help/en/howto/s3>.

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, see *Bucket and Key Naming Policies* on page 20.

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

2.2.1 Managing 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 <http://docs.aws.amazon.com/AmazonS3/latest/dev/Versioning.html>.

Bucket versioning is turned off by default. In CyberDuck, you can turn it on by selecting a checkbox in bucket properties. For example:




2.3 Mounting S3 Storage with Mountain Duck

Mountain Duck enables you to mount and access an S3 cluster in Virtuozzo Storage as a regular disk drive. Do the following:

1. If your service provider has provided you with an SSL certificate, install it.

2.3. Mounting S3 Storage with Mountain Duck

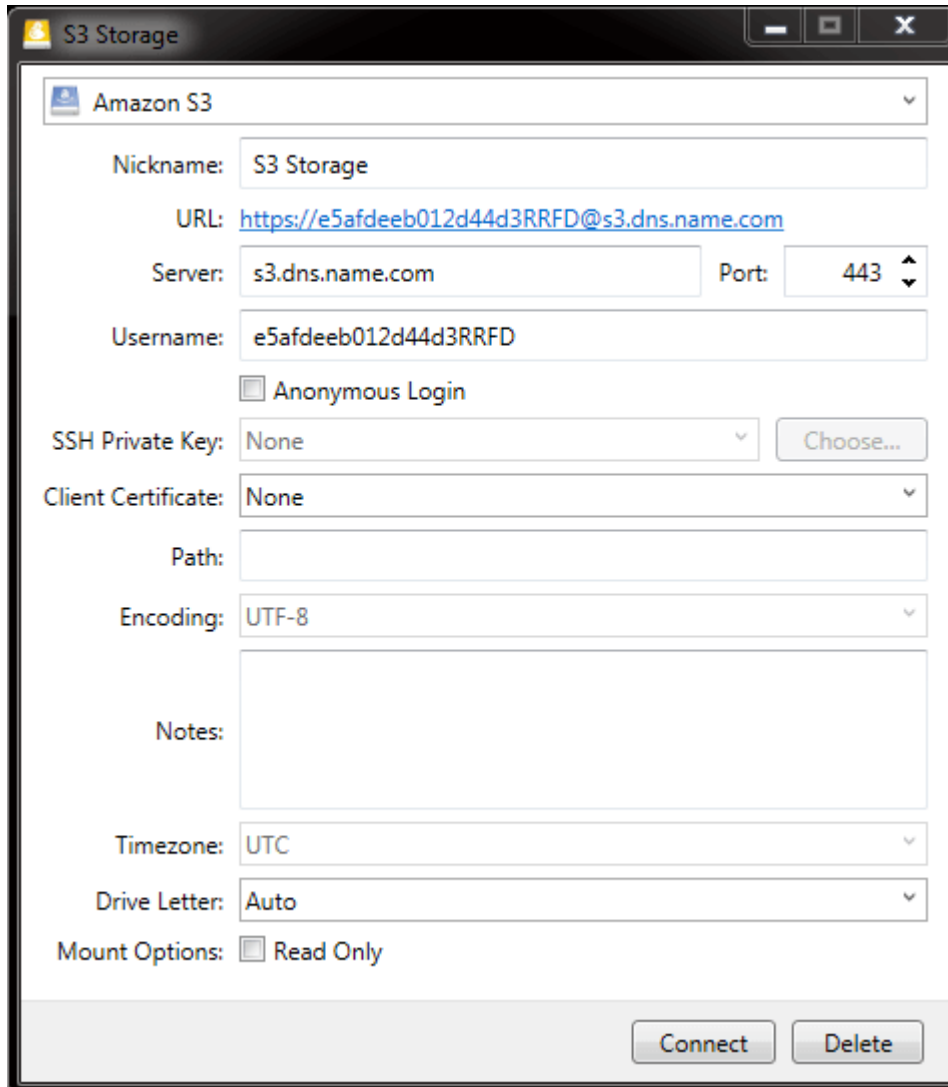
2. In Mountain Duck, click **New Bookmark**.



3. In the properties window, select **Amazon S3** profile from the first drop-down list and specify the following parameters:


- **Nickname** of the disk drive,
- endpoint DNS name in the **Server** field,
- access key ID in the **Username** field;

and click **Connect**.



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

2.4. Configuring Backup Exec to Keep Backups in S3 Storage



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

2.3.1 Creating Buckets on Mounted S3 Storage

Windows and Mac OS X, operating systems supported by MountainDuck, 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 (see *Bucket and Key Naming Policies* on page 20), 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.

2.4 Configuring Backup Exec to Keep Backups in S3 Storage

To store Backup Exec backups in S3 storage, do the following:

1. Create a bucket to store backups either using the Virtuozzo Storage user panel or another application.
2. Install Backup Exec. During installation, make sure so select all the components of Backup Exec and

check all the updates.

3. Run CLILauncher located in C:\Program Files\Veritas\Backup Exec.
4. In the Backup Exec command-line prompt, run the following command:

```
# New-BECloudInstance -Name "cloudinstance" -Provider "cloudian" \
-ServiceHost "<S3_DNS_name>" -SslMode "Disabled" -UrlStyle "Path"
```

The screenshot shows a Windows PowerShell window titled "Windows PowerShell". The text inside the window reads:

```
Welcome to the Backup Exec Management Command Line Interface
To get a list of Backup Exec commands, type:
    Get-BECommand
To launch the Backup Exec Management Command Line Interface Help, type:
    Show-BEHelp


Copyright (C) 2017 Veritas Technologies LLC. All rights reserved. Use of this product is subject to license terms.

BEMCLI> New-BECloudInstance -Name "cloudinstance" -Provider "cloudian" -ServiceHost "S3.DNS.name" -SslMode "Disabled" -UrlStyle "Path"

Name           : cloudinstance
Id              : 03353052-2567-4c5e-b928-52242763b868
Provider       : cloudian
ServiceHost    : s3.dns.name
SslMode        : Disabled
UrlStyle       : Path
HttpPort       : 80
HttpsPort      : 443
Endpoint       :


BEMCLI>
```

5. In Backup Exec, click **Configure Cloud Storage** on the **Storage** tab.



6. In the **Configure storage...** window, specify a name for Virtuozzo Storage S3 and click **Next**.

2.4. Configuring Backup Exec to Keep Backups in S3 Storage



Configure storage on WIN-1UMMOBTT4JM

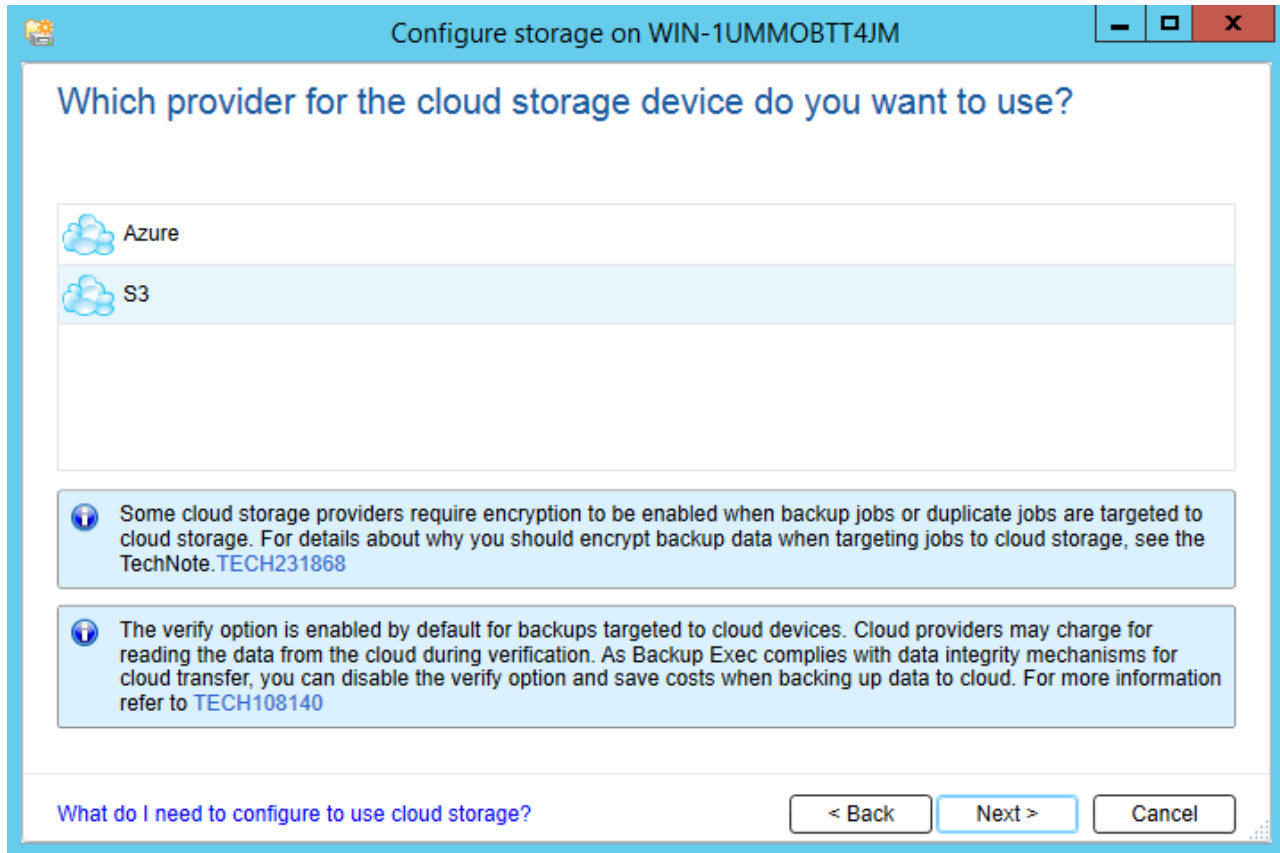
What name and description do you want to use for the cloud storage device?

Name:

Description:

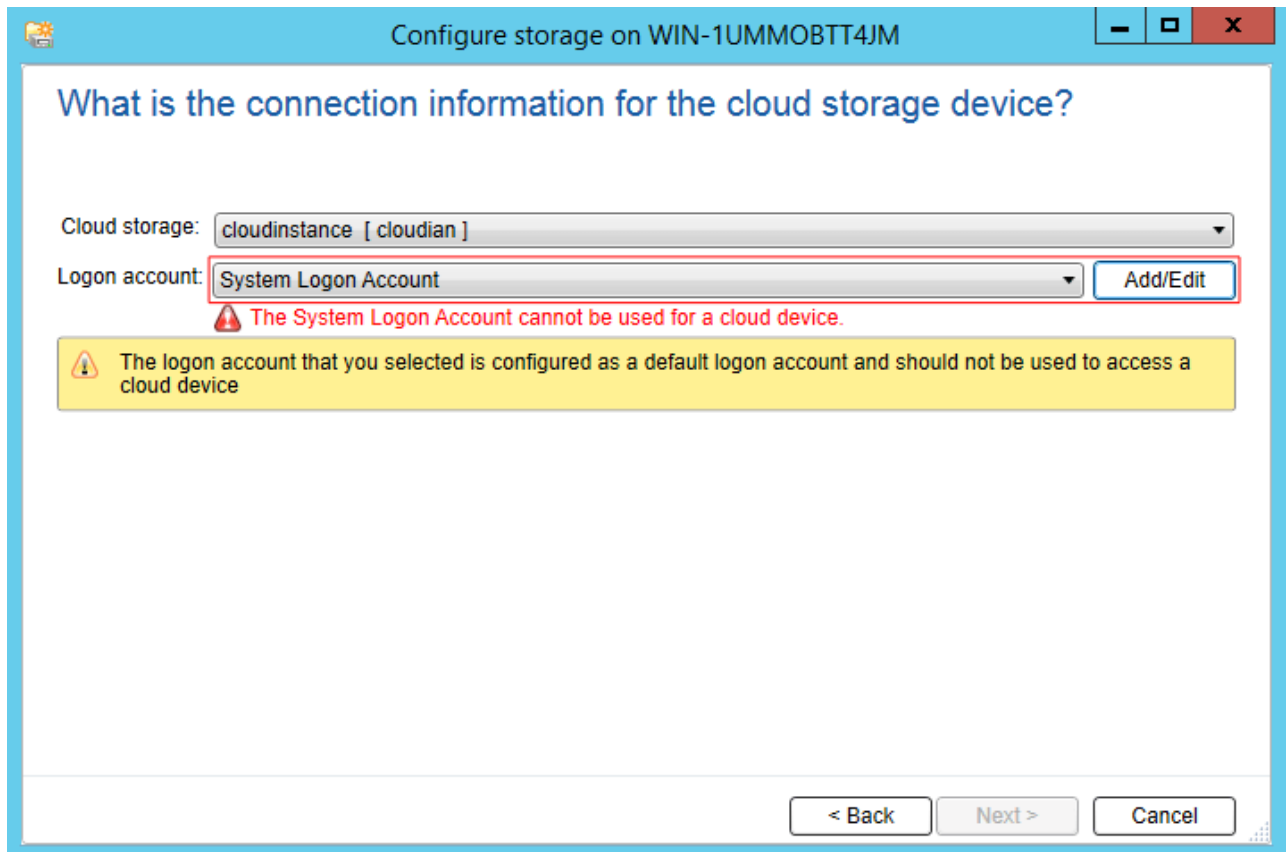
Next > Cancel

7. Select the **S3** device and click **Next**.

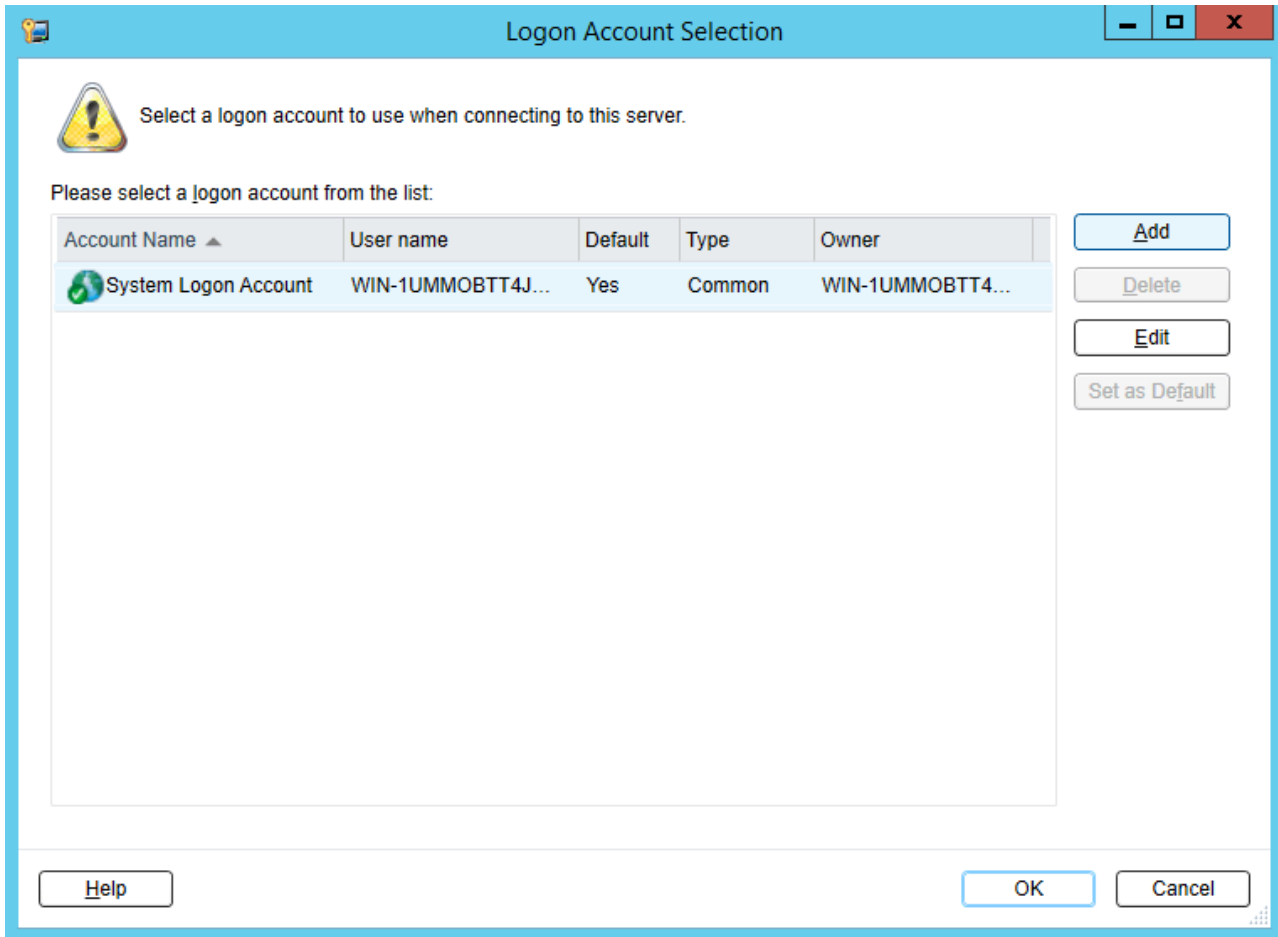


8. Select cloudinstance [cloudian] from the **Cloud storage** drop-down list.

2.4. Configuring Backup Exec to Keep Backups in S3 Storage



9. Click **Add/Edit** next to the **Logon account** drop-down list.
10. In the **Logon Account Selection** window, click **Add**.



11. In the **Account credentials** section, specify your credentials:
 - 11.1. S3 access key ID in the **User name** field.
 - 11.2. S3 secure access key in the **Password** field and confirm it.
 - 11.3. The username of your account in the **Account name** field.

2.4. Configuring Backup Exec to Keep Backups in S3 Storage

Add Logon Credentials

Account credentials

User name: e5afdeeb012d44d3RRFD

Password:

Confirm password:

Account name: s3user1@example.com

Notes:


This is a restricted logon account.

This is my default account.

Help OK Cancel

12. Uncheck all the checkboxes and click **OK**.

13. Back in the **Logon Account Selection** window, make sure the newly added user account is selected and click **OK**.



14. Back in the **Configure storage...** window, click **Next**.

15. Select a bucket and click **Next** twice.

16. On the summary screen, click **Finish**, **OK**, and **Yes**.

Once the Backup Exec services are restarted, the S3 storage will appear in the list on the **Storage** tab. Now you can create backup jobs and specify the S3 storage as destination.

2.5 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,

2.5. Bucket and Key Naming Policies

- 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.

CHAPTER 3

Accessing Virtuozzo Storage iSCSI Targets

This section describes ways to attach Virtuozzo Storage iSCSI targets to a number of operating systems and third-party virtualization solutions.

3.1 Accessing Virtuozzo Storage iSCSI Targets from CentOS 6

1. Make sure that the `iscsi-initiator-utils` package is installed.
2. Discover the required target by its IP address. For example:

```
# iscsiadm --mode discovery --type sendtargets --portal 192.168.10.100
```

3. Restart the `iscsid` service to rescan for newly added drives:

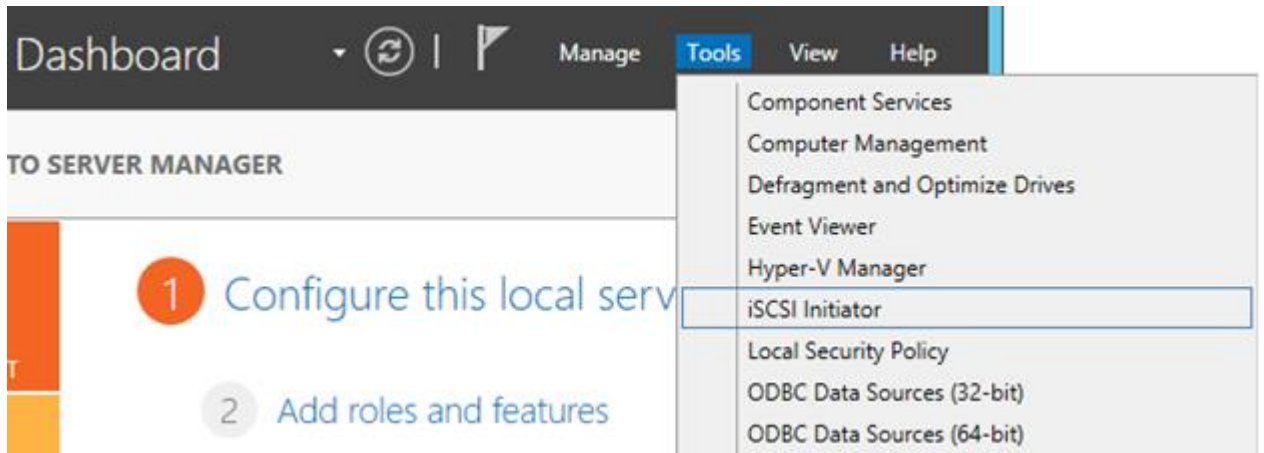
```
# service iscsi restart
```

To check that the new drive has appeared in the system, use `fdisk`, `parted` or similar tools.

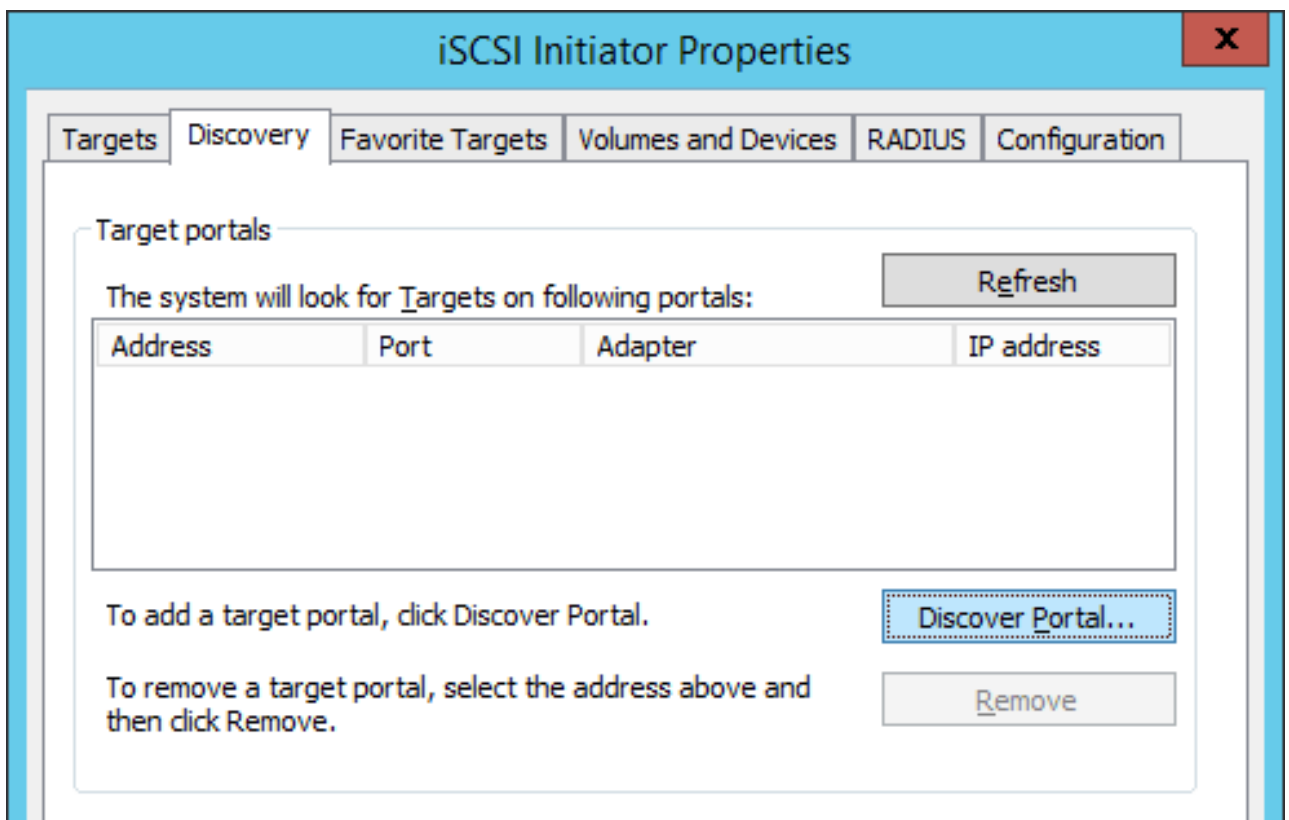
For more information, see the [Red Hat Enterprise Linux Storage Administration Guide](#).

3.2 Accessing Virtuozzo Storage iSCSI Targets from Microsoft Windows Server 2012 R2


1. In the **Server Manager Dashboard**, click the **Tools** menu in the toolbar and select **iSCSI Initiator**.



2. In the **iSCSI Initiator Properties**, switch to the **Discovery** tab and click **Discover Portal...**



3. In the **Discover Target Portal** window, enter the portal IP address and click **OK**.



Discover Target Portal

Enter the IP address or DNS name and port number of the portal you want to add.

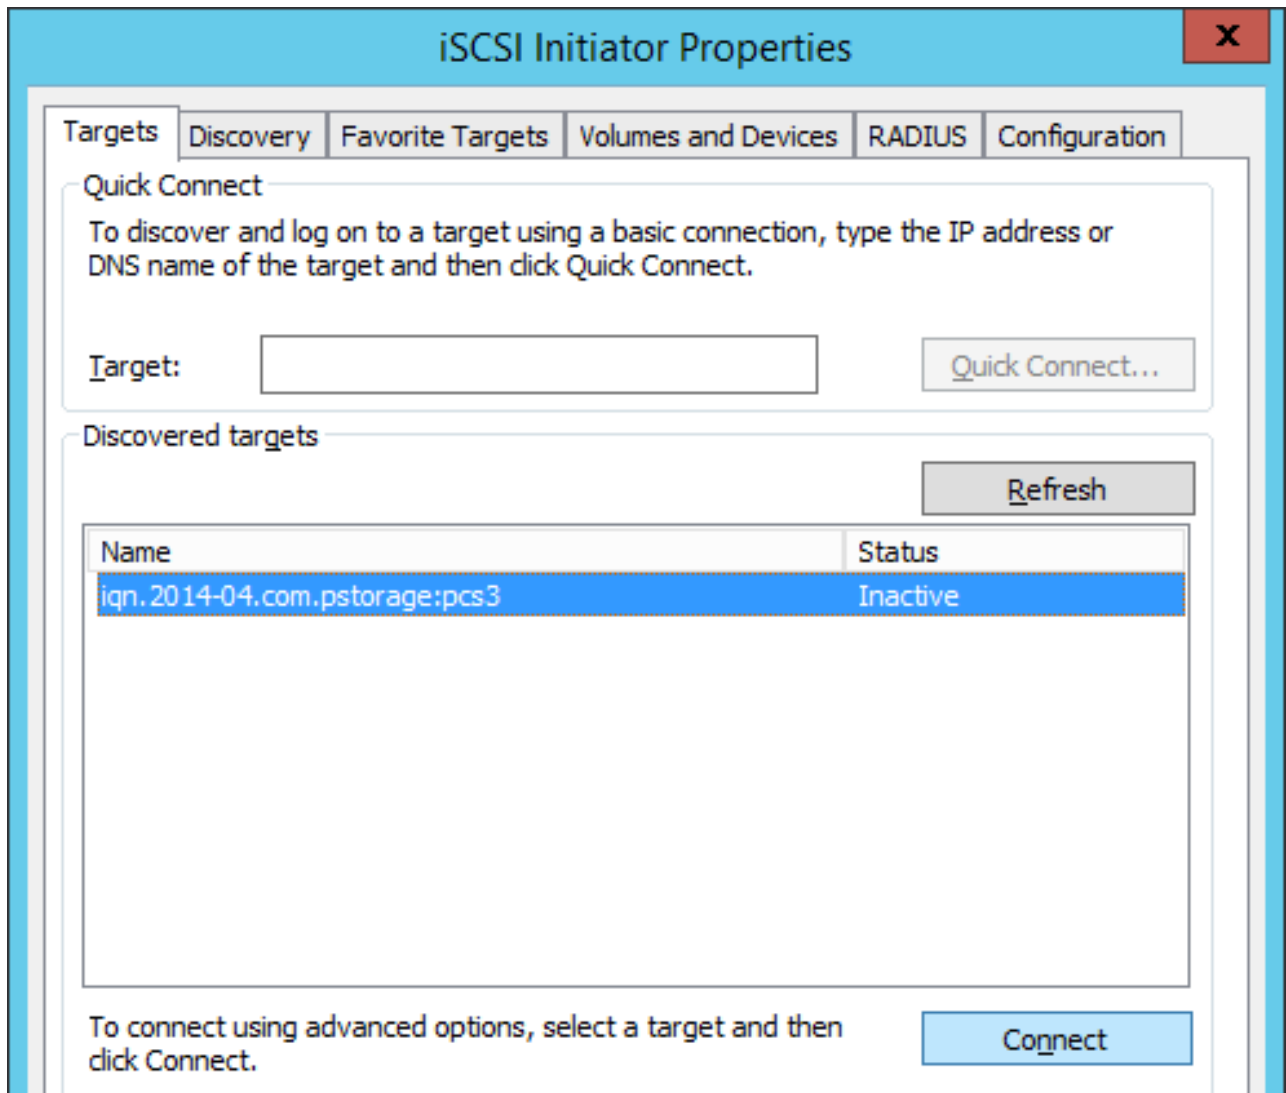
To change the default settings of the discovery of the target portal, click the Advanced button.

IP address or DNS name: Port: (Default is 3260.)

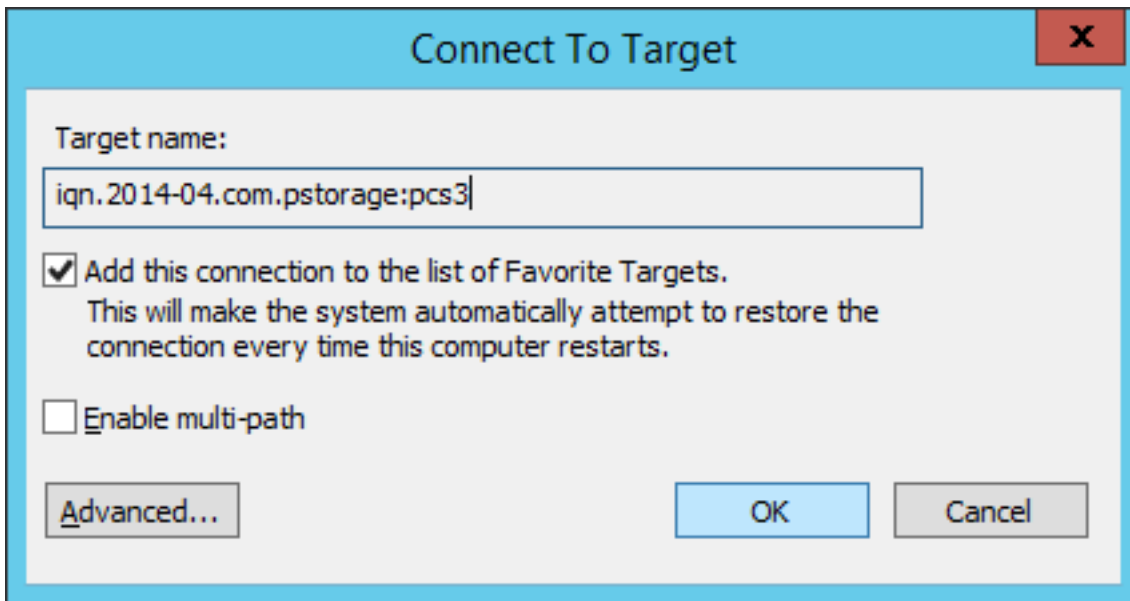
The newly added portal will appear in the **Target portals** section.

4. On the **iSCSI Initiator Properties > Targets** tab, select the new target in the **Discovered targets** section and click **Connect**.

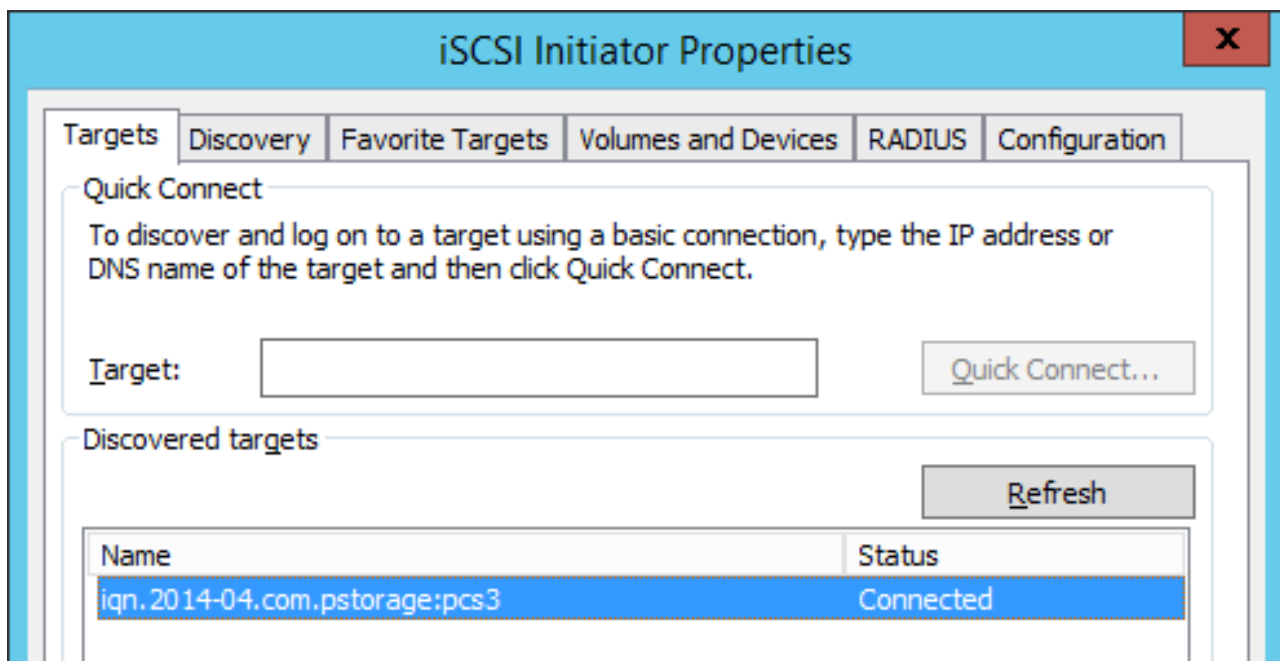
3.2. Accessing Virtuozzo Storage iSCSI Targets from Microsoft Windows Server 2012 R2



5. In the **Connect to Target** window, click **OK**.

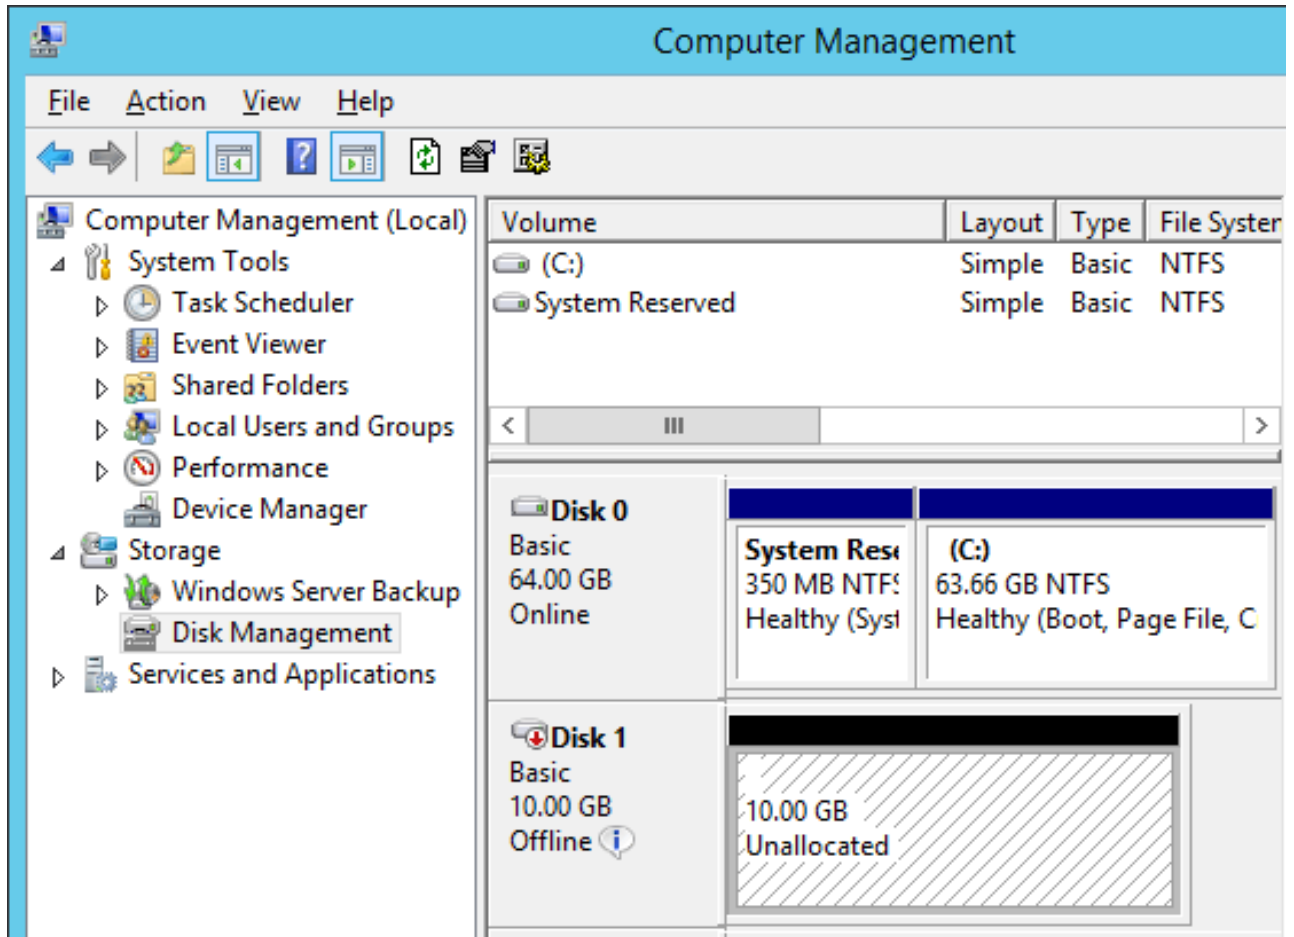


6. Target's **Inactive** status will change to **Connected**.




7. The newly attached disk will appear in **Server Manager Dashboard > Computer Management > Storage > Disk Management**.

3.2. Accessing Virtuozzo Storage iSCSI Targets from Microsoft Windows Server 2012 R2




8. Right-click the disk information section and select **Online**.

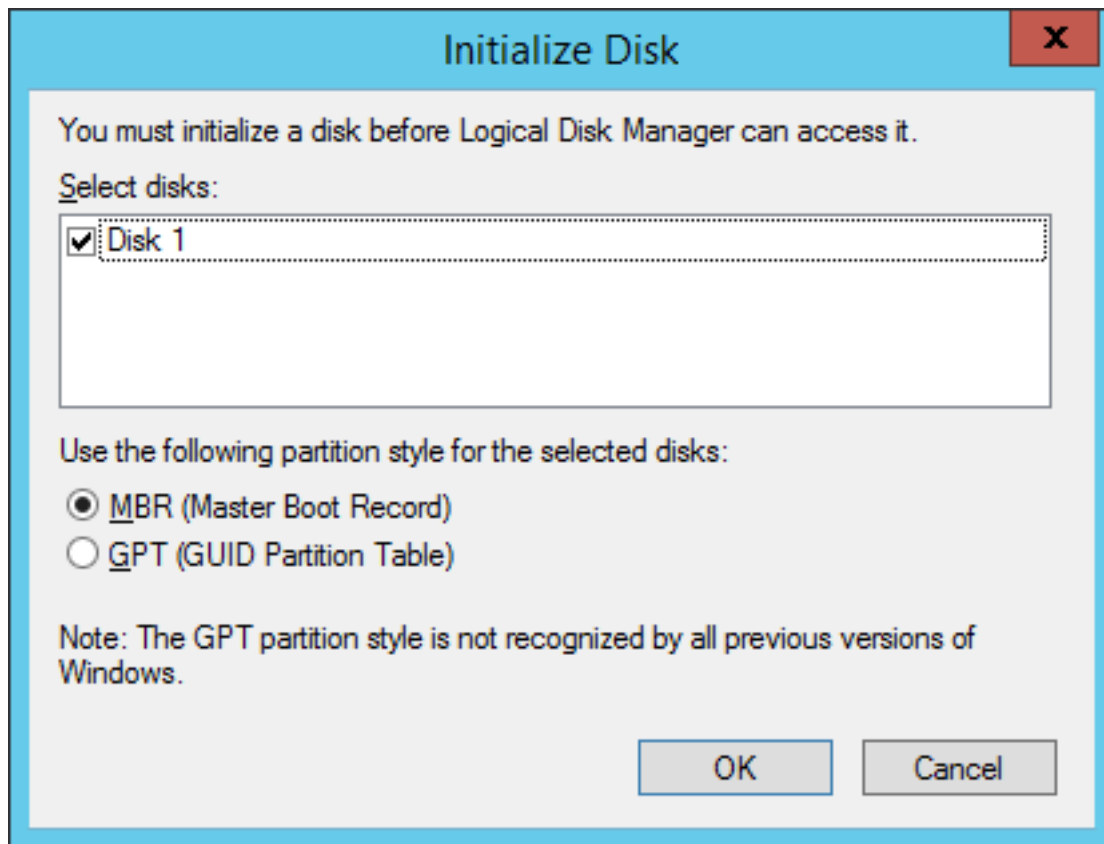


The disk status will change to **Online**.

9. Right-click the disk information section and select **Initialize Disk**.




10. In the Initialize Disk window, click **OK**.

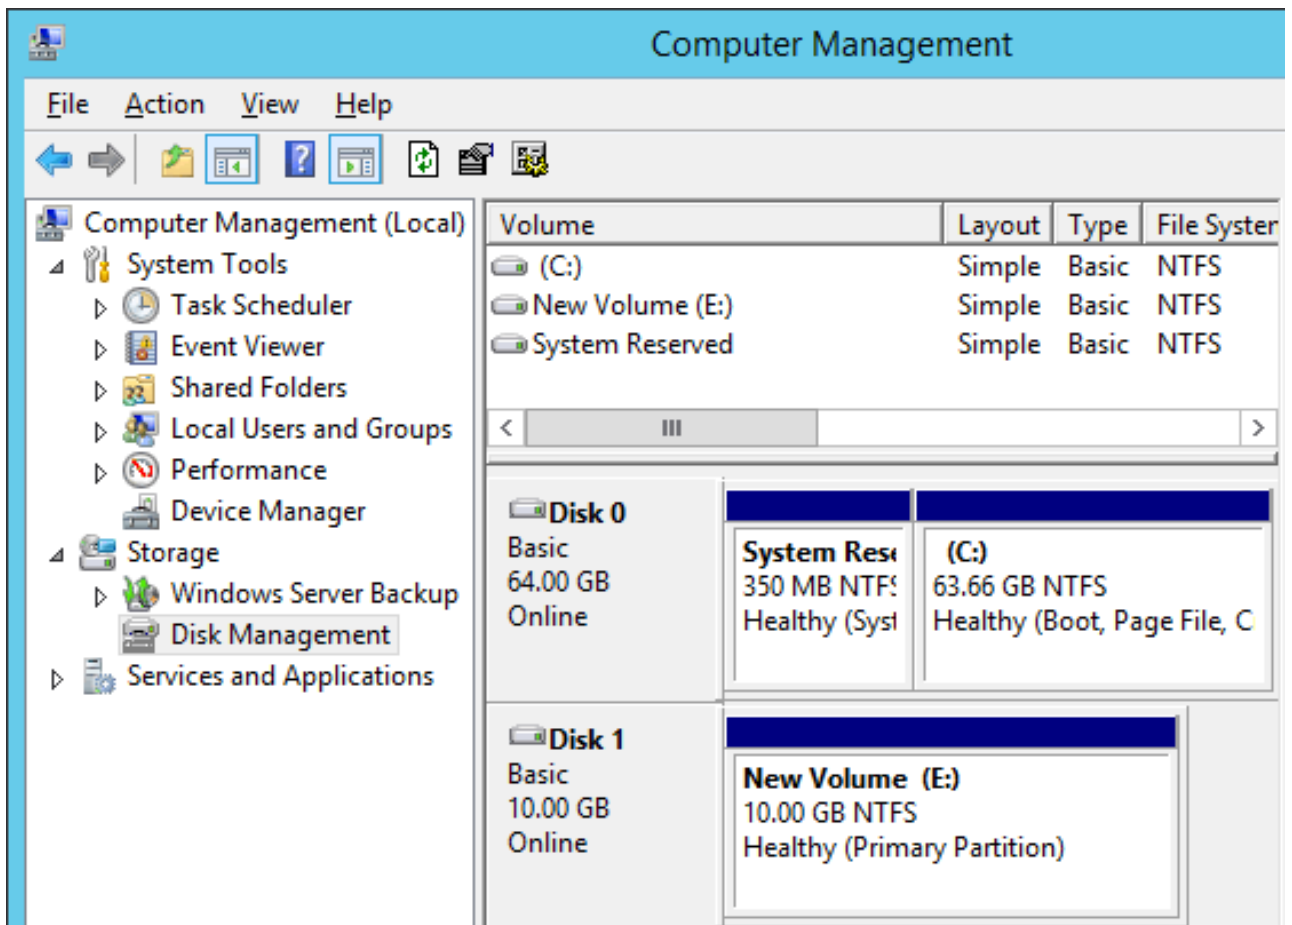


11. Right-click the disk space section, select **New Simple Volume...**, and follow the wizard's instruction to format the new disk to NTFS.

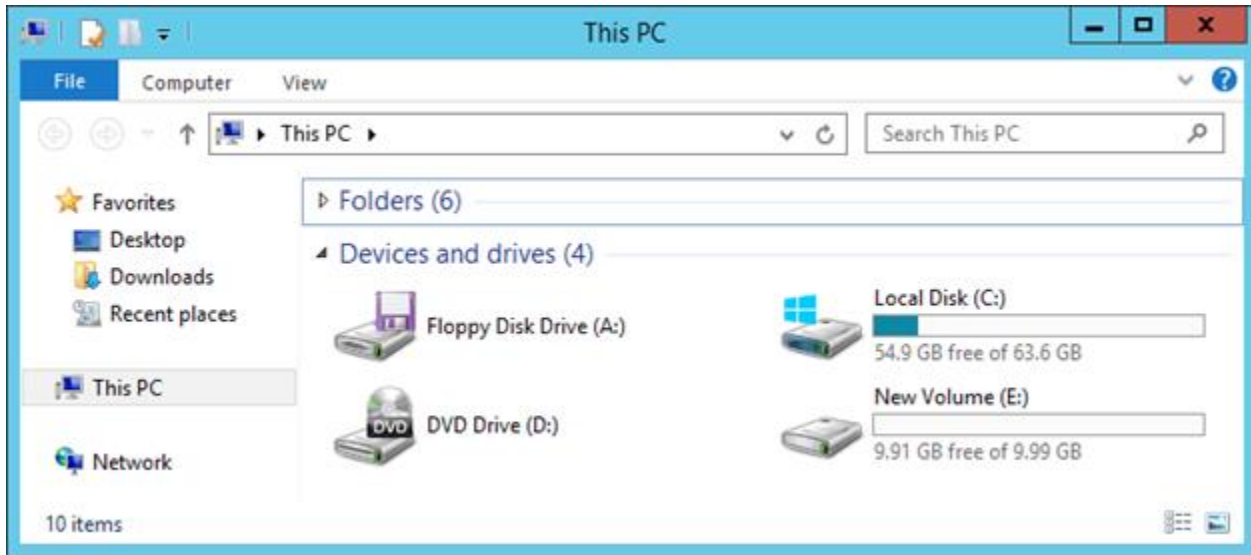
3.2. Accessing Virtuozzo Storage iSCSI Targets from Microsoft Windows Server 2012 R2



12. The disk state will change to **Healthy**.



The new disk will appear in Windows Explorer.



3.3 Accessing Virtuozzo Storage iSCSI Targets from VMware ESXi

1. In the **vSphere Client**, switch to the **Configuration** tab, and click **Storage Adapters** in the **Hardware** section.
2. If no software iSCSI adapters have been added, do so by right-clicking in the **Storage Adapters** section and selecting **Add Software iSCSI Adapter....**
3. Open the software iSCSI adapter's properties, switch to the **Static Discovery** tab and click **Add....**
4. In the **Add Static Target Server** window, enter the target's IP address and name.
5. Close the software iSCSI adapter's properties window and rescan the adapter as prompted.
6. The newly added iSCSI target will appear in the **Details** section of the software iSCSI adapter you have configured.

For more information, see the [VMware vSphere Storage Guide](#).

3.4 Accessing Virtuozzo Storage iSCSI Targets from Citrix XenServer 6.2

1. In XenCenter, switch to the **Storage** tab and click **New SR...**
2. In the **New Storage Repository** window:
 - 2.1. In the **Type** section, select the **Software iSCSI** option,
 - 2.2. In the **Name** section, provide a name or leave the default,
 - 2.3. In the **Location** section, enter target's IP address in the **Target Host** field, click **Discover IQNs** and select the desired target, then click **Discover LUNs** and select the desired LUN.
3. Click **Finish** to format the disk.

The new storage repository will appear in XenCenter.

For more information, see [XenCenter documentation](#).

3.5 Accessing Virtuozzo Storage iSCSI Targets from Microsoft Hyper-V

Note: Names of the targets to be mounted must not contain underscore characters.

1. Make sure that Microsoft iSCSI Initiator Service, MSiSCSI, is running.
2. Discover a new target portal. For example, for the portal 192.168.10.100, run:

```
PS C:\Users\Administrator>new-iscsitargetportal \  
-targetportaladdress 192.168.10.100  
Initiator Instance Name :  
Initiator Portal Address :  
IsDataDigest           : False  
IsHeaderDigest         : False  
TargetPortalAddress    : 192.168.10.100  
TargetPortalPortNumber : 3260  
.. _PSComputerName     : :
```

```
PSComputerName      :
```

3. Connect to the desired target. For example, for the target `iqn.2014-03.com.vstorage:test1`

```
PS C:\Users\Administrator> connect-iscsitarget
cmdlet Connect-IscsiTarget at command pipeline position 1
Supply values for the following parameters:
NodeAddress: iqn.2014-04.com.vstorage:test1
AuthenticationType      : NONE
InitiatorInstanceName   : ROOT\ISCSIPRT\0000_0
InitiatorNodeAddress    : iqn.1991-05.com.microsoft:win-l2dj7g36n7e...
InitiatorPortalAddress  : 0.0.0.0
InitiatorSideIdentifier : 400001370000
IsConnected             : True
IsDataDigest            : False
IsDiscovered            : True
IsHeaderDigest          : False
IsPersistent            : False
NumberOfConnections     : 1
SessionIdentifier       : fffffe00000b5e020-4000013700000005
TargetNodeAddress       : iqn.2014-04.com.vstorage:test1
TargetSideIdentifier    : 0001
.. _PSComputerName     ::
PSComputerName      :
```

4. To check that the disk has been connected, run

```
PS C:\Users\Administrator> get-disk
Number Friendly Name                OperationalStatus Total Size
-----
1      IET VIRTUAL-DISK SCSI Disk Device Offline          100 GB RAW
...
```

You can now initialise the newly mounted disk for use in Microsoft Hyper-V.

For more information, see **iSCSI Cmdlets in Windows PowerShell** at

<http://technet.microsoft.com/en-us/library/hh826099.aspx>.

CHAPTER 4

Accessing Virtuozzo Storage NFS Shares

This section describes ways to mount Virtuozzo Storage NFS shares on Linux and MacOS.

Note: Virtuozzo Storage currently does not support the Windows built-in NFS client.

4.1 Mounting NFS Exports on Linux

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

```
# mount -t nfs -o vers=4.0 -o sec=krb5 share1.example.com:/0200000000000002/export1 /mnt/nfs
```

where:

- `-o vers=4.0` is the NFS version to use. To use pNFS, change `-o vers=4.0` to `-o vers=4.1`. In all other cases, make sure to always specify NFS version 4.0 or newer.
- `-o sec=krb5` is required if your service provider uses Kerberos authentication in NFS shares.
- `share1.example.com` is the share hostname. You can also use IP addresses.
- `/0200000000000002/export1` is the path to export.
- `/mnt/nfs` is an existing local directory to mount the export to.

4.2 Mounting NFS Exports on MacOS

You can mount an NFS export created in Virtuozzo Storage like any other directory exported via NFS. You will need the share hostname (or IP address) and the volume identifier. You can mount the export either via a command-line prompt or in Finder:

- In the command-line prompt, run the following command:

```
# mount -t nfs -o vers=4.0 -o sec=krb5 share1.example.com:/020000000000002/export1 /r
```

where:

- `-o vers=4.0` is the NFS version to use.
 - `-o sec=krb5` is required if your service provider uses Kerberos authentication in NFS shares.
 - `share1.example.com` is the share hostname. You can also use IP addresses.
 - `/020000000000002/export1` is the path to export.
 - `/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://share1.example.com:/020000000000002/export1`
where:
 - `share1.example.com` is the share hostname. You can also use IP addresses.
 - `/020000000000002/export1` is the path to export.
 3. Click **Connect**.

The Finder will mount the export to `/Volumes/020000000000002/export1`.