

^zVirtuozzo

Virtuozzo Infrastructure Platform 2.5

User's Guide

January 31, 2019

Virtuozzo International GmbH

Vordergasse 59

8200 Schaffhausen

Switzerland

Tel: + 41 52 632 0411

Fax: + 41 52 672 2010

<https://virtuozzo.com>

Copyright ©2001-2018 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. Supported Storage Types	1
2. Accessing S3 Buckets	2
2.1 Managing Buckets via the Virtuozzo Infrastructure Platform User Panel	2
2.1.1 Logging in to User Panel	2
2.1.2 Adding, Deleting, and Listing S3 Buckets	3
2.1.2.1 Listing S3 Bucket Contents in a Browser	4
2.1.3 Creating, Deleting, and Listing Folders	4
2.1.4 Uploading and Downloading Files	4
2.1.5 Obtaining and Validating File Certificates	5
2.2 Accessing S3 Storage with CyberDuck	5
2.2.1 Managing S3 Bucket Versions	6
2.3 Mounting S3 Storage with Mountain Duck	7
2.3.1 Creating S3 Buckets on Mounted S3 Storage	9
2.4 Configuring Backup Exec to Keep Backups in S3 Storage	9
2.5 S3 Bucket and Key Naming Policies	13
3. Accessing iSCSI Targets	14
3.1 Accessing iSCSI Targets from VMware ESXi	14
4. Accessing NFS Shares	17
4.1 Mounting NFS Exports on Linux	17
4.2 Mounting NFS Exports on MacOS	18

CHAPTER 1

Supported Storage Types

Your service provider can configure Virtuozzo Infrastructure Platform 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 Infrastructure Platform in detail.

CHAPTER 2

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

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

2.1 Managing Buckets via the Virtuozzo Infrastructure Platform User Panel

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

2.1.1 Logging in to User Panel

To log in to the Virtuozzo Infrastructure Platform 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 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 S3 Buckets

On the **Buckets** screen:

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

✕ Add bucket

Bucket name

Add Cancel

Use bucket names that comply with DNS naming conventions. For more information on bucket naming, see *S3 Bucket and Key Naming Policies* (page 13).

- 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 S3 Bucket Contents in a Browser

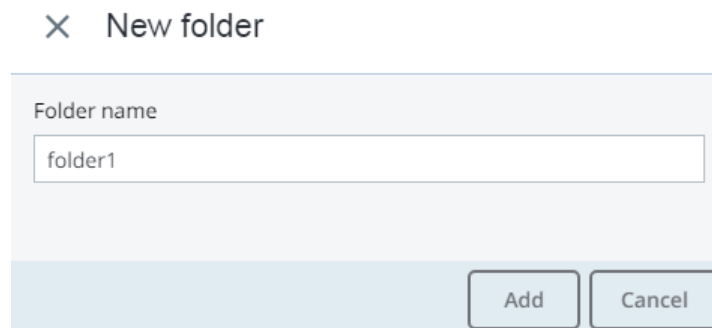
You can list bucket contents with a web browser. To do this, visit the URL that consists of the external DNS name for the S3 endpoint that you specified when creating the S3 cluster and the bucket name. For example, `mys3storage.example.com/mybucket`.

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**.




The image shows a 'New folder' dialog box. It has a title bar with a close button (X) and the text 'New folder'. Below the title bar is a text input field labeled 'Folder name' containing the text 'folder1'. At the bottom right of the dialog are two buttons: 'Add' and '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.



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

2.1.5 Obtaining and Validating File Certificates

Virtuozzo Infrastructure Platform 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.

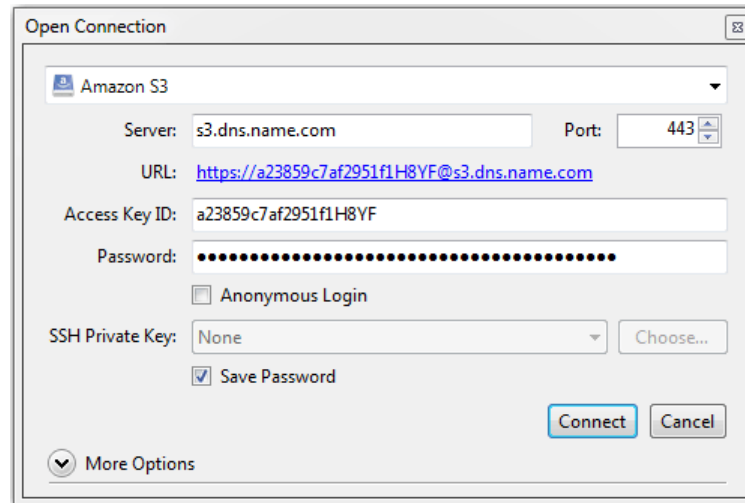
After that, you will be able to do the following:

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

2.2 Accessing S3 Storage with CyberDuck

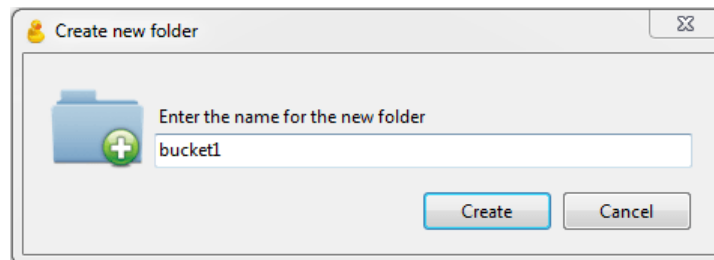
To access Virtuozzo Infrastructure Platform with CyberDuck, do the following:

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*.

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




- Specify a name for the new bucket, and then click **Create**. Use bucket names that comply with DNS naming conventions. For more information on bucket naming, see *S3 Bucket and Key Naming Policies* (page 13).

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

2.2.1 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:



2.3 Mounting S3 Storage with Mountain Duck

Mountain Duck enables you to mount and access Virtuozzo Infrastructure Platform 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:

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

and click **Connect**.

The screenshot shows the 'S3 Storage' configuration window. The fields are filled with the following values:

- Nickname: S3 Storage
- URL: <https://e5afdeeb012d44d3RRFD@s3.dns.name.com>
- Server: s3.dns.name.com
- Port: 443
- Username: e5afdeeb012d44d3RRFD
- Anonymous Login:
- SSH Private Key: None
- Client Certificate: None
- Path: (empty)
- Encoding: UTF-8
- Notes: (empty)
- Timezone: UTC
- Drive Letter: Auto
- Mount Options: Read Only

Buttons: Connect, Delete

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

The screenshot shows the 'Login s3.dns.name.com' dialog box. The fields are filled with the following values:

- Access Key ID: e5afdeeb012d44d3RRFD
- Secret Access Key: (masked with dots)
- Anonymous Login:
- SSH Private Key: None

Buttons: Save Password, Login, Cancel

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 S3 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 *S3 Bucket and Key Naming Policies* (page 13)), 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 Infrastructure Platform 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"
```

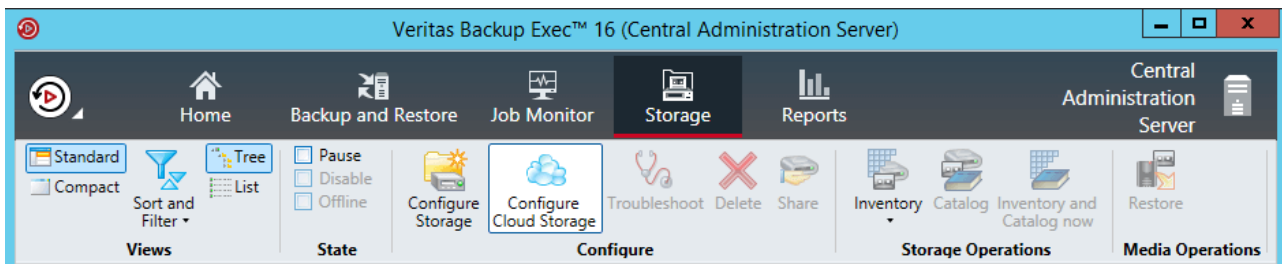
```

Windows PowerShell
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>

```

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



- In the **Configure storage...** window, specify a name for the S3 storage and click **NEXT**.

Configure storage on WIN-1UMMOBTT4JM

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


Name:

Description:


- Select the **S3** device and click **NEXT**.



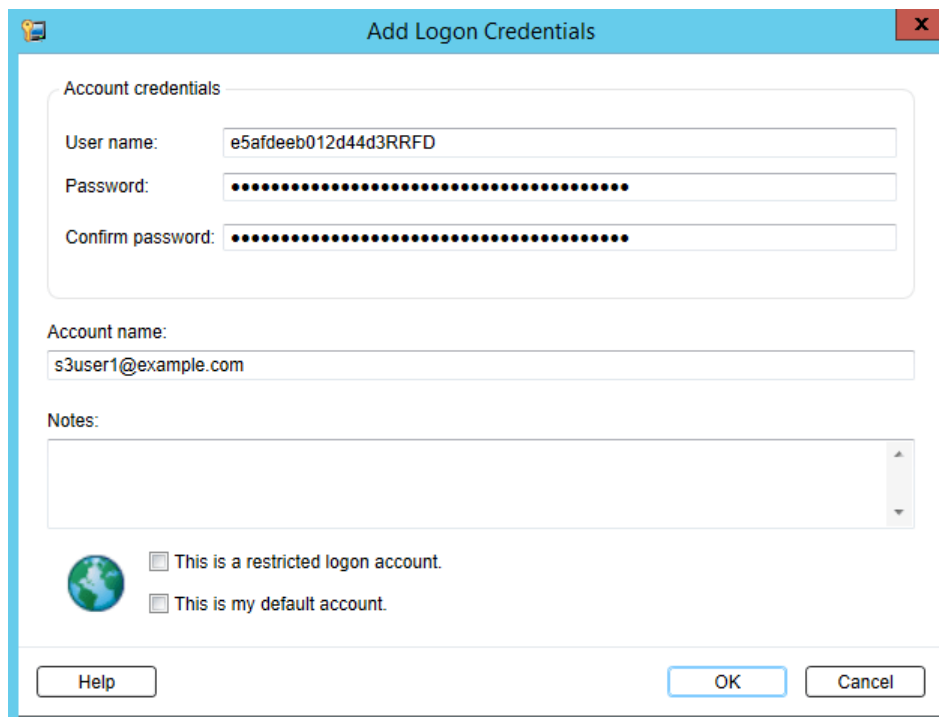
8. Select cloudinstance [cloudian] from the **Cloud storage** drop-down list.



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

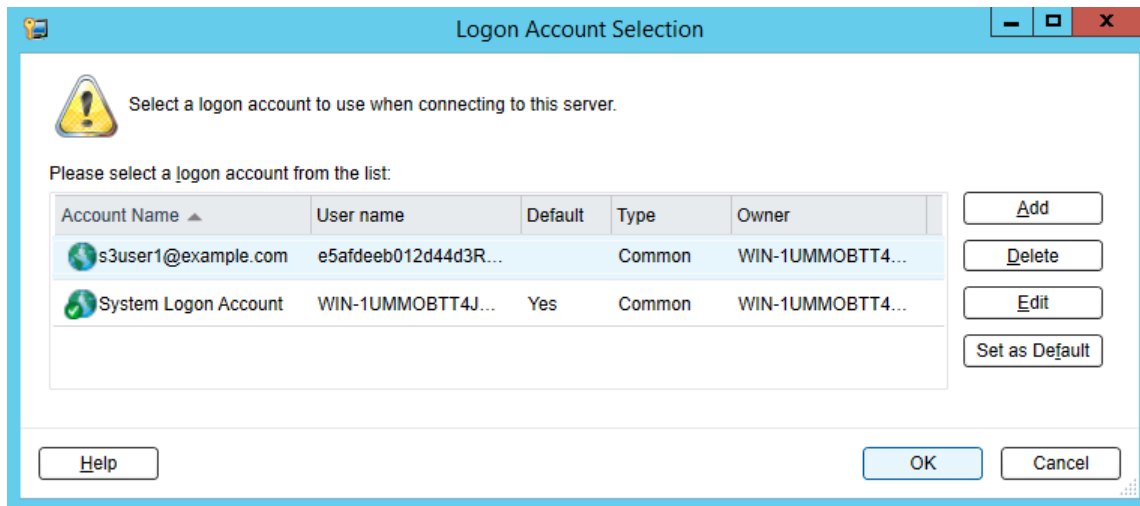


11. In the **Account credentials** section, specify your credentials:
 - 11.1. S3 access key ID in the **User name** field.
 - 11.2. S3 secure access key in the **Password** field and confirm it.
 - 11.3. The username of your account in the **Account name** field.



12. Clear 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 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.

CHAPTER 3

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.

3.1 Accessing iSCSI Targets from VMware ESXi

Before using Virtuozzo Infrastructure Platform volumes with VMware ESXi, you need to configure it to properly work with ALUA Active/Passive storage arrays. The default path selection policy (PSP) for ALUA devices is `VMW_PSP_MRU` that never falls back to standby paths. To use ALUA Active/Passive storage arrays, you need to set the default PSP to `VMW_PSP_FIXED`. For example, on VMware ESXi 6.5:

- to set the default PSP for all devices, run


```
# esxcli storage nmp satp set -b --satp=VMW_SATP_ALUA --default-bsp=VMW_PSP_FIXED
```

- to set the PSP for a specific device, run


```
# esxcli storage nmp device set --device=<device_name> --psp=VMW_PSP_FIXED
```

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

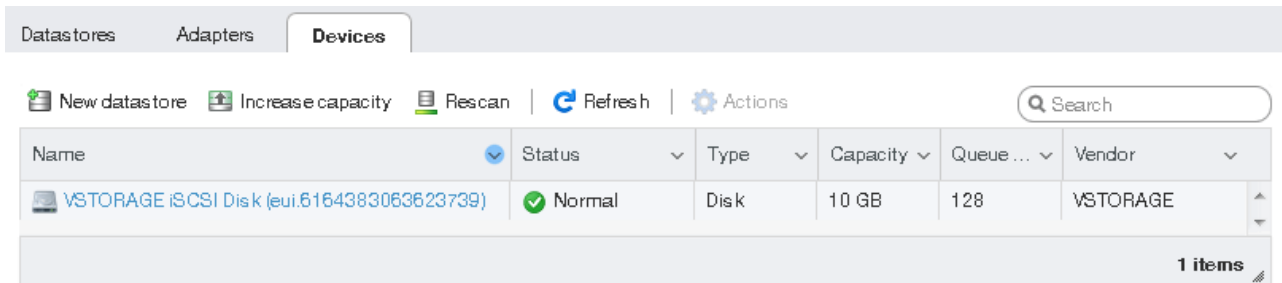
1. In the Navigator, go to the **Storage > Adapters** tab and click **Configure iSCSI**.



- 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.



- 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 datastores. You can now view its contents it with the datastore browser and provision it to VMs.

Datastores								
Adapters								
Devices								
New datastore Increase capacity Register a VM Datastore browser Refresh Actions								
Search								
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 datastores, follow the instructions in [KB article #2113956](#).

CHAPTER 4

Accessing NFS Shares

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

Note: Virtuozzo Infrastructure Platform 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 Infrastructure Platform like any other directory exported via NFS. You will need the share IP address (or hostname) and the volume identifier.

In console, run a command like the following:

```
# mount -t nfs -o vers=4.0 192.168.0.51:/<share_name>/ /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.

- `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`.
- `/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 Infrastructure Platform 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 a command like the following:

```
# mount -t nfs -o vers=4.0 192.168.0.51:/<share_name>/ /mnt/nfs
```

where:

- `-o vers=4.0` is the NFS version to use.
 - `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`.
 - `/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>/`.