# Scale-Out Functionality User Guide

## Important Notes:

1. **The Client mode is the default mode.**
2. **Once the system has enabled the server role of the Scale-Out function, all other services will be stopped.**
3. **It is recommend to use only identical HDD models in one NAS**
4. **If Auto Management is enabled, please ensure that the standby brick volume's size is equal or greater than the size of the damaged volume**
5. **It is recommended to use at least two Scale-Out systems to avoid a single point of failure.**
6. **Resetting a Scale-Out server will erase all data.**
7. **"Brick" means storage volume in the context of this manual, e.g. a formatted HDD or a RAID 1, etc.**
8. **Even though, theoretically a scale out compound can be realized over the internet with the use of VPNs, we recommend the usage in a local network environment with at least one Gigabit of bandwidth.**

# Table of Contents

# General Explanation of Scale Out

The Scale-Out function allows volume capacity being expanded dynamically through connecting numerous independent Thecus NAS systems in the same network subnet. The most remarkable advantage concerning the Thecus Scale-Out function is that the expansion impact is minimal. Added capacity will be seamlessly integrated into the network storage, the data will be automatically redistributed to the new storage capacity and even single disks in existing NAS can be added to the scale out compound easily.

Also, since the data is distributed across devices, RAID system's redundancy becomes obsolete in many applications as the Scale-Out system offers redundancy (with the number of copies being selectable) across devices with the added advantage that data would be still accessible, even if a whole NAS unit becomes unavailable.

Scale-Out relies on a client-server architecture and is usually requires at least two Thecus NAS systems. One acts as a client and the other one as server.

To understand better how Scale-Out works and learn how to easily set it up, let's systematically go through the steps below.

The Scale-Out function can be found under the "Storage" category in the "Setting" Panel.

# First Steps and Scale-Out Client Role

In this example, the Thecus N2810, N5810PRO and N12910SAS will be used for the setting steps and actual usage. The N5810PRO will serve as the Scale-Out client and the other 2 units will act as Scale-Out servers.

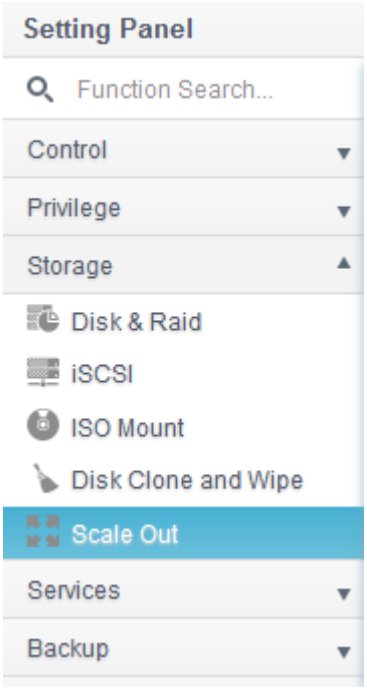


To become the Scale-Out server, the user must have created a volume, which can be a RAID over several disks or a single disk classified as a JBOD (these creation functions are listed together in the RAID menu of the OS). The (RAID) volume can be created by following the standard creation procedure (please refer to the user manual) or by enabling "Auto Management", located in the advanced settings. For this sample unit we had a volume created before. Below you can find a screenshot of the RAID creation menu showing the volume status, which is a JBOD over two disks in our example case.

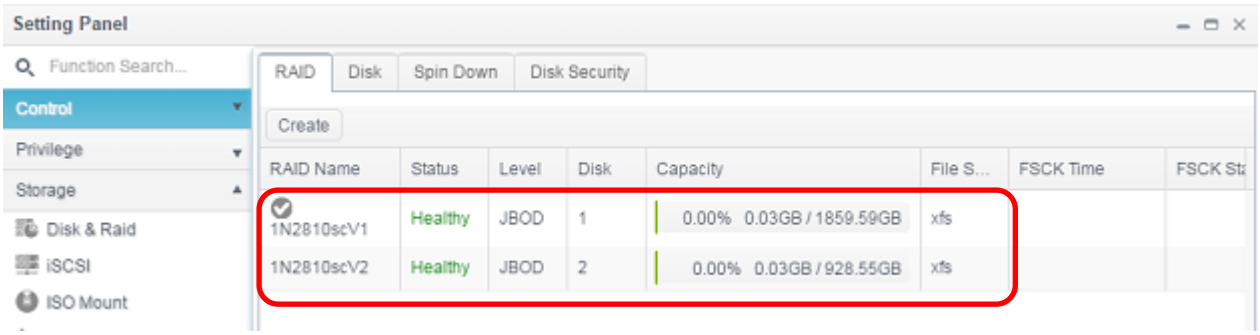


The system IP address is needed to let the Scale-Out client connect, so we have listed the system\s IP address. The example system IP for this unit is 172.16.65.143.

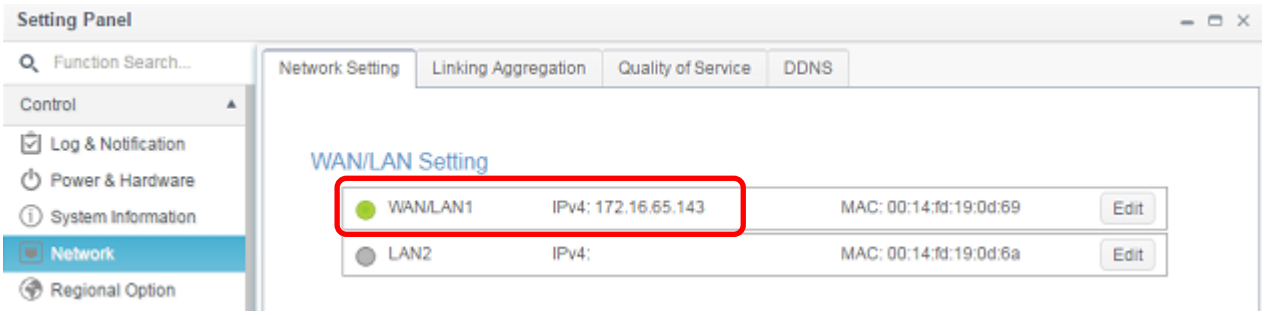


Now, we can continue to set the Scale-Out functional role for this candidate. The Scale-Out functions can be found under the “Storage” category. Click on Scale-Out and the settings screen will appear as below. The default role of Scale-Out is client mode.

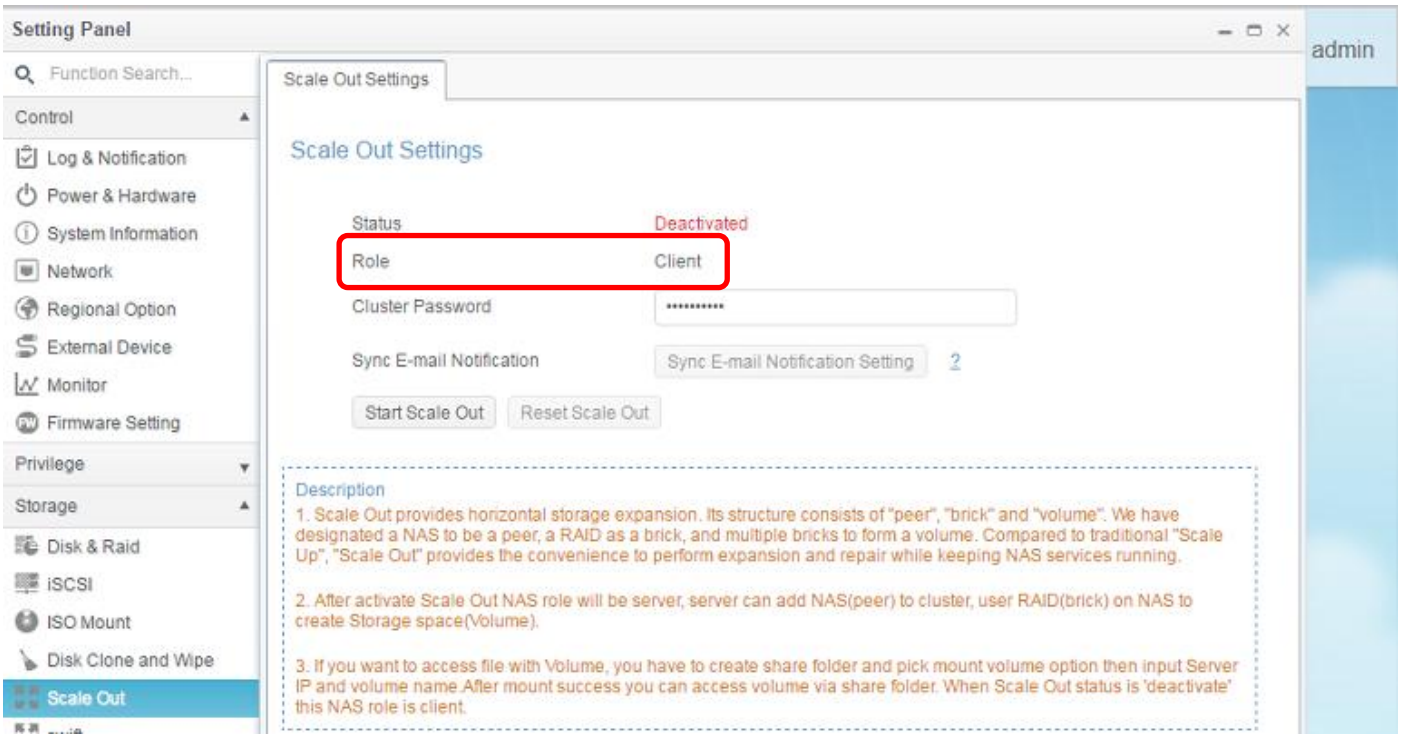

# Enabling Scale-Out Server Role

The default mode of Scale-Out is client mode. To select a server role, please input a password in “Cluster Password”. Then click on the “Start Scale-Out” button. This cluster password is going to be used for the scale our server members to know each other and distinguish it from other Scale-Out groups.

Once the Scale-Out server role has been created, the system will log out automatically. Please login again and since the Scale-Out server has been enabled, you will notice that many functions have been disabled, such as iSCSI, samba, afp ftp, etc. Now, when go to the Scale-Out settings page, you will notice that the status of Scale-Out will have been set to the “Server” role.

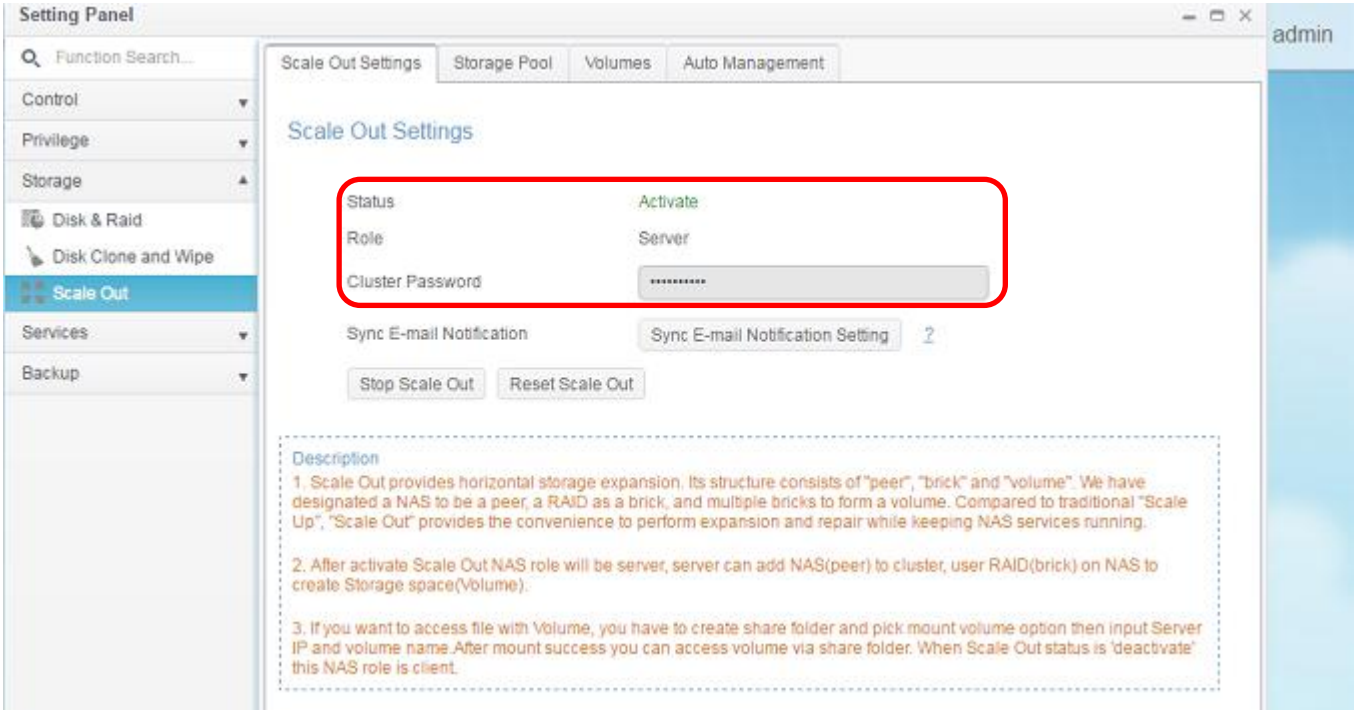


Other than then Scale-Out role being set to “Server”, there are a few more tabs available for further settings.

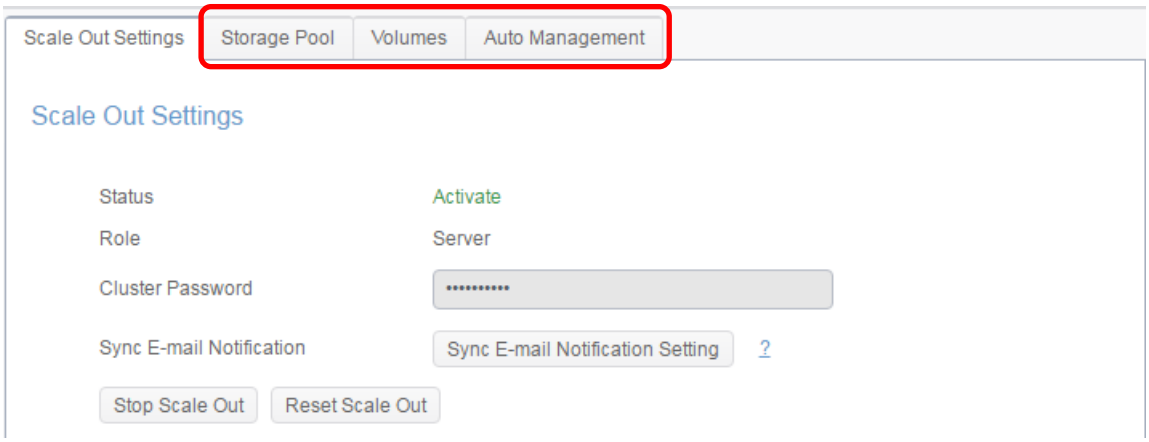

# Storage Pool

“Storage Pool” is the tab that lists the available “Peers” for Scale-Out server members. The “Peers” can be seen as a single systems. In the “Peer” information section, we can see that there are two “Bricks” available; the “Brick” can be seen as a “Volume” count of the associated system; this sample unit has 2 volumes created. Look at the screenshot below for details.

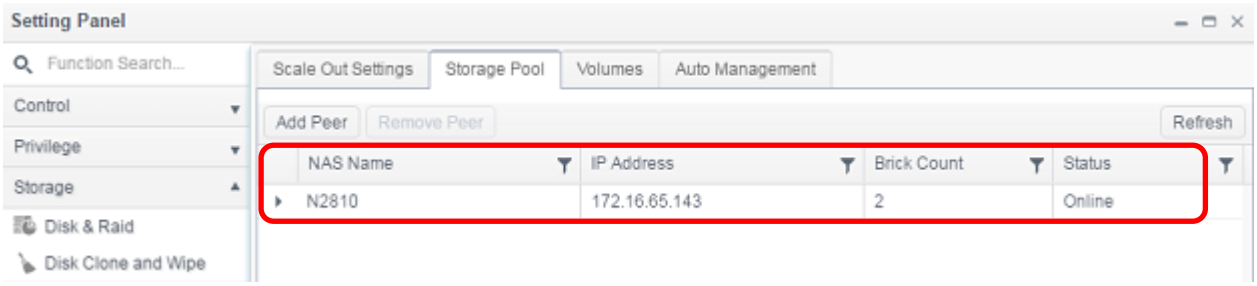


To see more “Bricks” from other systems, click on “Add Peer”, and the system will browse the local network and list all available “Peers” to choose from. See below for a sample list:

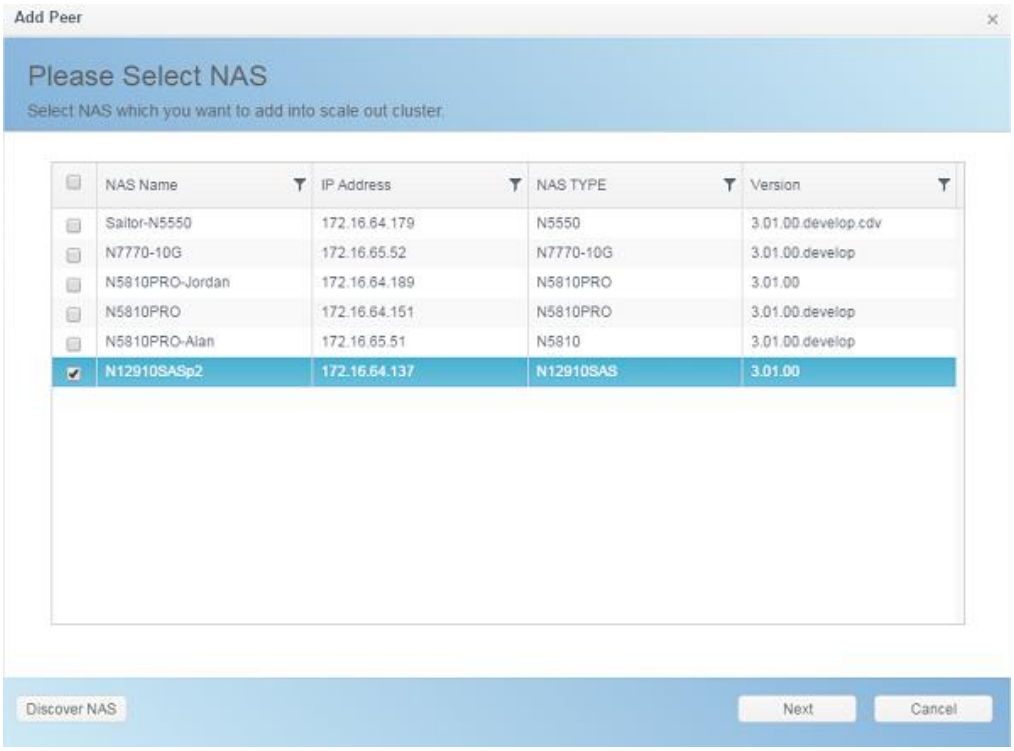


Let’s choose N12910SASp2, then click “Next” and the system will require the admin’s password to gain the permission to add.

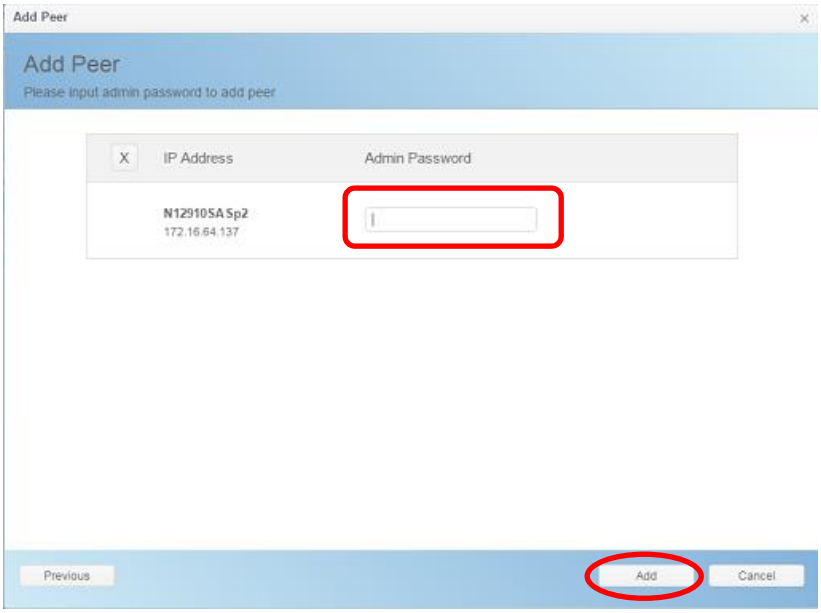


Input the admin's password of the associated system and then click the "Add" button.

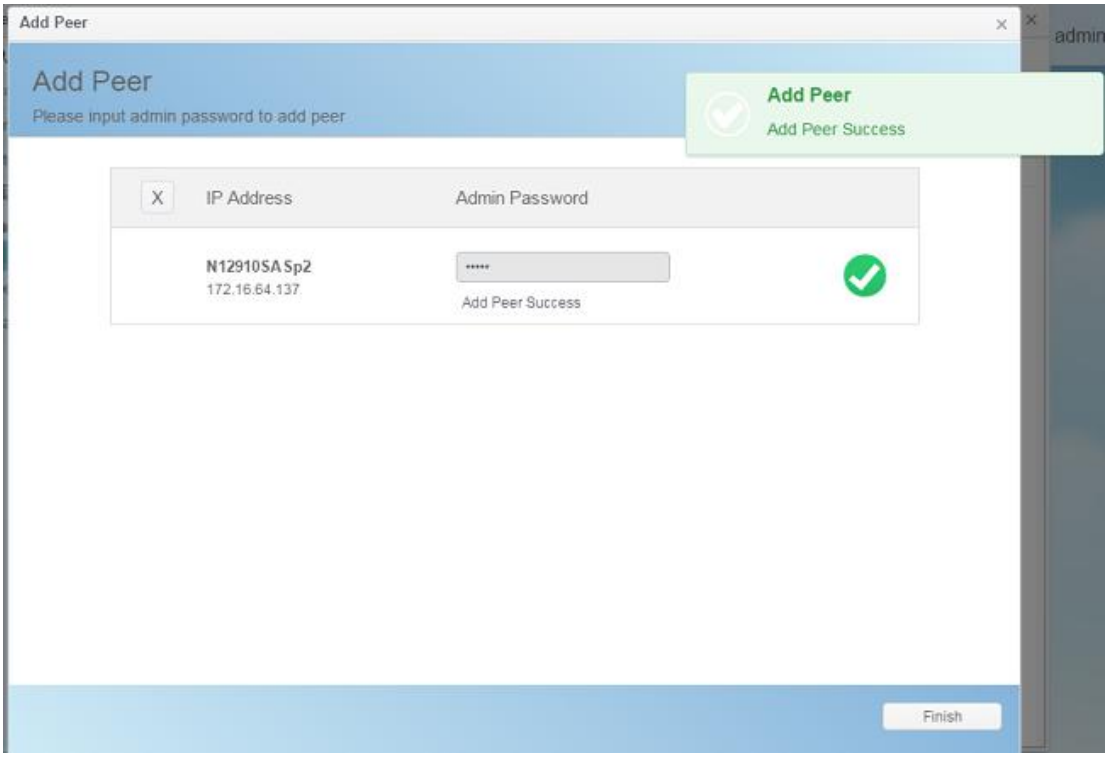


The system will start to communicate with the selected candidate and display a pop-up message once the operation has been completed. Since the system will need to enable the Scale Out role at candidate side, it may take a little while; therefore please be patient while the task is being completed.

Now the storage pool will have a new Peer that will join four additional bricks to the storage pool. Please see below.

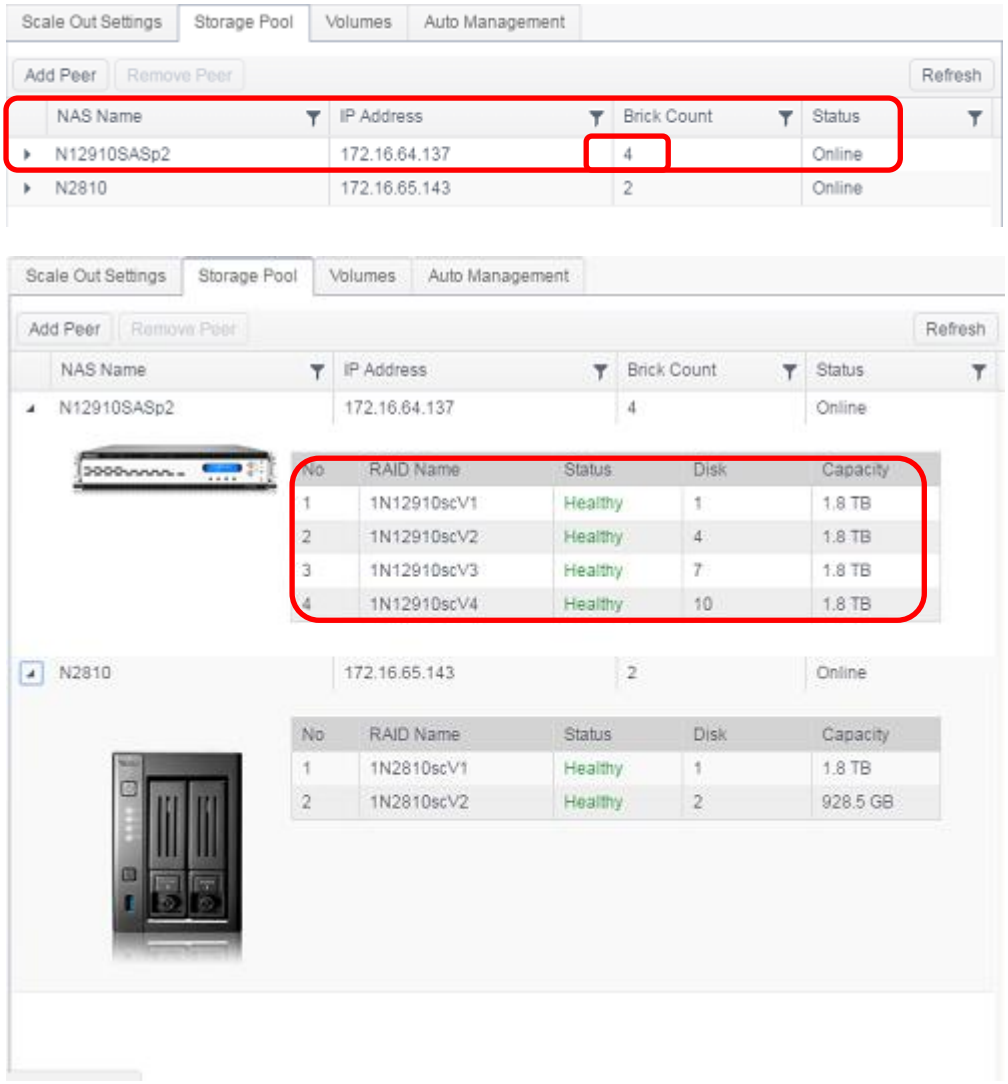


Let's check where these bricks have come from. Login to the just added NAS (http://172.16.64.137) (only in this example, your IP will differ) and check the RAID volumes and you can find that the volumes have joined the list of our Scale-Out bricks.


For the same Scale-Out server group, members will sync their settings periodically. As you can see, the first Scale-Out sample unit N2810@172.16.65.143 and the joined peer N12910SASp2@172.16.64.137 have the same 'Storage Pool' lists.

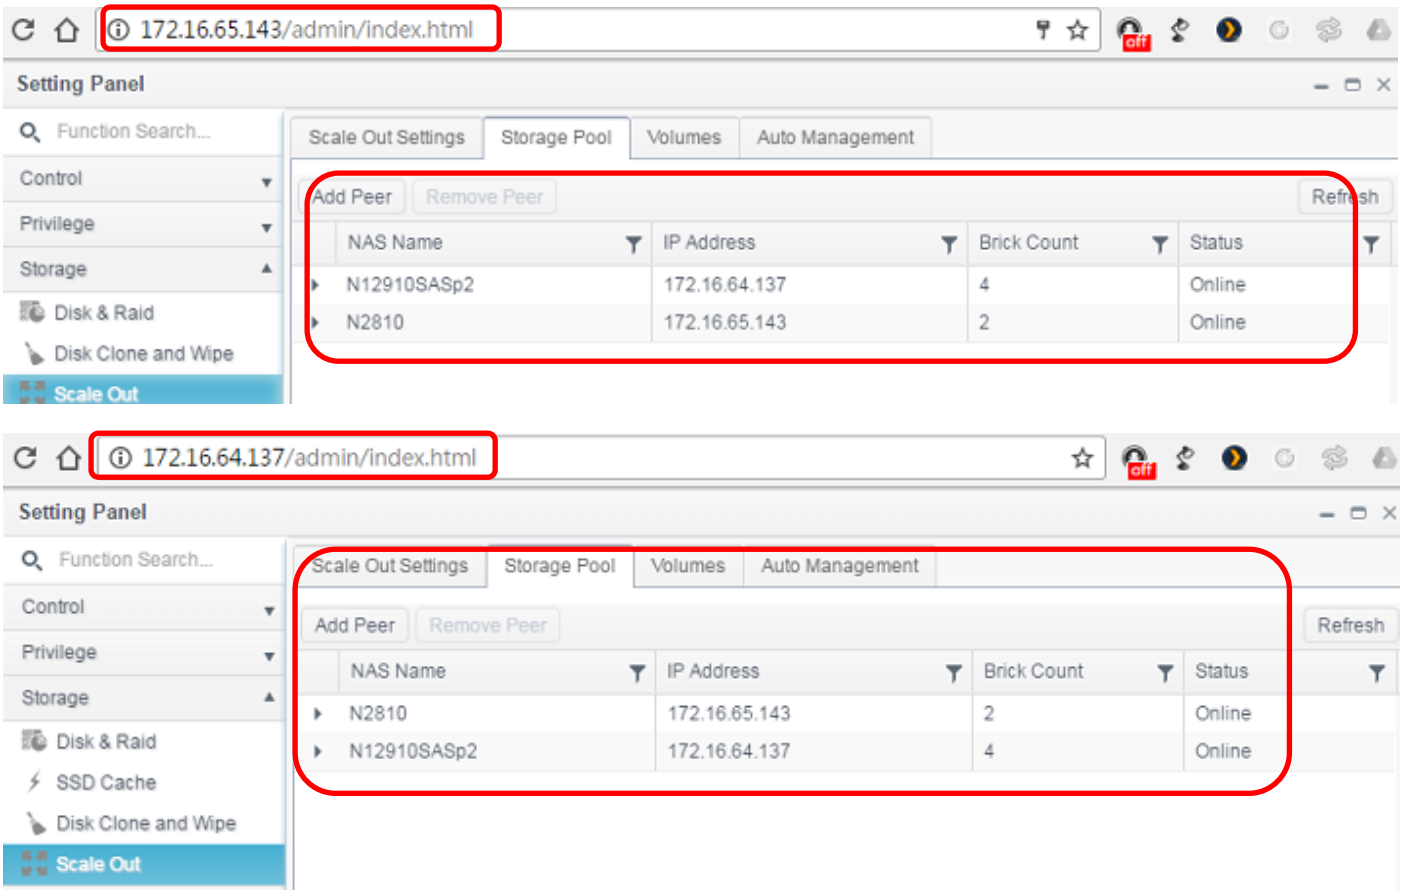


**NOTE**

The System will fail to add a Peer if:
1. The selected Peer is already being used in other scale out server group.
2. The input admin password is incorrect.
3. The selected system does not support scale out function.

To remove a “Peer” from Scale-Out server group, select the associated “Peer” then click the “Remove Peer” button, then confirm.

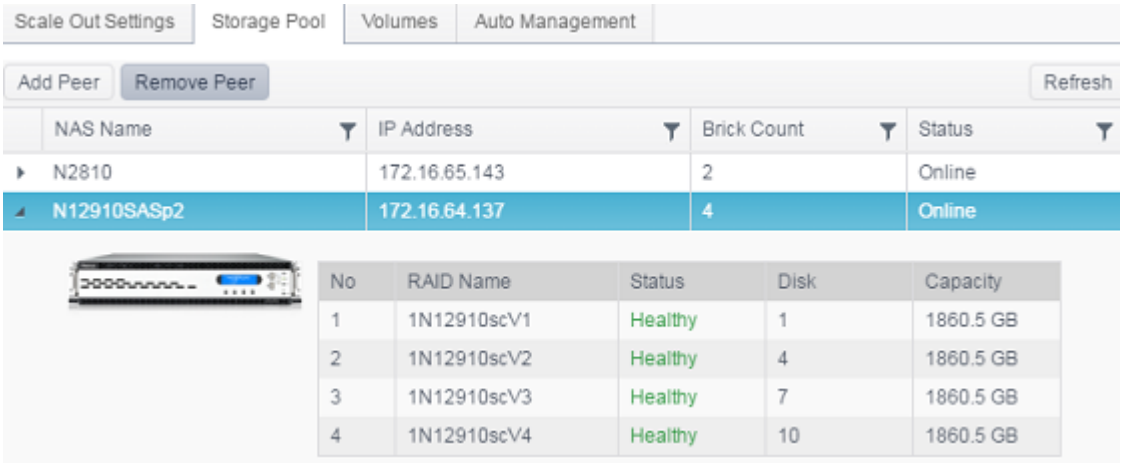


**NOTE**

The Peer cannot be removed if it is at the localhost level.

# Volumes

After the “Peer” and “Brick” setup has been completed, now we can create a “Scale-Out Volume” for Scale-Out to connect with.

**Add Scale-Out Volume:**

The Scale-Out volume can be created in ANY Scale-Out server as long as it is in the same group. Let’s take our Scale-Out server with the IP http://172.16.65.143 (example) to create a first Scale-Out volume.


Click on “Add” and the “General Settings” screen will appear as below:

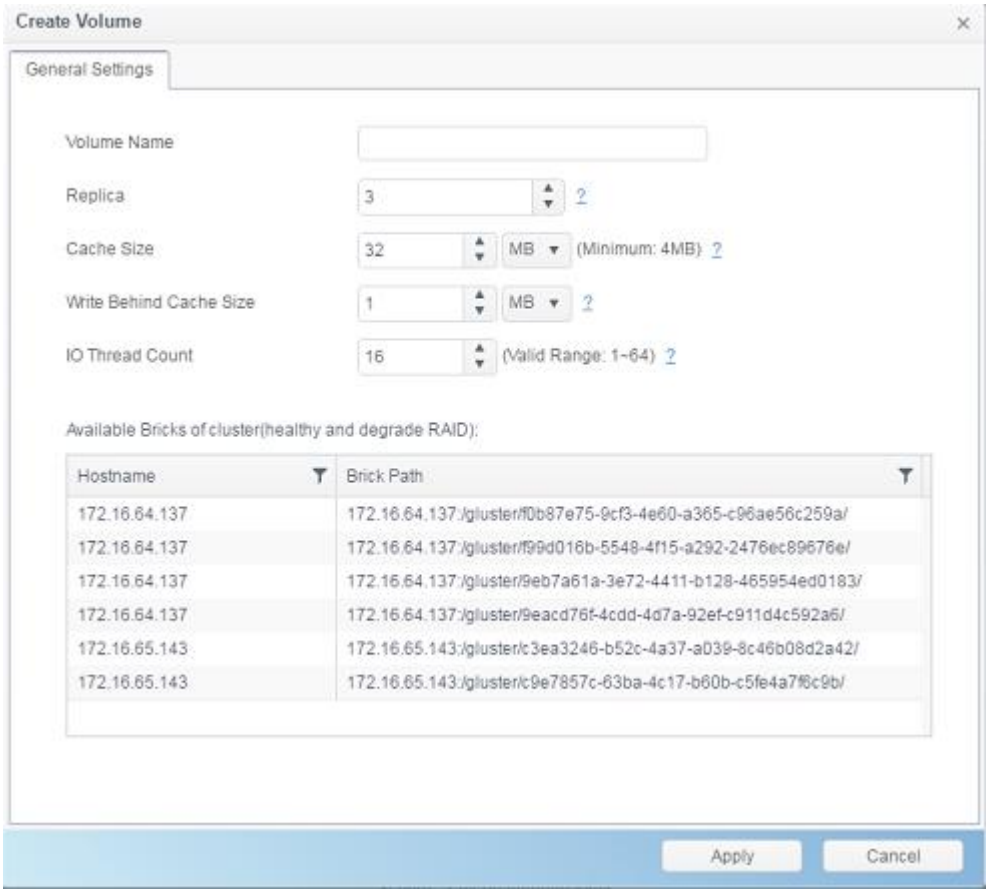

**Steps:**

1. Input Volume Name: It is going to use this name to create the shared folder.
Let's take "firstSCvolume" input as an example.


2. Replica: This is the setting for how many data copies that are going to be created per volume group. The default value is 3. So from this example, the total bricks count is 6 (N2810 x2 and N12910SASp2 x4). If using the default value 3, then this volume will have 2 groups and each group having 3 data copies.


**NOTE**

The example above features 6 available bricks but you could choose 2 or 3 or 4 or other replica values.
1. The available brick must be greater or equal than the replica or the volume won't be created.
2. Setup three replicas and the volume will have two groups. Data I/O will be read/written to this volume of two groups at the same time with three data copies.
3. If the replica value 2 is chosen, then the volume will have three groups with two respective data copies but I/O are done to the volumes in three groups simultaneously.

3. Cache size: Size of the read cache, default value is 32MB.

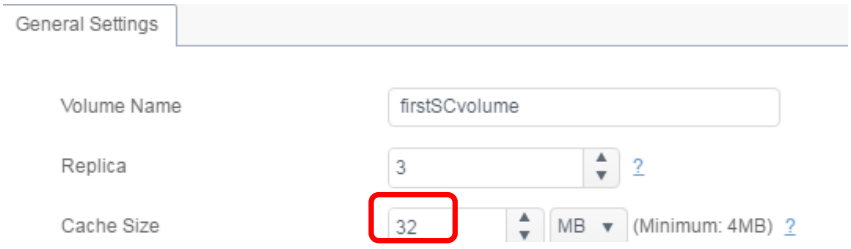

4. Write Behind Cache Size: Size of the write-behind bufferDefault value is 1MB.


5. IO (Input/Output) Thread Count: Number of threads in IO threads translator which are concurrently performed at a given time, default value is 16.

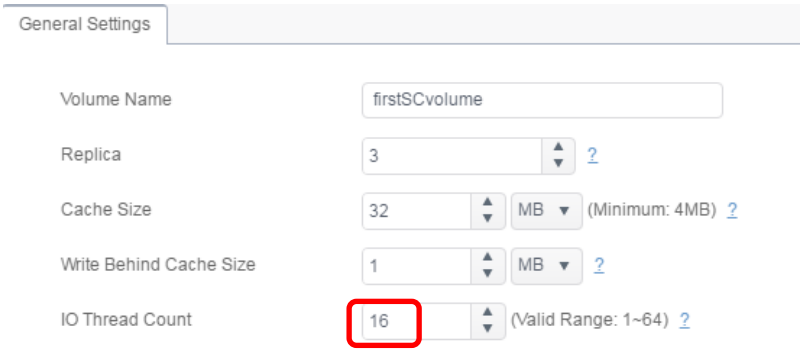


There are available bricks listed for the associated Scale-Out server group.


Click Apply to confirm the settings, then you will have your first Scale-Out volume created, please look at the screenshot below.

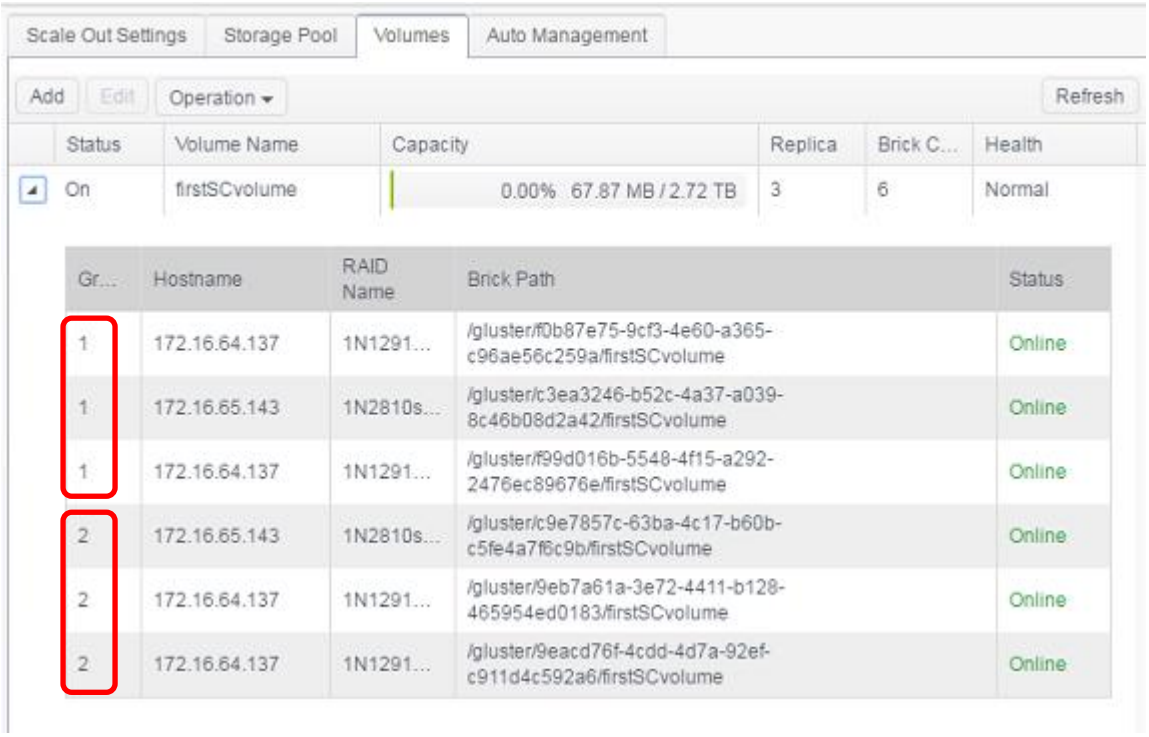


This volume has 2 groups and each group contains three data copies.

**NOTE**

The system will allocate bricks to volume groups automatically. This cannot be assigned manually.

The Scale-Out volume capacity is optimized through thin-provisioning. Users can create as many Scale-Out volumes as needed.

Let's create the 2nd Scale-Out volume with the name "secondSCvolume" at replica level 2. Same steps as above have been taken as can be seen below.

First Scale-Out volume "firstSCvolume" in two groups, three data copies.

| Status | Volume Name | Capacity | Replica | Brick C... | Health |
|---|---|---|---|---|---|
| On | firstSCvolume | 0.00% 67.93 MB / 2.72 TB | 3 | 6 | Normal |

| Gr... | Hostname | RAID Name | Brick Path | Status |
|---|---|---|---|---|
| 1 | 172.16.64.137 | 1N1291... | /gluster/f0b87e75-9cf3-4e60-a365-c96ae56c259a/firstSCvolume | Online |
| 1 | 172.16.65.143 | 1N2810s... | /gluster/c3ea3246-b52c-4a37-a039-8c46b08d2a42/firstSCvolume | Online |
| 1 | 172.16.64.137 | 1N1291... | /gluster/f99d016b-5548-4f15-a292-2476ec89676e/firstSCvolume | Online |
| 2 | 172.16.65.143 | 1N2810s... | /gluster/c9e7857c-63ba-4c17-b60b-c5fe4a7f6c9b/firstSCvolume | Online |
| 2 | 172.16.64.137 | 1N1291... | /gluster/9eb7a61a-3e72-4411-b128-465954ed0183/firstSCvolume | Online |
| 2 | 172.16.64.137 | 1N1291... | /gluster/9eacd76f-4cdd-4d7a-92ef-c911d4c592a6/firstSCvolume | Online |

Second Scale-Out volume “secondSCvolume” in three groups, two data copies.

| Status | Volume Name | Capacity | Replica | Brick C... | Health |
|---|---|---|---|---|---|
| On | firstSCvolume | 0.00% 67.92 MB / 2.72 TB | 3 | 6 | Normal |
| On | secondSCvolume | 0.00% 68.74 MB / 4.54 TB | 2 | 6 | Normal |

| Gr... | Hostname | RAID Name | Brick Path | Status |
|---|---|---|---|---|
| 1 | 172.16.64.137 | 1N1291... | /gluster/f0b87e75-9cf3-4e60-a365-c96ae56c259a/secondSCvolume | Online |
| 1 | 172.16.65.143 | 1N2810s... | /gluster/c3ea3246-b52c-4a37-a039-8c46b08d2a42/secondSCvolume | Online |
| 2 | 172.16.64.137 | 1N1291... | /gluster/f99d016b-5548-4f15-a292-2476ec89676e/secondSCvolume | Online |
| 2 | 172.16.65.143 | 1N2810s... | /gluster/c9e7857c-63ba-4c17-b60b-c5fe4a7f6c9b/secondSCvolume | Online |
| 3 | 172.16.64.137 | 1N1291... | /gluster/9eb7a61a-3e72-4411-b128-465954ed0183/secondSCvolume | Online |
| 3 | 172.16.64.137 | 1N1291... | /gluster/9eacd76f-4cdd-4d7a-92ef-c911d4c592a6/secondSCvolume | Online |

# Volume Expansion and Repairing

The great thing about Scale-Out is the capability of dynamic expansion. There is a tab on the menu bar named “Operation”. Click on it and it will show options dedicated to global capacity expansion and repairing.

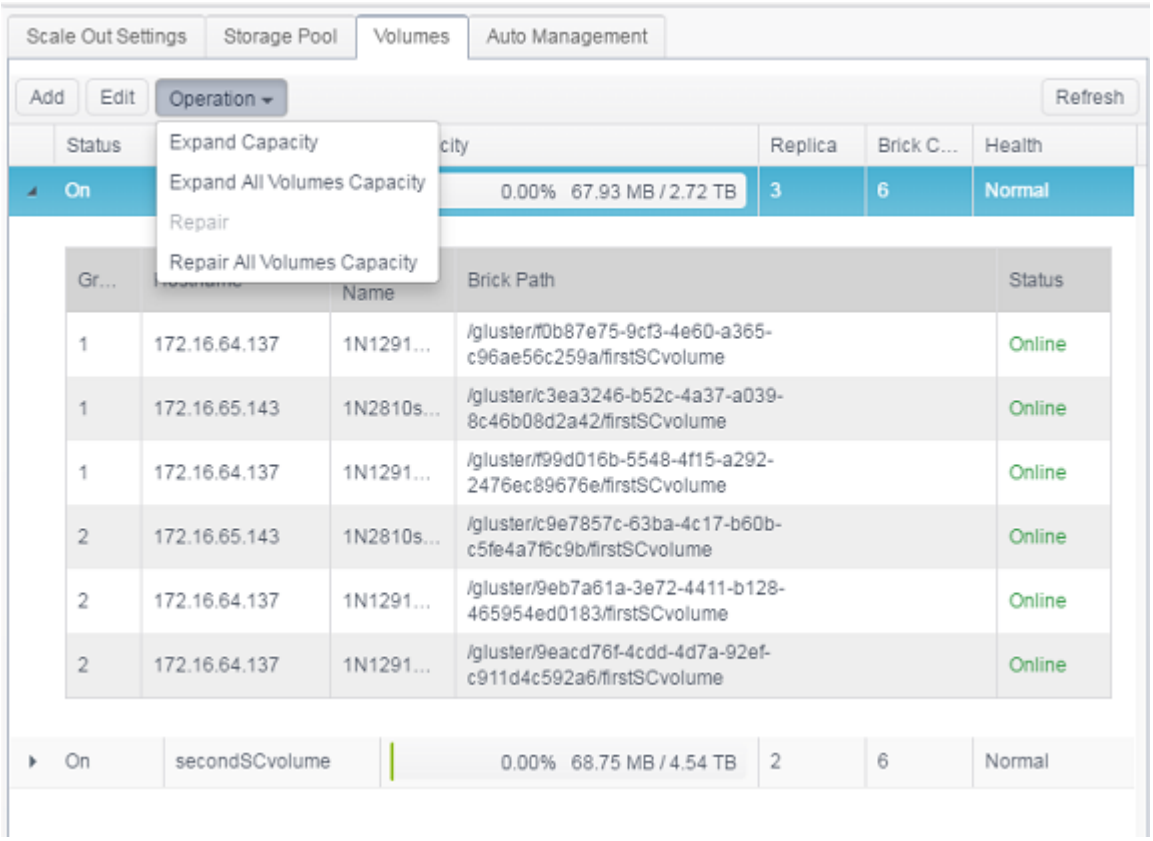


If the expansion is performed for a particular volume, select the associated volume and click on “Expand Capacity”, then the system will auto check available bricks and proceed with the volume expansion. If it is required, expand capacity for all volumes by choosing “Expand All Volume Capacity”.

Other than using available bricks to expand volume capacity, it can be used to repair damage in bricks. Following the same scenario as above, the user can select a particular one or perform the operation for all volumes.

In the example below, one bricks is damaged on Peer 172.16.65.143 and the system has detected the issue, the volume status will show that the Scale-Out volume is ‘Abnormal” and also list the damaged brick.

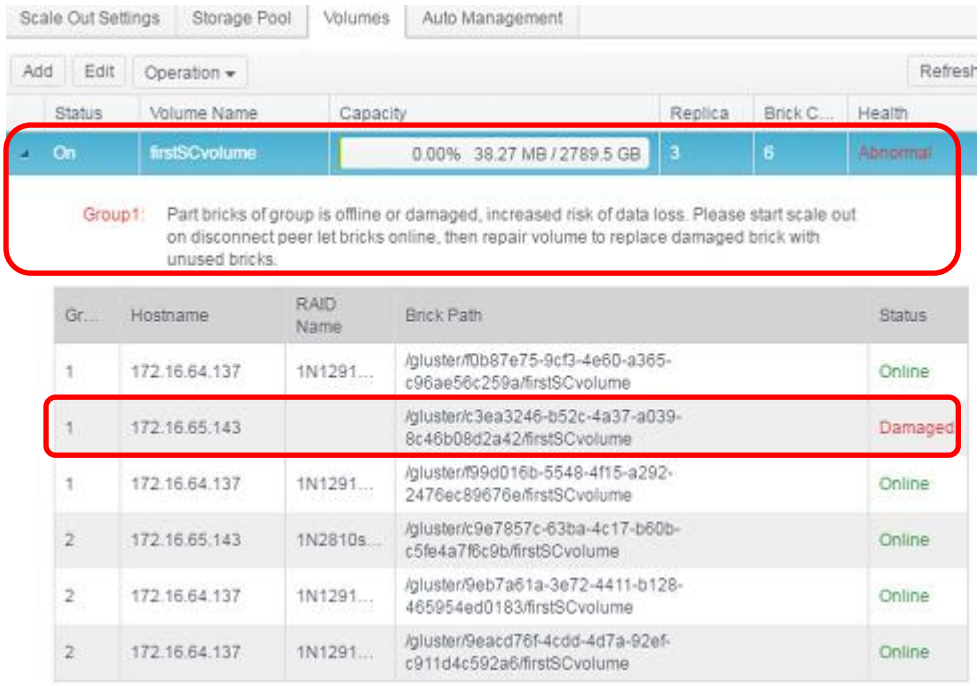


Let's repair the brick by replacing it with a new one. So we have created a new volume named"1N2810SCrepair".

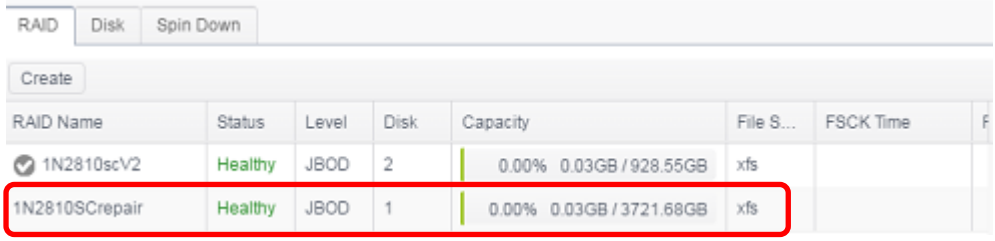


Now we could go to the Scale-Out volume setting page and choose "Repair All Volumes Capacity".

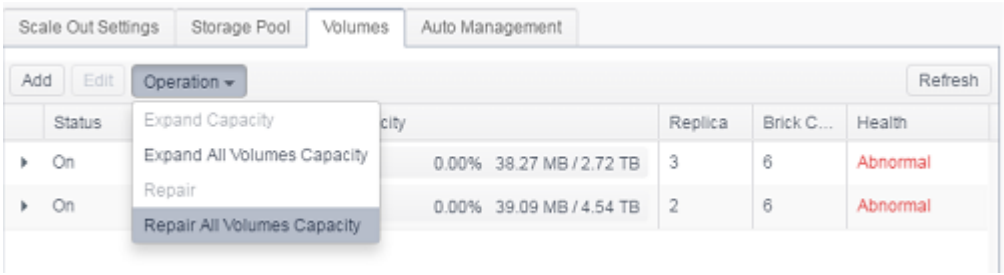


Now the volume has been repaired and is back to a healthy status.

| Status | Volume Name | Capacity | Replica | Brick C... | Health |
|---|---|---|---|---|---|
| On | 1sSCvolume | 0.00% 68.95 MB / 3.63 TB | 2 | 6 | Normal |

| Gr... | Hostname | RAID Name | Brick Path | Status |
|---|---|---|---|---|
| 1 | 172.16.64.137 | 1N1291... | /gluster/f0b87e75-9cf3-4e60-a365-c96ae56c259a/1sSCvolume | Online |
| 1 | 172.16.65.143 | 1N2810s... | /gluster/d9bf57a4-ce52-4429-96ef-fe670039aab9/1sSCvolume | Online |
| 2 | 172.16.64.137 | 1N1291... | /gluster/f99d016b-5548-4f15-a292-2476ec89676e/1sSCvolume | Online |
| 2 | 172.16.65.143 | 1N2810s... | /gluster/c9e7857c-63ba-4c17-b60b-c5fe4a7f6c9b/1sSCvolume | Online |
| 3 | 172.16.64.137 | 1N1291... | /gluster/9eb7a61a-3e72-4411-b128-465954ed0183/1sSCvolume | Online |
| 3 | 172.16.64.137 | 1N1291... | /gluster/9eacd76f-4cdd-4d7a-92ef-c911d4c592a6/1sSCvolume | Online |

**Scale-Out Volume Editing:**

The created Scale-Out volume can be edited by selecting it from the available list then clicking on the “Edit” button. The Scale-Out volume can be stopped by clicking on the Stop Volume” button. Once confirmed, the Scale-Out volume status will change to “Off” and it is inaccessible from the Scale-Out client. Alternatively, this Scale-Out volume can be changed through some advanced setting pertaining to cache size, Write Behind Cache Size and IO Thread.

Rules to accept or reject connections can also be specified for the Scale-Out volume
. Click on “Security Setting” and the screen below appears.

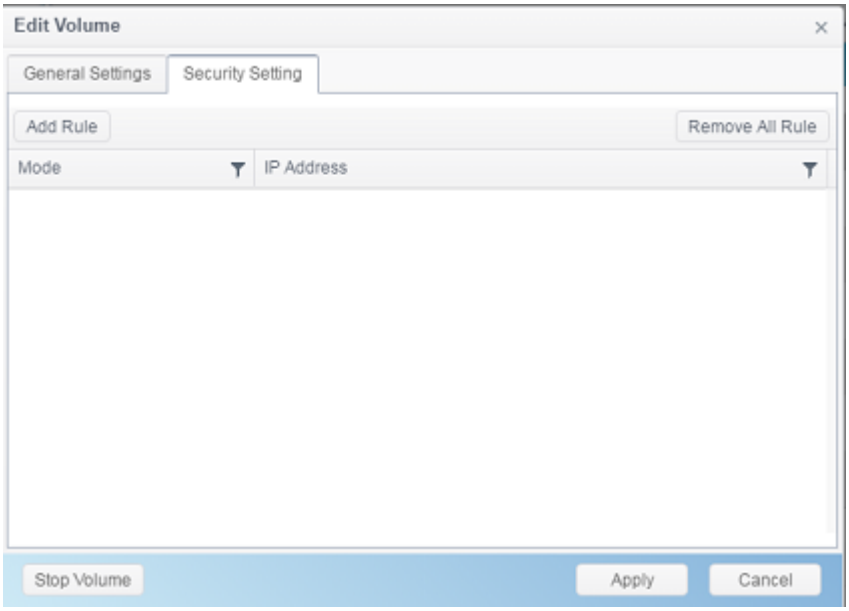


Click on “Add Rule” to add a new connection definition or “Remove All Rules” to clear the list.

# Auto Management

Scale-Out can be empowered to work smartly by enabling “Auto Management”. There are 3 different settings that can be enabled.

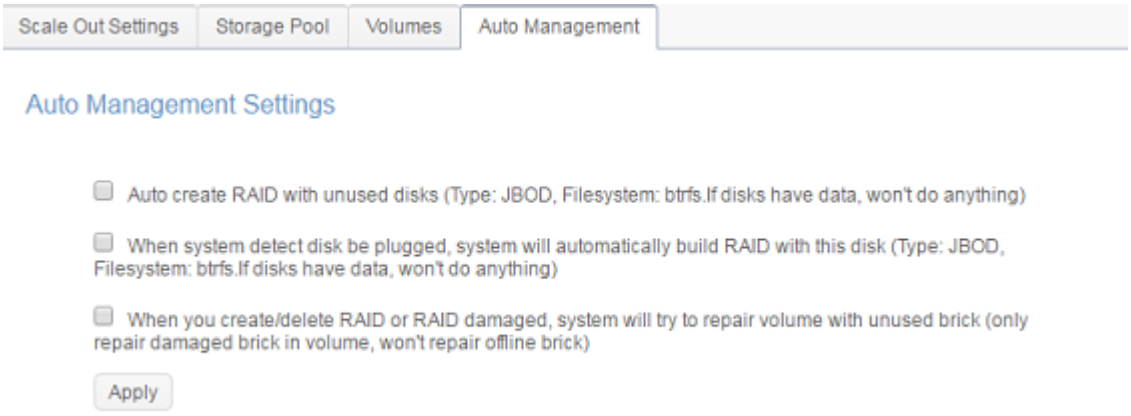


1. If scenario one is enabled, if the system is cold booted with a disk installed (clean), then the RAID volume will be crated automatically by default and this RAID volume will become an unused brick.
2. If the 2nd scenario is enabled, if a disk is plugged in (hot plug-in, disk clean), then the system will create a RAID volume automatically and this RAID volume will become unused brick.
3. If any brick from volume group is damaged, the system will use unused bricks to repair it.

# Stop Scale-Out

In any case, if the Scale-Out server needs to stop, click on the “Stop Scale-Out” button. The status for the Scale-Out server will change to “Offline”.

Example for stopping Scale-Out server on system IP 172.16.65.143:

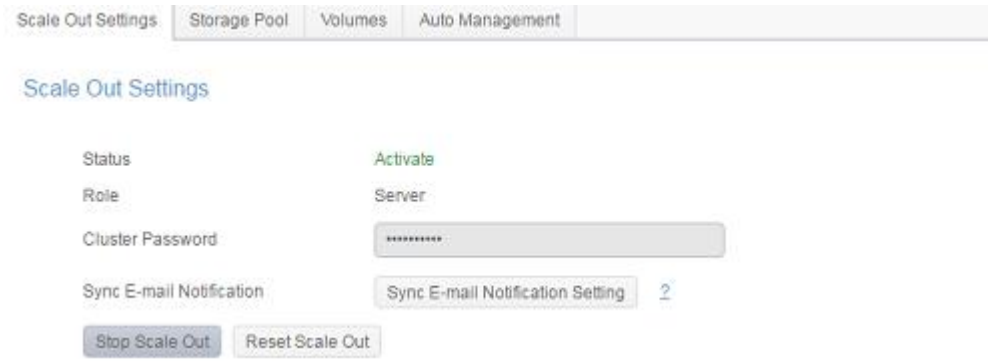


If you now check on any member of the associated Scale-Out server group ex. 172.16.64.137 the status of the system IP 172.16.65.143 will be shown as “Offline”.

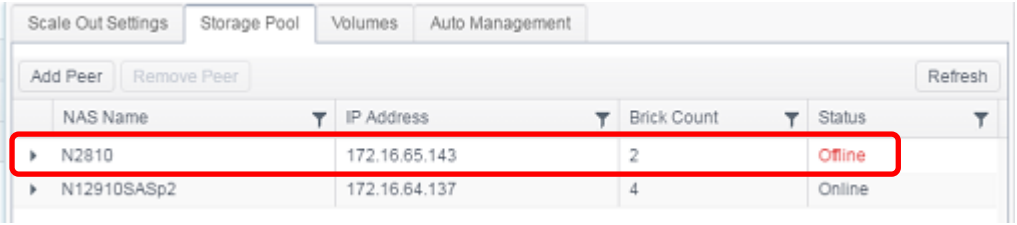


Stopping the Scale-Out server won't affect any data existing in the Scale-Out server volumes. Re-start the Scale-Out server will put this Scale-Out member back to online status.

## Reset Scale-Out

If the Scale-Out server needs to be removed from the Scale-Out group, click on "Reset Scale-Out" to apply.

**WARNING: Once it has been confirmed, all data inside the Scale-Out volumes will be destroyed completely and there is no way to get it back.**

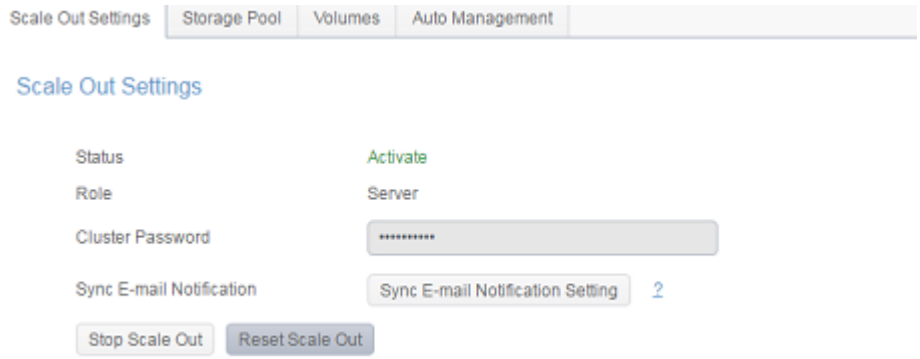


## Create Scale-Out Shared Folder

After the Scale-Out server group setup has been completed, the Scale-Out volume will be ready to use. Next, we will go through how a Scale-Out client can connect to a Scale-Out volume.

Let's take an N5810PRO with the IP http://172.16.64.185 in Scale-Out client mode. The system default mode for the Scale-Out function is set to client, so there is no need to set it up separately. To connect with the Scale-Out volume, please go to "Share Folder" under "Privilege" in the Control Panel.

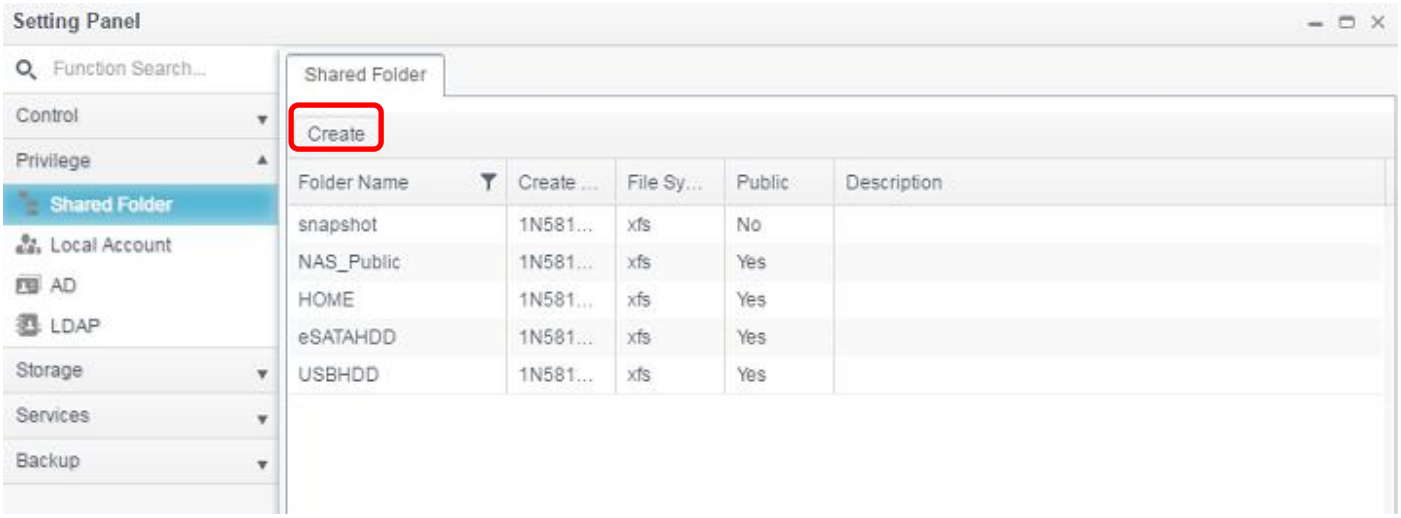


Click on “Create” and screen appear as below.

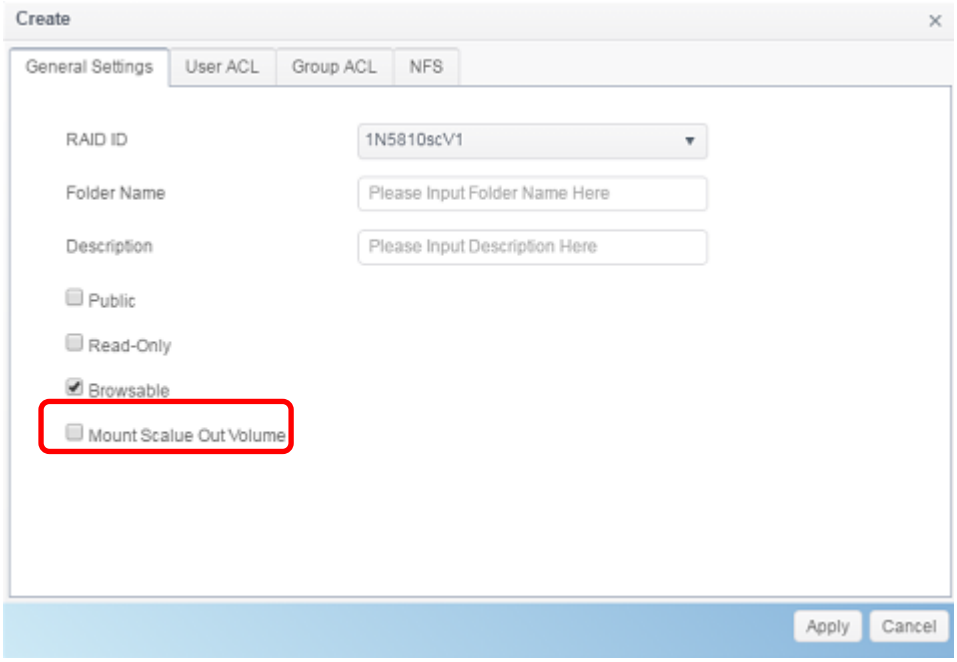


**Steps:**

1. Folder name: To be seen on file protocol level access, let’s choose “firstSCfolder” as a folder name.

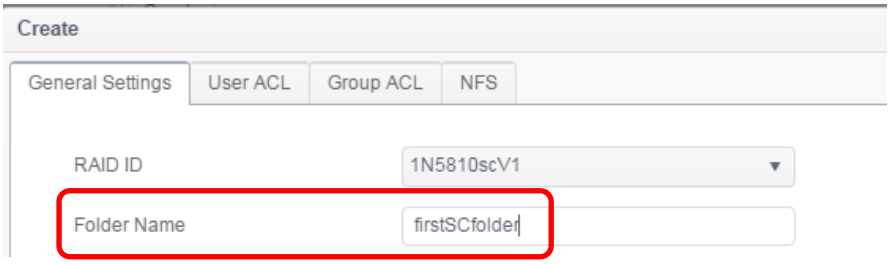


2. Description: Fill in if needed.
3. Public/Read-Only/Browseable: enable/disable as needed.

4. Mount Scale-Out Volume: **Enable** this one and input the Scale-Out Volume IP address and Volume Name.

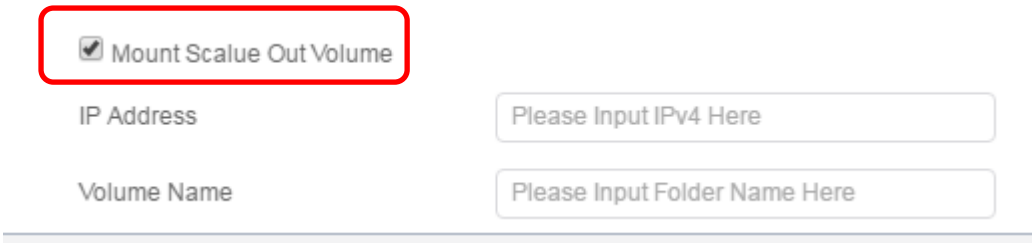


The IP address can be that of any Scale-Out server member, we have 172.16.64.137 and 172. 16.65.143. The volume names are those we have created earlier; they are “firstSCvolume” and “secondSCvolume”. Let’s fill in 172.16.65.143 and “firstSCvolume” then click Apply

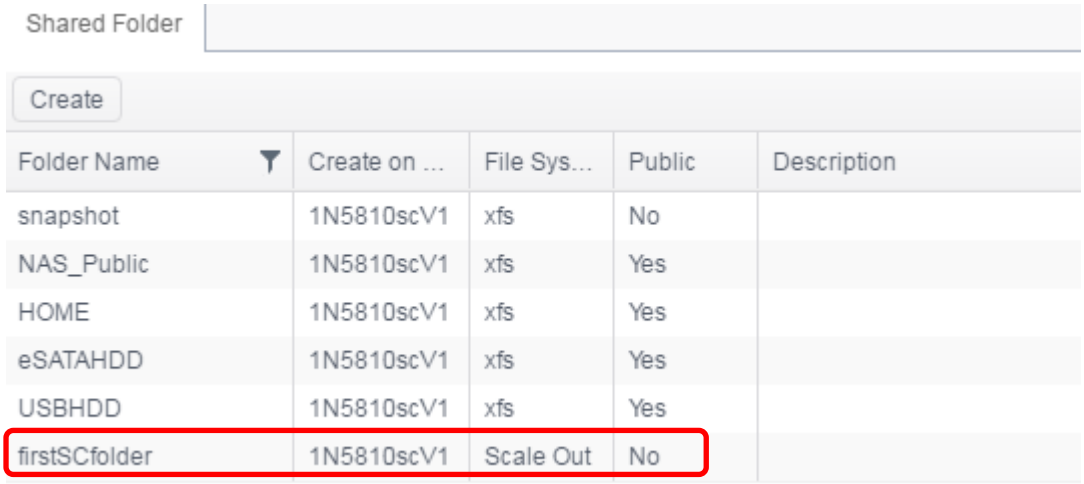


Now in the shared folder list,”firstSCfolder” has been added. It can be used just like a standard shared folder. Let’s connect via Windows and we see that “firstSCfolder” is there; please refer to the screenshot below. Surely you could use the same steps to create more Scale-Out folders, even using the same storage resources.

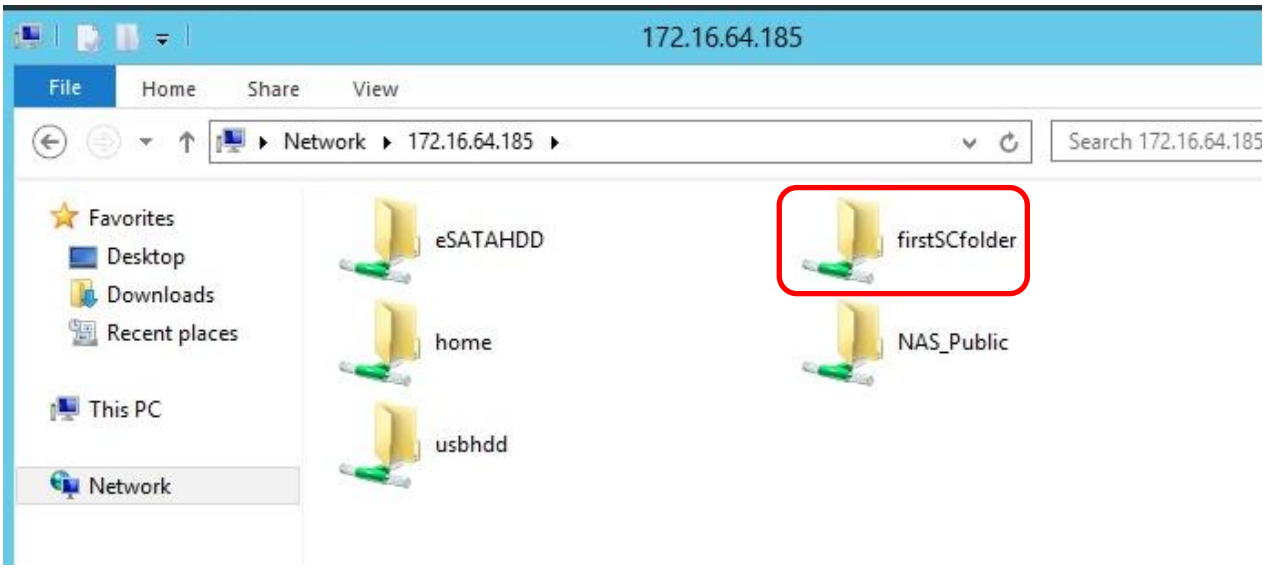

# Support

If you have questions, encounter technical difficulties or need assistance with your NAS or setting up Scale-Out, please contact Thecus technical support under: http://www.thecus.com/sp_tech.php