



# Administrator Guide

Version 2.0



## Revision History

Date	Version	Description
07-06-2017	1.0	Initial version which is based on PetaSAN Release 1.3.0
29-10-2017	1.1	Adding the cluster benchmark and update the Dashboard based on PetaSAN release 1.4.0
08-01-2018	1.2	Add cluster maintenance and path assignment features, update Dashboard ,physical disks list and cluster settings based on PetaSAN release 1.5.0
31-10-2018	1.3	Add CRUSH and Pools Customization based on release 2.1.0
06-12-2018	1.4	Add VLAN support, EC Pools support and SMART monitoring based on release 2.2.0
28-04-2019	1.5	Add Replication Support based on Release 2.3.0
29-07-2020	1.6	Updates based on release 2.5.0 and 2.6.0 mainly the CIFS and NFS support
27-06-2021	1.7	Updates based on Release 2.7 and 2.8 mainly support S3
14-12-2021	1.8	Updates based on Release 3.0
17-10-2023	2.0	Updates based on Release 3.1 and 3.2

## Contents

1. Purpose .....	1
2. Pre-requisites .....	1
3. Dashboard .....	1
➤ Ceph Health.....	1
➤ Storage .....	2
➤ Maintenance .....	2
➤ OSD Usage.....	2
➤ Ceph OSD Status .....	2
➤ View Chart.....	3
4. Configure the cluster.....	4
4.1. Download Certificate .....	4
➤ Download Certificate .....	4
4.2. EC Profiles .....	4
➤ View EC Profiles .....	4
➤ Add EC Profile.....	5
4.3. Manage Pools.....	7
➤ Pools List .....	7
➤ Add Pool.....	8
4.4. Manage CRUSH .....	10
➤ Buckets Tree.....	10
➤ Add Bucket .....	10
➤ Move Bucket .....	11
➤ Delete Bucket.....	12
➤ Rules List .....	12
➤ Add Rule .....	13
4.5. Ceph Configuration .....	14
➤ View Ceph Configuration .....	14
➤ Add Ceph Configuration Key .....	16
➤ Delete Ceph Configuration Key Value.....	16

4.6.	NUMA Pinning.....	17
➤	View NUMA Pinning.....	17
➤	Update Pinning.....	17
4.7.	File Systems.....	18
➤	View File Systems.....	18
➤	Add File System.....	18
➤	Edit File System.....	19
➤	Add layout.....	20
4.8.	S3 Configuration.....	21
➤	View Zonegroups.....	21
➤	Pull.....	22
➤	Import Peer Certificate.....	23
➤	Add Zonegroup.....	23
➤	Zones View List.....	25
➤	Add Zone.....	26
4.9.	General Settings.....	29
➤	Update General Settings.....	29
4.10.	iSCSI Settings.....	30
➤	Update iSCSI Settings.....	30
4.11.	CIFS Settings.....	31
➤	Update CIFS Settings.....	31
➤	Active Directory Settings.....	33
4.12.	NFS Settings.....	35
➤	Update NFS Settings.....	35
4.13.	S3 Settings.....	36
➤	Update S3 Settings.....	36
5.	Manage Nodes.....	40
➤	Node List.....	40
➤	Node Physical Disk List.....	41
➤	Manage Roles.....	43
➤	View Log.....	44
6.	Manage iSCSI Disks.....	45

6.1.	iSCSI Disks.....	45
➤	Add iSCSI Disk.....	45
➤	iSCSI Disk List.....	46
6.2.	Path Assignment .....	53
6.3.	Replication .....	56
➤	Add Replication Job.....	56
➤	Jobs List .....	60
➤	Active Jobs List .....	61
➤	Add Replication User.....	62
➤	Replication Users List .....	63
➤	Add Destination Cluster .....	64
➤	Destination Clusters List.....	65
7.	Manage CIFS.....	66
7.1.	CIFS Shares .....	66
➤	Add CIFS Share .....	66
➤	View CIFS Shares .....	67
7.2.	CIFS Status.....	71
➤	View CIFS Status.....	71
➤	View Connections .....	71
7.3.	Reassign IPs.....	72
8.	Manage NFS .....	72
8.1.	NFS Exports .....	72
➤	Add NFS Export .....	72
➤	View NFS Exports .....	74
8.2.	NFS Status .....	77
➤	View NFS Status .....	77
➤	View Connections .....	77
8.3.	Reassign IPs.....	78
9.	Manage S3.....	78
9.1.	S3 Users.....	78
➤	View S3 Users List .....	78
➤	Add S3 User.....	79

➤ Edit S3 User .....	81
➤ Add Subuser .....	82
➤ Edit Subuser .....	83
9.2. S3 Status.....	84
➤ View S3 Status.....	84
9.3. Reassign IPs .....	85
10. Cluster Benchmark.....	85
➤ Benchmark .....	85
11. Maintenance .....	88
11.1. Cluster Maintenance.....	88
11.2. Backfill Speed .....	89
11.3. Scrub Speed.....	89
11.4. Ceph Balancer .....	90
11.5. OSD CRUSH Weight.....	90
12. Manage Users .....	91
➤ Add User.....	91
➤ User List.....	91
➤ Change Password .....	92

## 1. Purpose

The purpose of this guide is to describe how you can manage your PetaSAN cluster using the web based Management Application.

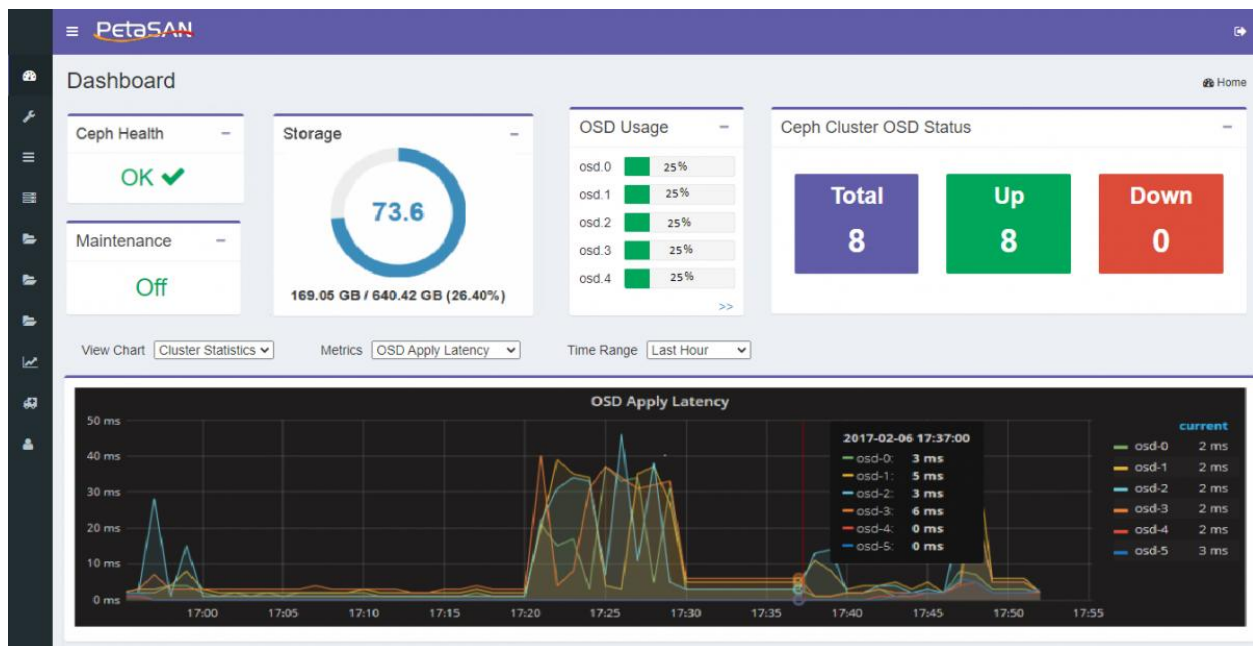
## 2. Pre-requisites

This guide assumes you have already built your cluster and you are logged in to the Management Application via one of the management URLs.

- *Note: PetaSAN's default administrator password is "password" which can be changed once you are logged in.*

## 3. Dashboard

The home page of the Management Application is a dashboard which provides a quick look at the state of the cluster.



### ➤ Ceph Health

Ceph Health will be described by one of the following icons according to the current Ceph status:

- **OK**  
The cluster is healthy
- **Warning**  
The cluster has warnings, which can be viewed by clicking on the icon.
- **Error**  
The cluster has errors, which can be viewed by clicking on the icon.

➤ **Storage**

This chart shows the free as well as the total storage available in the cluster.

➤ **Maintenance**

The maintenance section shows whether the system is under maintenance or not, the system is under maintenance if any of the maintenance settings is off.

For example during system upgrade it is recommended to turn off the fencing.

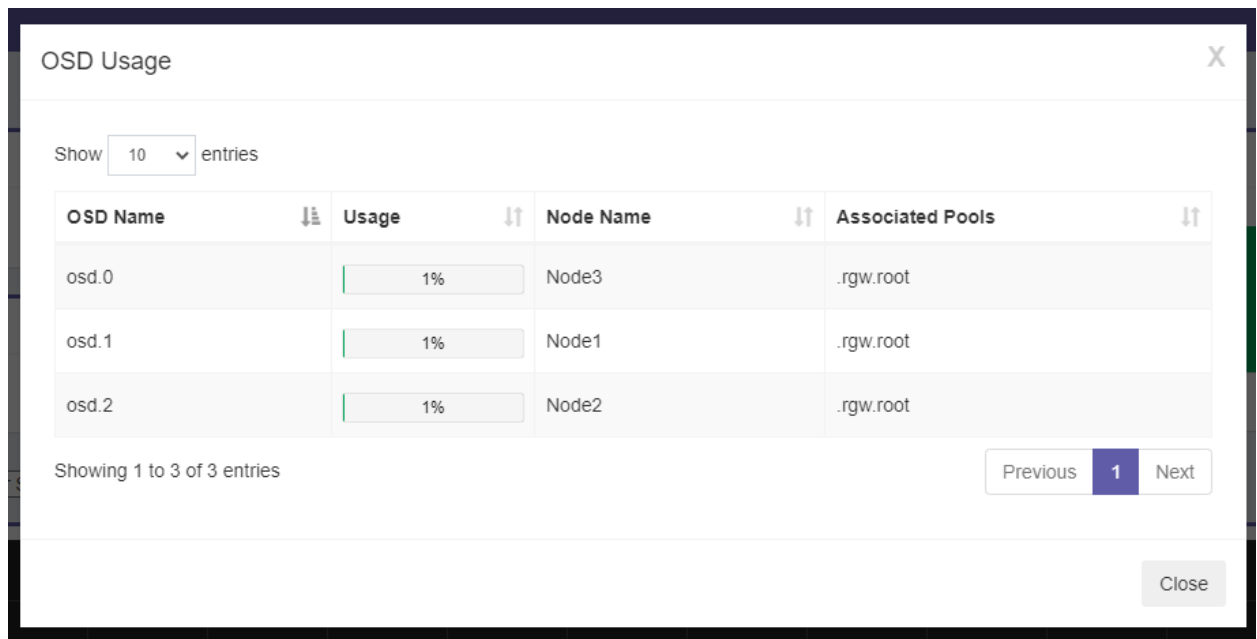
➤ **OSD Usage**

The OSD usage section shows the top 5 used OSDs and the percentage of used space for each OSD

You can view the full list of OSD usage by clicking on >>

Which will shows the following

- OSD name
- Usage: percentage used of the OSD
- Node Name: the node which hosts the OSD
- Associated Pools: Pool names that uses the OSD



➤ **Ceph OSD Status**

The dashboard displays the following information about the Ceph OSDs:

- **Total**  
The total number of OSDs in the cluster
- **Up**  
The number of running OSDs
- **Down**



The number of down OSDs, in case there are any down OSDs you can view the list of down OSDs and the node(s) hosting them.

### ➤ **View Chart**

You can view the following charts which displays various cluster metrics for various time ranges starting from one hour and ending with one year:

- **Cluster Storage**

This chart represents the used and free storage available.

- **Throughput**

This chart represents the Read/Write bandwidth for the selected pool.

- **IOPS**

This chart represents the Read/Write operations per second for the selected pool.

- **Monitor Status**

This chart represents the number of Ceph monitor nodes and their status.

- **OSD Status**

This chart represents the number of Ceph OSDs and their status.

- **OSD Commit Latency**

This chart represents the time latency during OSD commit operations.

- **OSD Apply Latency**

This chart represents the time latency during OSD apply operations.

- **PG Status**

This chart shows the number of Ceph PGs and their state.

- **Node Statistics**

Select the node you want to view its statistics and then select the chart you want to view from one of the following options:

- **CPU**

Total CPU percentage utilization for node.

- **Memory**

Total RAM percentage utilization for node.

- **Disk Utilization**

Disk percentage utilization, indicates how busy your disks are.

- **Disk Throughput**

Disk Read/Write transfer rate.

- **Disk IOPS**

Disk total operations per second.

- **Disk Reallocated Sectors**

Disk reallocated sectors using SMART Tools.

- **Disk Power On Hours**

Disk number of power on hours read using SMART Tools.

- **Disk Temperature**

Disk temperature read using SMART Tools.

- **Network Utilization**  
Network percentage utilization, indicates how busy your interface cards are.
- **Network Throughput**  
Network Read/Write transfer rate.

## 4. Configure the cluster

### 4.1. Download Certificate

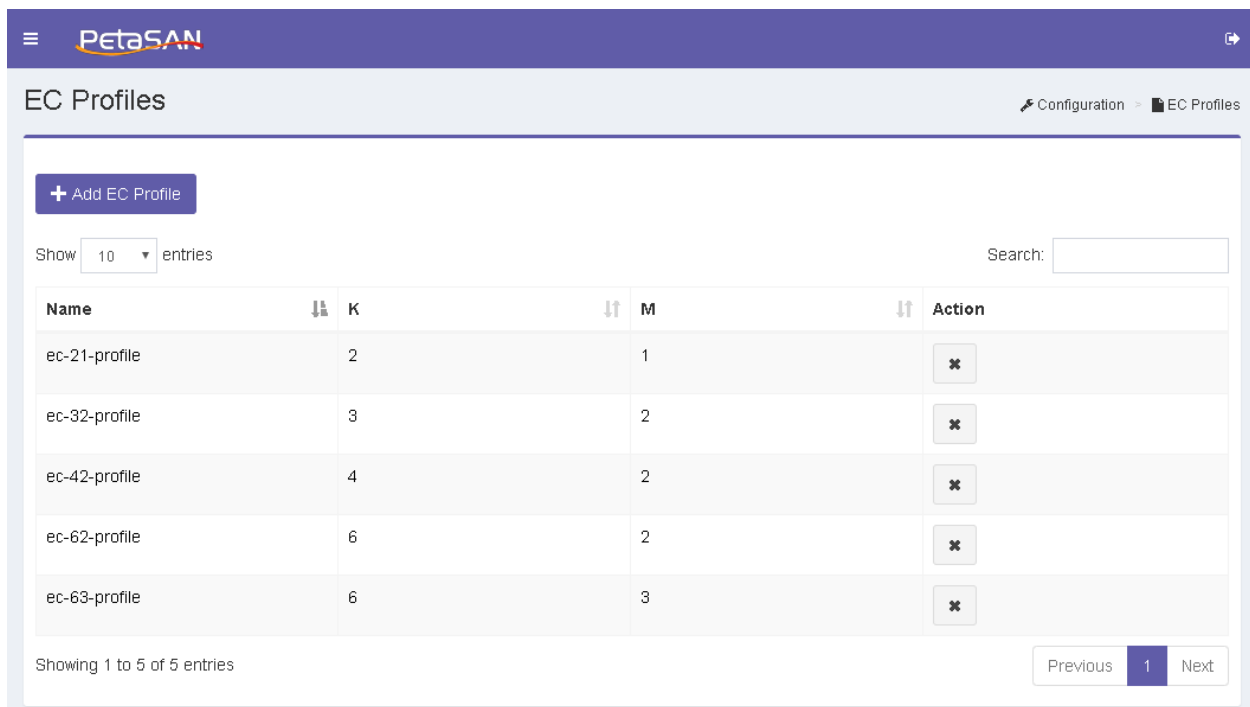
#### ➤ Download Certificate

The download certificate link enables you to download the browser certificate that you should install to connect securely using https.

### 4.2. EC Profiles

#### ➤ View EC Profiles

The system views the list of predefined EC profiles and enables you to add new profiles:



The screenshot shows the 'EC Profiles' page in the PetaSAN web interface. At the top, there is a navigation bar with the PetaSAN logo and a breadcrumb trail: 'Configuration > EC Profiles'. Below the navigation bar, there is a '+ Add EC Profile' button. The main content area features a table with the following columns: Name, K, M, and Action. The table contains five entries:

Name	K	M	Action
ec-21-profile	2	1	
ec-32-profile	3	2	
ec-42-profile	4	2	
ec-62-profile	6	2	
ec-63-profile	6	3	

Below the table, there is a pagination control showing 'Showing 1 to 5 of 5 entries' and buttons for 'Previous', '1', and 'Next'. A search bar is located at the top right of the table area.

The EC Profiles List displays the following information for each Profile:

- **Name**  
The EC profile name
- **K**  
Number of data chunks used to divide each stored object, each chunk is stored on a different OSD.
- **M**

Number of coding chunks computed for each stored object, each chunk is stored on a different OSD. This number determines how many OSDS can simultaneously fail without loss of data.

- **Action**

The column action displays the actions that can be done on each profile; each action has a tool tip that is displayed when the mouse is over the action button.

- Delete

Deletes the selected profile

➤ **Add EC Profile**

By clicking on Configuration->EC Profiles-> Add Profile, the system will open the Add EC Profile form

The screenshot shows the 'Add EC Profile' form in the PetaSAN administrator interface. The form is titled 'Add EC Profile' and has a breadcrumb trail: Configuration > EC Profiles > Add EC Profile. The form contains the following fields:

- Name:** A text input field with a red asterisk indicating it is required.
- K:** A text input field with a red asterisk and an information icon.
- M:** A text input field with a red asterisk and an information icon.
- Advanced:** A collapsed section containing:
  - Plugin:** A dropdown menu with 'Jerasure' selected.
  - Technique:** A dropdown menu.
  - Stripe unit:** A text input field with an information icon.
  - Packet Size:** A text input field.

At the bottom right of the form, there are two buttons: 'Cancel' and 'Save'.

The Add EC Profile form contains the following fields:

- **Name**

The EC profile name

- **K**

Number of data chunks used to divide each stored object, each chunk is stored on a different OSD.

- **M**

Number of coding chunks computed for each stored object, each chunk is stored on a different OSD. This number determines how many OSDS can simultaneously fail without loss of data.

### Advanced fields

- **Plug in**  
The erasure code plugin used to compute coding chunks and recover missing chunks; you can select one of the plugins: Jerasure, ISA, Locally Repairable or SHEC.
- **Technique**  
The techniques available in the selected plugin.
- **Stripe Unit**  
The amount of data stored per stripe, you should enter numeric value in bytes, usage of K or M is allowed to specify values in Kilobytes or megabytes respectively.value should be multiple of 4K bytes
- **Packet Size**  
The encoding will be done on packets of bytes size at a time. Appear in case plugin selected is Jerasure.
- **Locality**  
Group the coding and data chunks into sets of size locality. Appear in case plugin selected is Locally Repairable.
- **Durability Estimator**  
The number of parity chunks each of which includes each data chunk in its calculation range, appear in case plugin selected is SHEC.

## 4.3. Manage Pools

### ➤ Pools List

Clicking on Configuration ->Pools, the system will view list of existing pools.

You can filter the pools by its usage by selecting the required usage from the combobox.

Name	Type	Usage	PGs	Size	Min Size	Rule Name	Used Space	Available Space	Active OSDs	Status	Action
rgw.root	replicated	radosgw	32	3	2	replicated_rule	2.25 MB	281.79 GB		Active	
default.rgw.buckets.data	replicated	radosgw	64	3	2	replicated_rule	6.0 MB	281.79 GB		Active	
default.rgw.buckets.index	replicated	radosgw	16	3	2	replicated_rule	0 Bytes	281.79 GB		Active	
default.rgw.control	replicated	radosgw	16	3	1	replicated_rule	0 Bytes	281.79 GB		Active	
default.rgw.custom.buckets.data	replicated	radosgw	64	3	2	replicated_rule	0 Bytes	281.79 GB		Active	
default.rgw.custom.buckets.index	replicated	radosgw	64	3	2	replicated_rule	0 Bytes	281.79 GB		Active	
default.rgw.custom.glacier.buckets.data	replicated	radosgw	64	3	2	replicated_rule	0 Bytes	281.79 GB		Inactive	
default.rgw.glacier.buckets.data	replicated	radosgw	64	3	2	replicated_rule	0 Bytes	281.79 GB		Active	
default.rgw.log	replicated	radosgw	16	3	1	replicated_rule	0 Bytes	281.79 GB		Active	
default.rgw.meta	replicated	radosgw	16	3	1	replicated_rule	1.12 MB	281.79 GB		Active	

The Pools List displays the following information for each Pool:

- **Name**  
The pool name
- **Type**  
The pool type, Replicated or EC
- **Usage**  
rbd , cephfs or radsgw, rbd is used for iSCSI, cephfs for CIFS/NFS and radosgw for S3.
- **No of PGs**  
Number of PGs in the Pool
- **Size**
  - In case of replicated pool: Number of replicas per object.
  - In case of EC pool: Number of chunks per object
- **Min Size**
  - In case of replicated pool: The minimum number of replicas required for the pool to remain active and keep serving client I/O requests.
  - In case of EC pool: The minimum number of chunks required for the pool to remain active and keep serving client I/O requests.
- **Rule Name**

The rule used to distribute the replicas

- **Used Space**

Space used in the pool

- **Available Space**

Remaining space in the pool.

- **Active OSDs**

Number of active OSDs in the pool

- **Status**

Displays the current status of the pool which can be one of the following:

- Active  
Pool can serve client I/O requests.
- Inactive  
Pool cannot serve client I/O requests.
- Deleting  
Pool is being deleted

*Note:*

*Deleting a pool will delete all its stored data, please be careful using this action.*

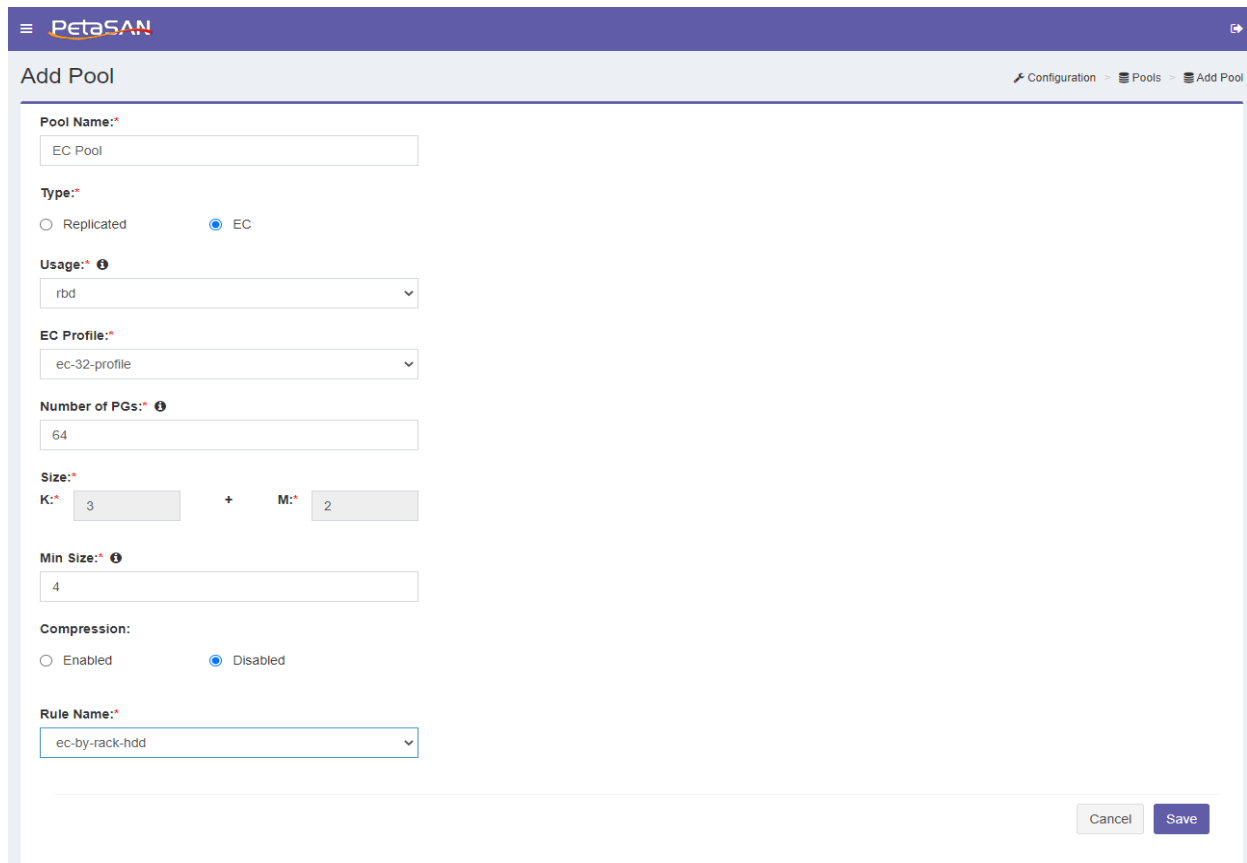
## ➤ Add Pool

By clicking on Configuration->Pools->Add Pool the system will open the Pool form

The screenshot shows the 'Add Pool' configuration form in the PetaSAN administrator interface. The form is titled 'Add Pool' and includes the following fields and options:

- Pool Name:** Text input field containing 'Fast\_VM\_Pool'.
- Type:** Radio button options for 'Replicated' (selected) and 'EC'.
- Usage:** Dropdown menu set to 'rbd'.
- Number of PGs:** Text input field containing '256'.
- Size:** Dropdown menu set to '3'.
- Min Size:** Dropdown menu set to '2'.
- Compression:** Radio button options for 'Enabled' (selected) and 'Disabled'.
- Compression Algorithm:** Dropdown menu set to 'snappy'.
- Rule Name:** Dropdown menu set to 'by-rack-hdd'.

At the bottom right of the form, there are 'Cancel' and 'Save' buttons.



**Add Pool**

Configuration → Pools → Add Pool

**Pool Name:**  
EC Pool

**Type:**  
 Replicated     EC

**Usage:**  
rbd

**EC Profile:**  
ec-32-profile

**Number of PGs:**  
64

**Size:**  
K: 3 + M: 2

**Min Size:**  
4

**Compression:**  
 Enabled     Disabled

**Rule Name:**  
ec-by-rack-hdd

Cancel Save

The Add Pool form contains the following fields:

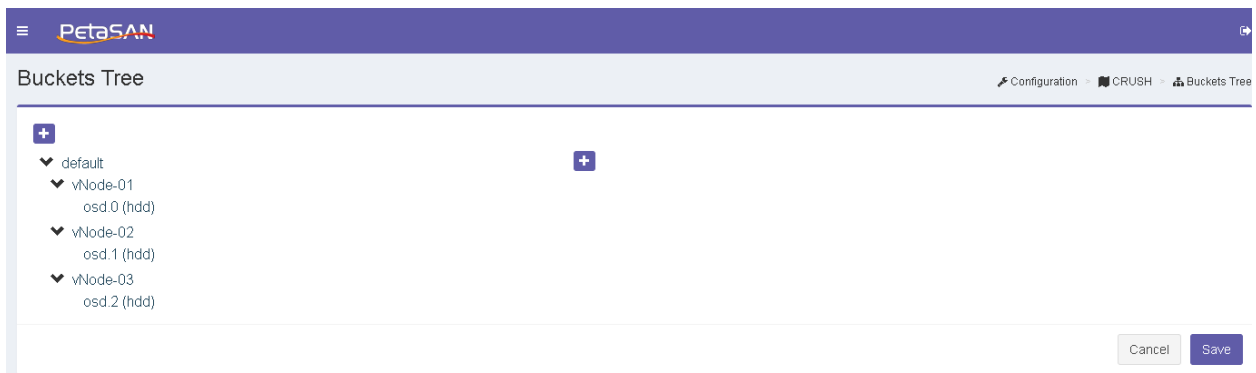
- **Pool Name**  
Name of the pool, name cannot contain spaces
- **Pool Type**  
Type of pool which can be replicated or EC
- **Usage**  
rbd , cephfs or radosgw, rbd is used for iSCSI, cephfs for CIFS/NFS and radosgw for S3.
- **EC Profile**  
Appear in case of EC pool, to enable selecting an EC profile
- **Number of PGs**  
Number of PGs in the pool, you should choose this value so that ideally each OSD will handle a total PG count of approximately 100 PGs (including replicas) across all its pools
- **Size**
  - In case of replicated pool: Number of replicas in the pool.
  - In case of EC Pool: Total number of chunks per object. Which equals K+M where K is the number of data chunks and M is the number of parity chunks.
- **Min Size**

- In case of replicated pool: The minimum number of replicas required for the pool to remain active and keep serving client I/O requests.
- In case of EC Pool: The minimum number of chunks required for the pool to remain active and keep serving client I/O requests.
- **Compression**  
If you want to enable compression select “Enabled”
- **Compression Algorithm**  
If you enabled compression, select the compression algorithm to use.
- **Rule Name**  
The rule to use which defines how stored data is distributed.

## 4.4. Manage CRUSH

### ➤ Buckets Tree

By clicking on Configuration ->CRUSH->Buckets Tree, the system will open a tree of existing buckets



You will be able to manage buckets as following:

### ➤ Add Bucket

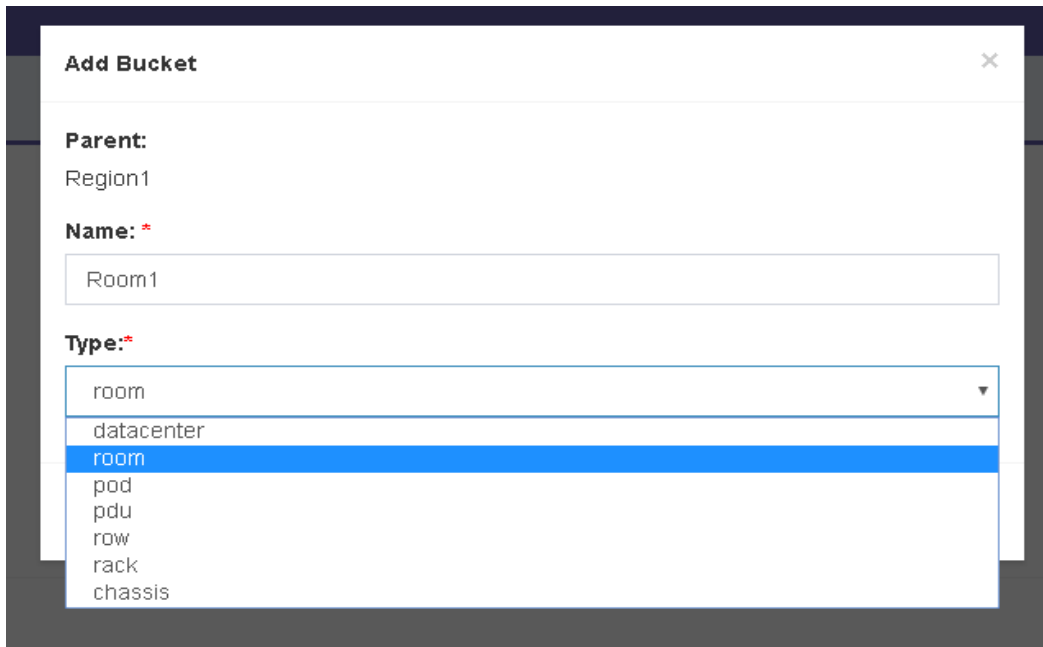
Select a parent bucket then click on the Add Bucket button; this will open the Add New bucket Form.

The newly created bucket will be added under the selected Parent bucket

You can select to add a new bucket as the top level bucket using the Add button on the top left of the tree.

The Add bucket Form contains the following fields:





- **Parent**

The name of the selected parent bucket.

- **Name**

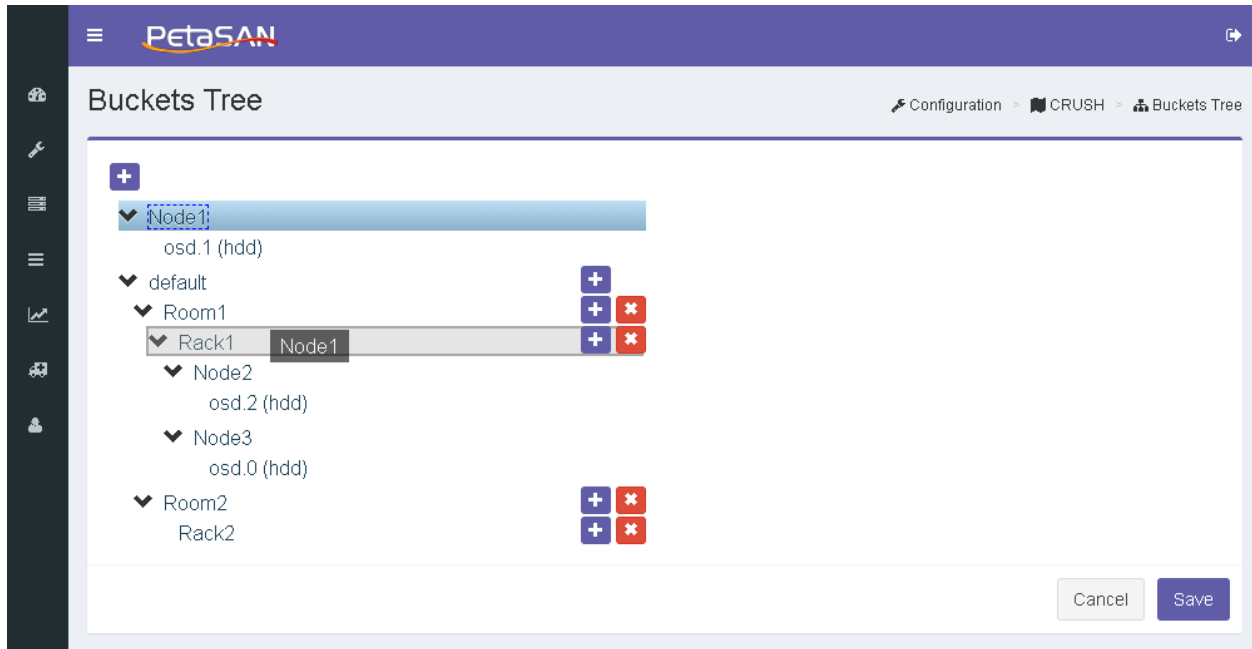
The new bucket name

- **Type**

The new bucket type, you will be able to select types that are allowed Childs of the parent bucket, after entering the bucket information and clicking Add, the bucket will be added to the temporary buckets tree. To save your changes to the CRUSH map, click Save button on the Buckets Tree Form.

➤ **Move Bucket**

You can move a bucket with all of its children by dragging and dropping it under a new parent (Highlighted with grey)



Any movements will not be saved to the CRUSH map until you click Save on the Buckets Tree form.

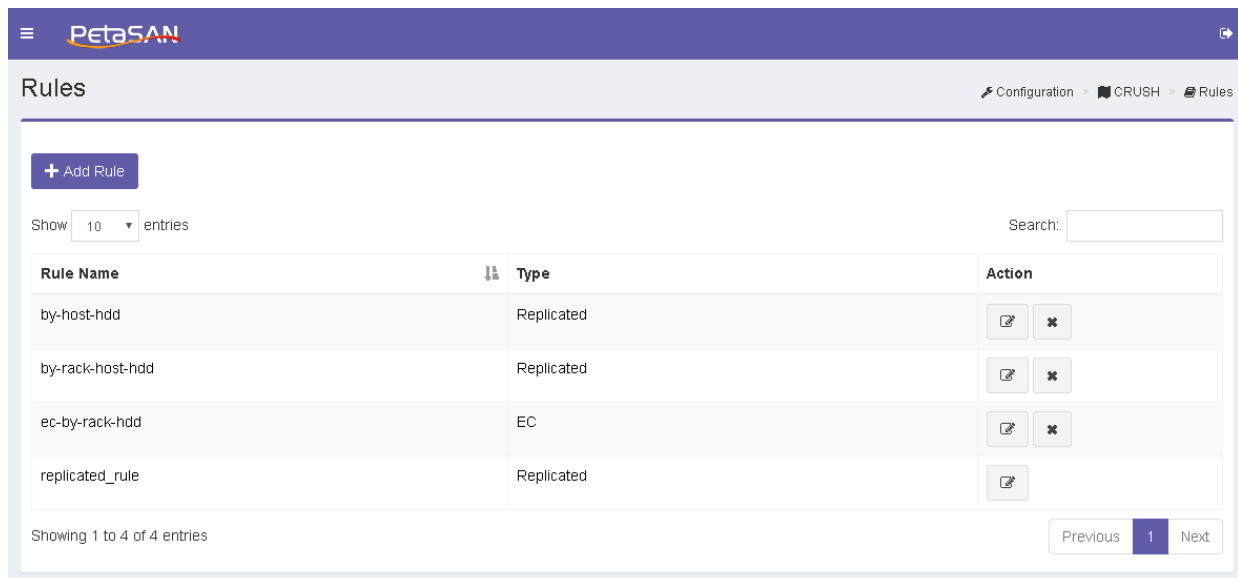
### ➤ Delete Bucket

You can only delete a bucket if it has no children

Any deletions will not be saved to the CRUSH map until you click Save on the Buckets Tree form.

### ➤ Rules List

Configuration ->CRUSH->Rules, the system will view list of all existing rules

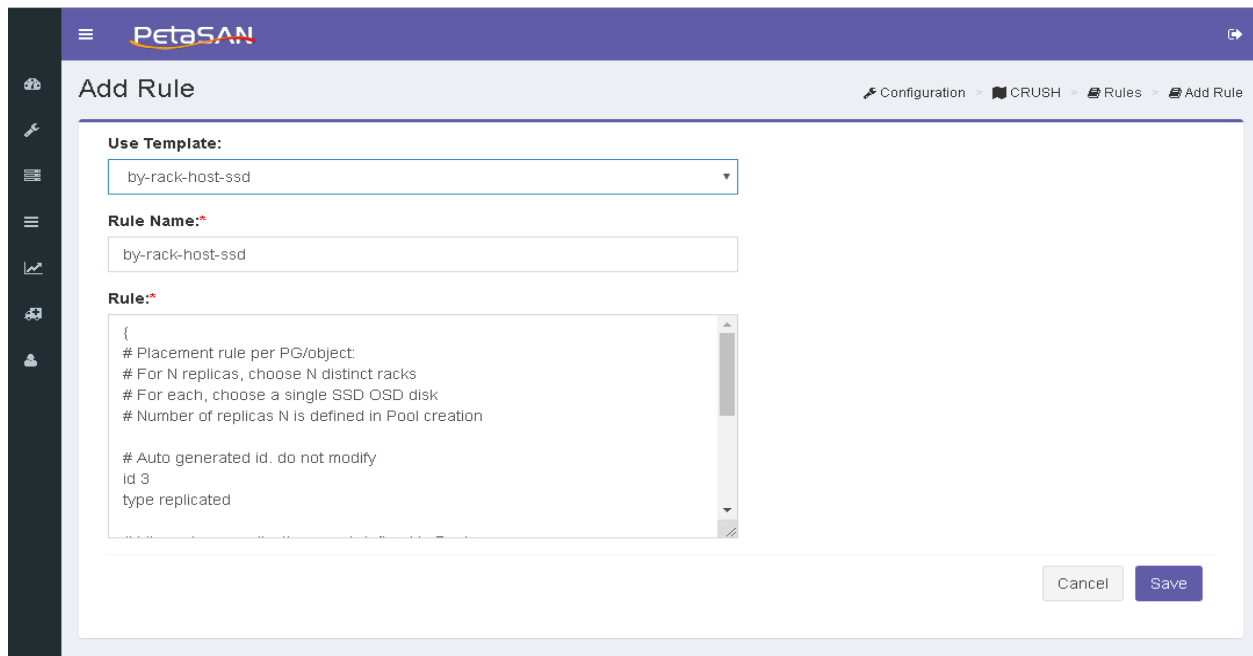


The Rules List displays the following information for each Rule:

- **Rule Name**  
The rule name
- **Type**  
The rule type, replicated or EC
- **Action**  
The column action displays the actions that can be done on each rule; each action has a tool tip that is displayed when the mouse is over the action button.
  - Edit  
Opens the Edit Rule form which allows updating the rule body only.
  - Delete  
Deletes the selected rule with the exception of the default rule “replicated rule” which cannot be deleted

### ➤ Add Rule

By clicking on Configuration->CRUSH->Rules-> Add rule, the system will open the Add Rule



The screenshot shows the 'Add Rule' form in the PetaSAN interface. The breadcrumb trail is Configuration > CRUSH > Rules > Add Rule. The form contains the following fields:

- Use Template:** A dropdown menu with 'by-rack-host-ssd' selected.
- Rule Name:** A text input field containing 'by-rack-host-ssd'.
- Rule:** A text area containing the following rule definition:
 

```
{
# Placement rule per PG/object
# For N replicas, choose N distinct racks
# For each, choose a single SSD OSD disk
# Number of replicas N is defined in Pool creation

# Auto generated id. do not modify
id 3
type replicated
```

At the bottom right of the form, there are 'Cancel' and 'Save' buttons.

The Add Rule form contains the following fields:

- **Use Templates**  
You can select a rule template from list of predefined templates, templates represents some of the most commonly used rules
- **Rule Name**  
The rule name (cannot contain spaces)
- **Rule**  
The body of the rule, the rule id is auto generated automatically by the system.

## 4.5. Ceph Configuration

### ➤ View Ceph Configuration

By clicking on the Ceph configuration you will be able to view configuration keys that are saved in the common cluster database. These apply to all instances of services such as OSDs, Monitors, MDSs..etc however, these will be over-ridden by any settings set for specific instances as well as by any settings defined in local configuration files.

When the mouse is hovered over any key the system displays a help description of that key

You can filter keys by:

- **Section**

The current sections available options are

- All
- global
- osd
- mon
- mgr
- mds
- client

- **Level**

By default system will all the levels or you can select the configuration level from one of the following values:

- All
- Basic
- Advanced
- Dev

The screenshot shows the PetaSAN Ceph Configuration interface. At the top, there is a navigation bar with the PetaSAN logo and a menu icon. Below the navigation bar, the page title is "Ceph Configuration". There are two breadcrumb links: "Configuration" and "Ceph Configuration". A blue information banner at the top states: "View keys that are saved in the common database not including the keys that are set for specific OSD, Monitor, or in the local configuration file and it will override the built-in keys". Below the banner, there is a "View by:" section with two radio buttons: "Section" (selected) and "Category". To the right of this section is an "Add" button. Below the "View by:" section, there are two dropdown menus: "Section:" with "osd" selected and "Level:" with "Basic" selected. A note below the dropdowns reads: "To edit double click the key row and press Enter key or click out of the row to save". Below the note, there is a table with one row: "osd\_memory\_target = 4294967296". To the right of the value is a small square button with an "x" icon.

- **Category**

The current categories options are

- Recovery
- Scrubbing

The screenshot shows the PetaSAN Ceph Configuration interface with the "Category" view selected. The "View by:" section now has "Category" selected. The "Category:" dropdown menu is set to "recovery". The table below the note shows five rows of configuration keys: "osd\_max\_backfills = 1", "osd\_recovery\_max\_active = 1", "osd\_recovery\_priority = 1", "osd\_recovery\_op\_priority = 1", and "osd\_client\_op\_priority = 63". Each row has a small square button with an "x" icon to its right.

## ➤ Add Ceph Configuration Key

You can add new Ceph configuration key

- **Key**

Select the key name you want to add from list of available keys, you can type any characters and system will filter the keys contain the typed characters

System will view the description of the selected key

- **Section**

Select the section you want to add the key to it

- **Value**

Enter the value of the selected key

## ➤ Delete Ceph Configuration Key Value

You can delete one of the existing Ceph configuration keys

## 4.6. NUMA Pinning

### ➤ View NUMA Pinning

You can view the list of OSDs and their NUMA node pinning.

NUMA Pinning

Configuration > NUMA Pinning

Node1

OSD Number	NUMA Node
OSD.1	0 0-1

Node2

Node3

Copyright 2020 PetaSAN www.petasan.org All rights reserved. Version 2.6.0

### ➤ Update Pinning

You can update the pinning for OSDs on specific NUMA nodes

NUMA Pinning

Configuration > NUMA Pinning > Update NUMA Pinning

**Info**  
OSD service will be restarted for any updated OSDs ...

Node2:

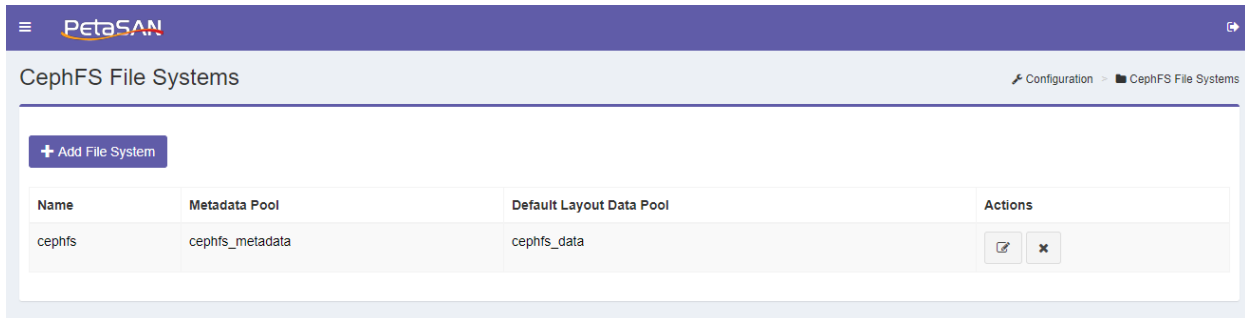
OSD Number	NUMA Node
OSD.2	0 0,1
OSD.3	

Copyright 2020 PetaSAN www.petasan.org All rights reserved. Version 2.6.0

## 4.7. File Systems

### ➤ View File Systems

By Default the cephfs file system is created but you can delete it and create a more customized one.



System views the following columns for the file systems

- **Name**  
The file system name
- **Metadata Pool**  
The file system Metadata Pool
- **Actions**
  - Edit  
Edits the selected file system
  - Delete  
Deletes the selected file system including its metadata and data pools and if there are any layouts, they too will be deleted including their data pools

➤ Note:

Currently you are not allowed to add multiple file systems

### ➤ Add File System

You can add a new file system with its default layout, if you currently have no file system created System requires entering the following fields

- **Name**  
The file system name
- **Metadata Pool**  
Select if from list of existing pools of type cephfs

### Default layout

- **Name**  
Fixed as default
- **Data Pool**  
Select the default layout data pool from list of existing pools of type cephfs



- **Actions**

Select if from list of existing pools of type cephfs

- Save

Saves the file system only if there is currently no other file system exists

The screenshot shows the 'Add File System' configuration page in the PetaSAN interface. The breadcrumb trail is 'Configuration > CephFS File Systems > Add File System'. The form contains the following fields:

- Name:** A text input field containing 'cephfs'.
- Metadata Pool:** A dropdown menu with 'cephfs\_metadata' selected.
- Default Layout:** A section with a **Name:** field containing 'Default'.
- Data Pool:** A dropdown menu with 'cephfs\_data' selected.

At the bottom right of the form are 'Cancel' and 'Save' buttons.

➤ **Edit File System**

After saving the file system, the system will redirect you to the Edit File System form where you will be able to add or delete layouts

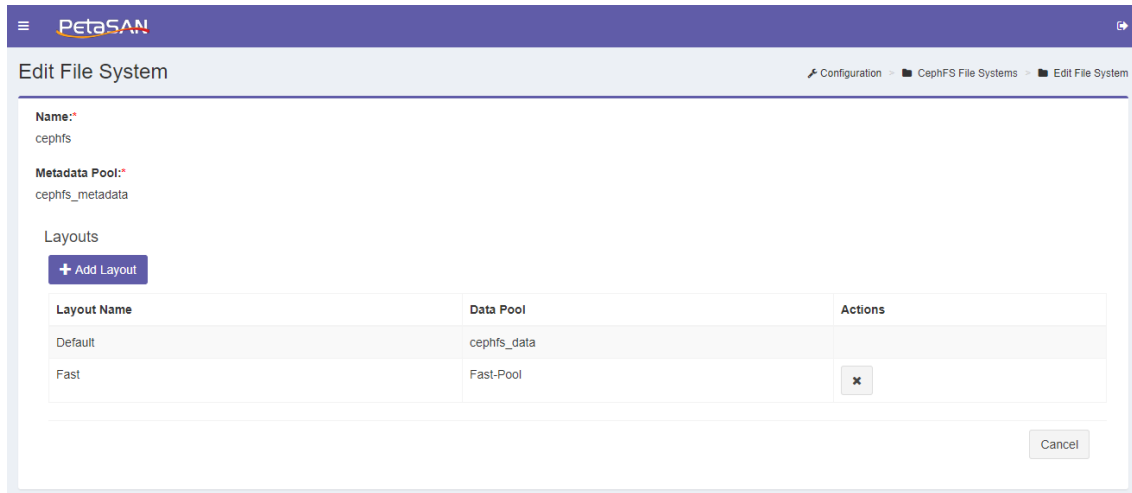
- **Actions**

- Add layout

Enables you to add a new layout

- Delete layout

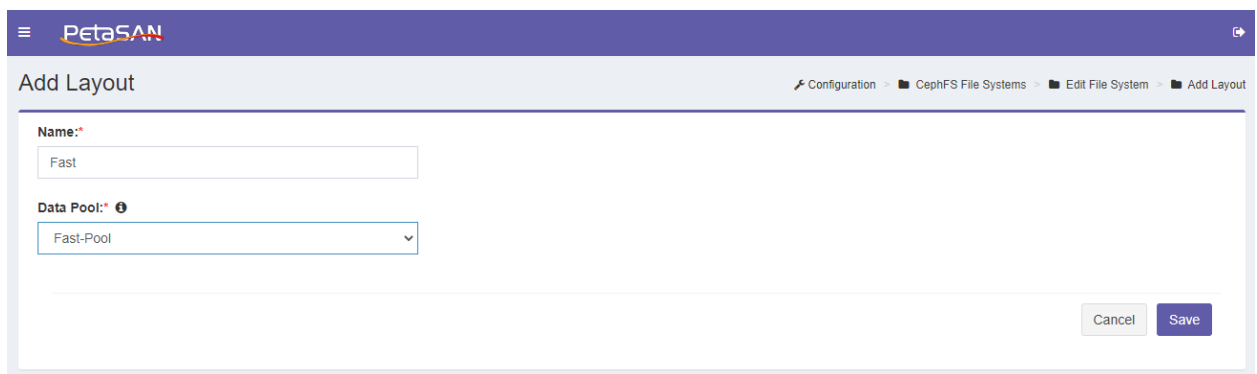
Enables you to delete the selected layout including its data pool .



### ➤ Add layout

You can add a new layout under the existing file system; the system requires entering the following fields:

- **Name**  
The layout name
- **Data Pool**  
Select the default layout data pool from list of existing pools of type cephfs.
- **Actions**
  - Save  
Saves the entered layout.

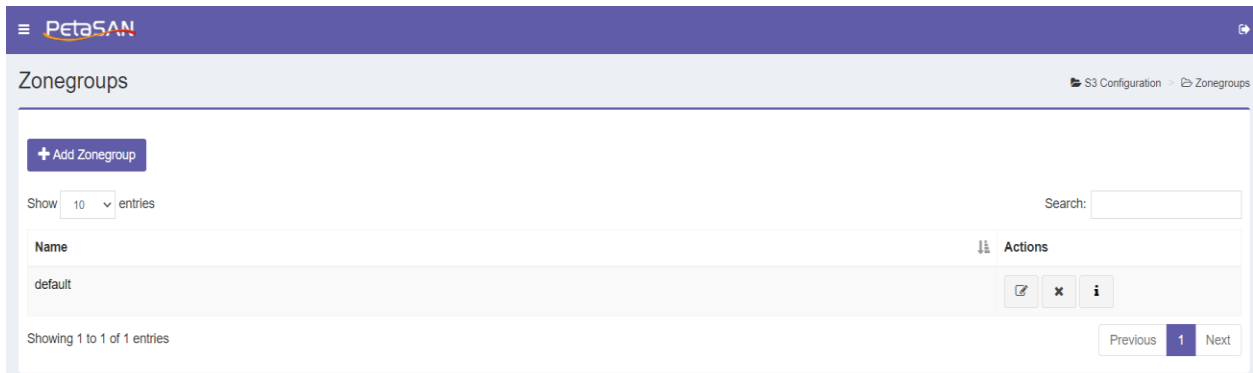


## 4.8. S3 Configuration

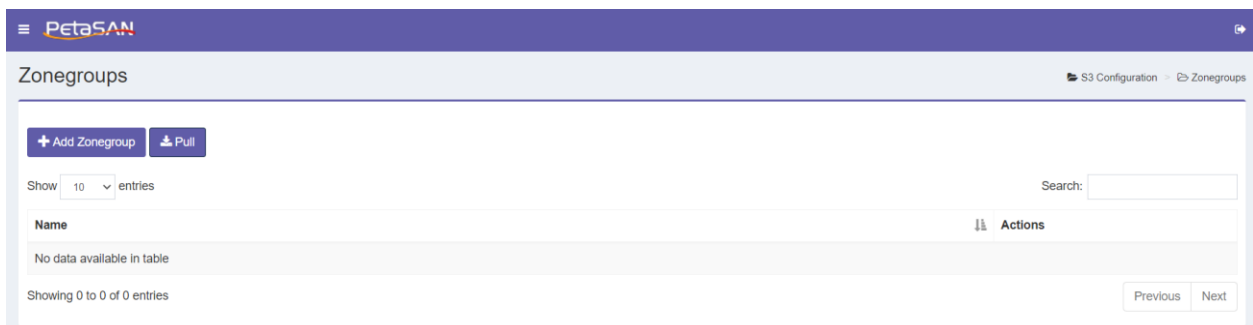
PetaSAN supports multi zones but with only one Realm and one zonegroup, If you have any other setup done manually the system will show an error message that current setup in not supported.

### ➤ View Zonegroups

System views list of zonegroups, for a supported system only one zonegroup should exist



In case the site is not master site, system will view the Pull button as follows



System views the following columns for the zonegroup.

- **Name**  
The zonegroup name
- **Actions**
  - Edit  
Edits the selected zonegroup, the name and rgw.root pool information can't be modified
  - Delete  
Deletes the selected zonegroup
  - Info  
Views the zonegroup information.
  - Pull  
Opens the Pull form which pulls the master zone configuration .

## ➤ Pull

You can get the master zone configuration in case you are creating a secondary site, system will open the Pull form which contains the following fields:

The screenshot shows a 'Pull' dialog box with the following fields and values:

- Master Zone Endpoint URL:\***
- Master Zone Access Key:\***
- Master Zone Secret Key:\***

Buttons: Cancel, Pull

- **Master Zone URL**
  - You need to specify the master zone's URL
- **Master Zone Access Key**
  - You need to specify the system user key for the master zone.
- **Master Zone Secret Key**
  - You need to specify the access key for the master zone.
- **Actions**
  - Pull
    - Pulls the master zone's configuration.
  - Cancel
    - Cancels the operation

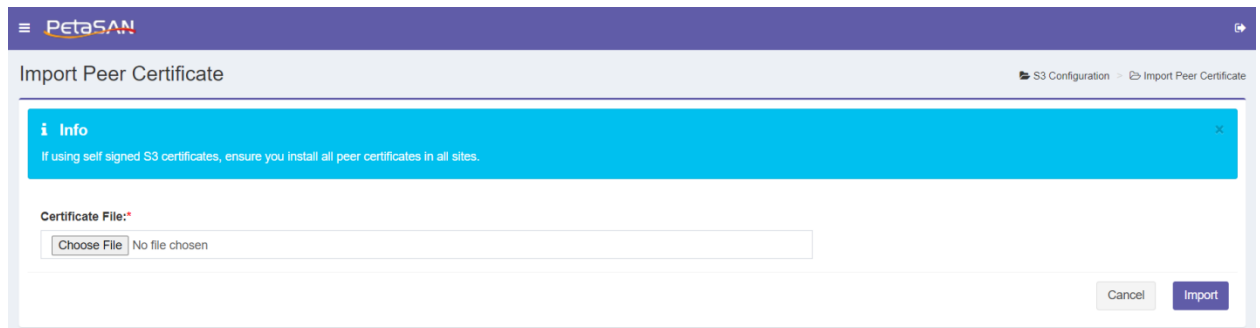
In Case you are using self signed certificate, import the peer certificate of the master zone before trying to pull the master zone configuration.

## ➤ Import Peer Certificate

In case you are using https using self signed certificates, you can import certificate of peer sites you will connect to; the system requires uploading the peer certificate .cert file:

- **Certificate File**

Upload the peer certificate .cert file



- **Actions**

- Imports  
Imports the peer certificate.
- Cancel  
Cancels the import.

## ➤ Add Zonegroup

You can add a new zonegroup; the system requires entering the following fields:

- **Name**

The zonegroup name

- **.rgw.root pool**

You need to specify the rgw.root pool's no of PGs, placement rule and size.

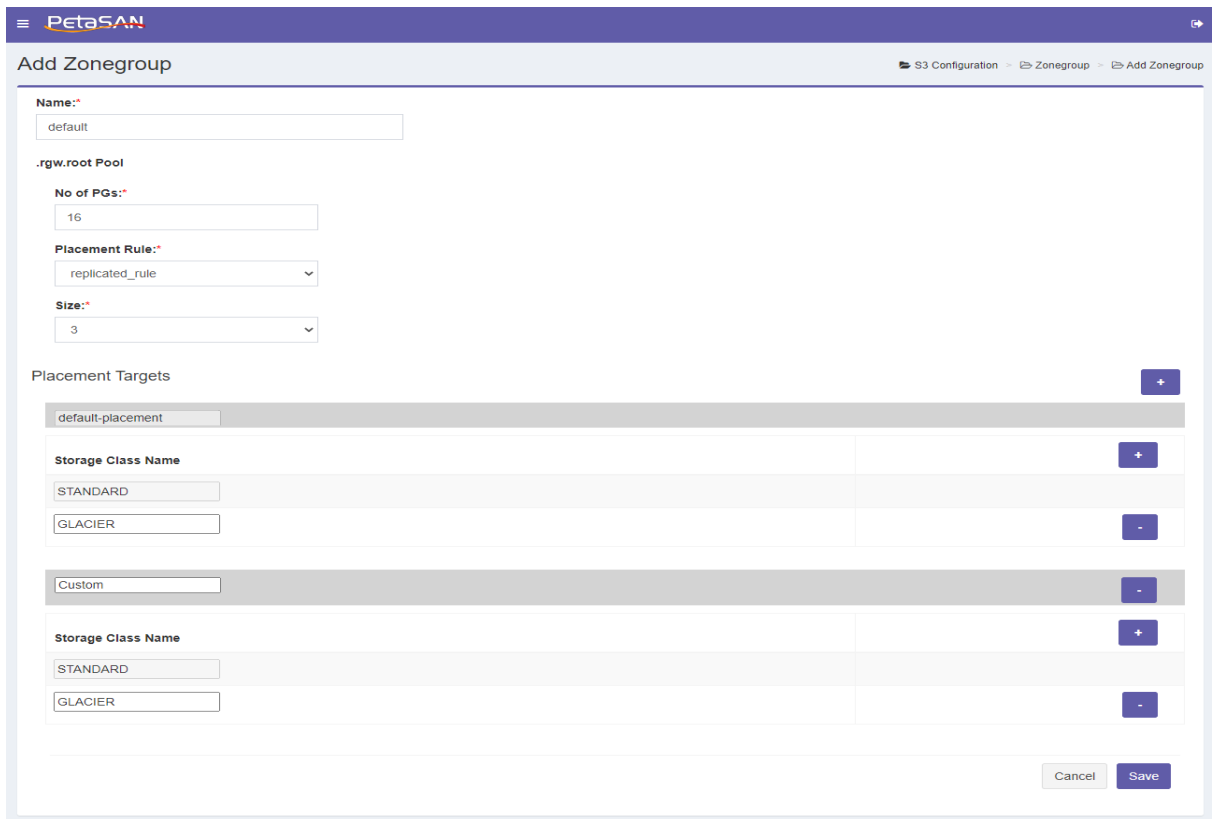
- **Placement targets**

By default you will have the default-placement target with STANDARD storage class and they can't be deleted, but you can add more placement targets or storage classes.

- **Actions**

- Save  
Saves the entered zonegroup.
- Cancel  
Cancel and returns to the zonegroup view list.
- Add Placement Target  
Adds new placement target other than the default-placement
- Add Storage Class

Adds new storage class under one of the existing placement targets



**Add Zonegroup**

Name: default

.rgw.root Pool

No of PGs: 16

Placement Rule: replicated\_rule

Size: 3

Placement Targets

Storage Class Name	
STANDARD	+
GLACIER	-
Custom	-
STANDARD	+
GLACIER	-

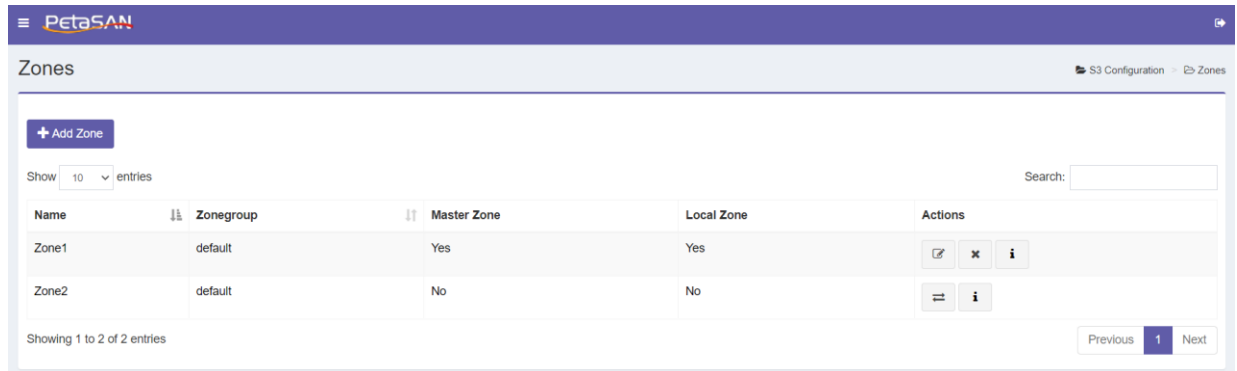
Cancel Save

➤ **Note:**

- You must enter the S3 Settings and make sure at least one node has the S3 role before adding new zonegroup.
- You may find .rgw.root pool created before adding a zonegroup, this is because ceph creates it if any rgw command has been run, when you add a new zonegroup system will delete the existing one and creates it with you selected specification.
- You can add only one zonegroup (as this is the supported setup)

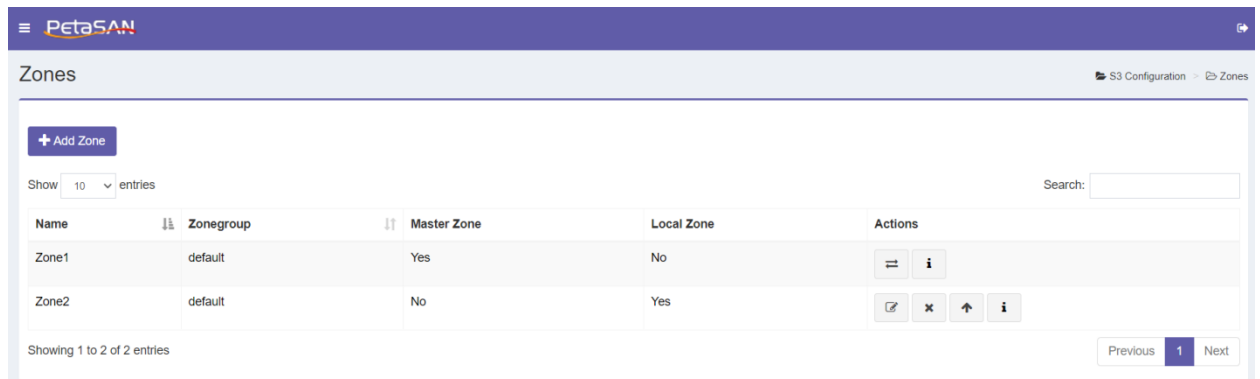
## ➤ Zones View List

In case of the master zone, system views list of zones as follows



Name	Zonegroup	Master Zone	Local Zone	Actions
Zone1	default	Yes	Yes	
Zone2	default	No	No	

And in case of non master zone, system will view the zones list as follows:



Name	Zonegroup	Master Zone	Local Zone	Actions
Zone1	default	Yes	No	
Zone2	default	No	Yes	

System views the following columns for the zones

- **Name**  
The zone name
- **Zonegroup**  
The zonegroup name
- **Master Zone**  
Yes if the zone is a master zone
- **Local Zone**  
Yes if the zone is the local zone
- **Actions**
  - Edit  
Edits the selected local zone, the main pools pool information can't be modified
  - Delete  
Deletes the selected local zone and all its pools
  - Info  
Zone information
  - Edit Endpoints

Enables the administrator to update the endpoint of a remote site

- Promote Zone

Promotes a non master to a master zone

Note: promoting a zone while not all metadata has been synched to it yet, may cause losing these metadata.

### ➤ **Add Zone**

You can add a new zone; the system requires entering the following fields:

- **Name**

The zone name

- **Main Pools**

The system shows the default setup of the control, meta and log pools, you are able to modify this information by selecting to Modify Main Pools.

- **Placement Pools**

By default, you have to specify a bucket index pool for the default-placement and the bucket data pool for the STANDARD storage class bucket data pool. The bucket index pool should be specified on a fast SSD storage pool.

You can add bucket index pool for any added placement and bucket data pool for each added storage class.

In case you selected an EC Rule for the Bucket Data Pool, the system will ask you to select the EC Profile and in this case you will have to select the Data Extra Pool information to be able to create the Buckets.non-ec pool which is required for proper functionality

- **Actions**

- Save

Saves the entered zone.

- Cancel

Cancel and returns to the zone view list.

- Add Placement Target

Enables the user to select placement target other than the default-placement from list of defined placement targets added in the zone group and create its Buckets Index Pool.

- Add Storage Class

Enables the user to select Storage Class other than the STANDARD class from list of storage classes defined in the zone group and creates its Buckets Data Pool.



☰ PetaSAN
➔

Add Zone
Manage Zones > Zones > Add Zone

**Zone Group Name:**  
default

**Zone Name:**

**End Points:**

**Main Pools** Modify Main Pools

Function	Pool Name	PGs	Rule Name	Size
Control	<input type="text" value="Zone1.rgw.control"/>	<input type="text" value="16"/>	<input type="text" value="replicated_rule"/>	<input type="text" value="3"/>
Meta	<input type="text" value="Zone1.rgw.meta"/>	<input type="text" value="16"/>	<input type="text" value="replicated_rule"/>	<input type="text" value="3"/>
Log	<input type="text" value="Zone1.rgw.log"/>	<input type="text" value="16"/>	<input type="text" value="replicated_rule"/>	<input type="text" value="3"/>

**Default Placement** +

**Buckets Index Pool:**

Function	Pool Name	PGs	Rule Name	Size
Buckets Index	<input type="text" value="Zone1.rgw.buckets.index"/>	<input type="text" value="16"/>	<input type="text" value="replicated_rule"/>	<input type="text" value="3"/>

**Storage Classes Buckets Data Pools:** +

Storage Class	Pool Name	PGs	Rule Name	Size	Action
STANDARD	<input type="text" value="Zone1.rgw.buckets.data"/>	<input type="text" value="64"/>	<input type="text" value="replicated_rule"/>	<input type="text" value="3"/>	
GLACIER	<input type="text" value="Zone1.rgw.glacier.buckets.data"/>	<input type="text" value="64"/>	<input type="text" value="replicated_rule"/>	<input type="text" value="3"/>	-

Custom Placement -

**Buckets Index Pool:**

Function	Pool Name	PGs	Rule Name	Size
Buckets Index	<input type="text" value="Zone1.rgw.custom.buckets.index"/>	<input type="text" value="64"/>	<input type="text" value="replicated_rule"/>	<input type="text" value="3"/>

**Storage Classes Buckets Data Pools:** +

Storage Class	Pool Name	PGs	Rule Name	Size	Action
STANDARD	<input type="text" value="Zone1.rgw.custom.buckets.data"/>	<input type="text" value="64"/>	<input type="text" value="replicated_rule"/>	<input type="text" value="3"/>	
GLACIER	<input type="text" value="Zone1.rgw.custom.glacier.buckets.data"/>	<input type="text" value="64"/>	<input type="text" value="replicated_rule"/>	<input type="text" value="3"/>	-

Cancel Save

In case of adding a secondary zone you will enter the zone information as in the master zone plus the access and secret keys of the master zone

PetaSAN
Manage Zones > Zones > Add Zone

### Add Zone

**Zonegroup Name:**  
default

**Zone Name:**

**Master Zone Access Key:**

**Master Zone Secret Key:**

**End Points: (For Multisite)**

**Main Pools** Modify Main Pools

Function	Pool Name	PGs Autoscale	Rule Name	Size
Control	<input type="text" value="Zone2.rgw.control"/>	<input type="text" value="on"/>	<input type="text" value="replicated_rule"/>	<input type="text" value="3"/>
Meta	<input type="text" value="Zone2.rgw.meta"/>	<input type="text" value="on"/>	<input type="text" value="replicated_rule"/>	<input type="text" value="3"/>
Log	<input type="text" value="Zone2.rgw.log"/>	<input type="text" value="on"/>	<input type="text" value="replicated_rule"/>	<input type="text" value="3"/>

**Placement Targets** +

Default Placement

**Buckets Index Pool:**

Function	Pool Name	PGs Autoscale	Rule Name	Size
Buckets Index	<input type="text" value="Zone2.rgw.buckets.index"/>	<input type="text" value="on"/>	<input type="text" value="replicated_rule"/>	<input type="text" value="3"/>

**Storage Classes Buckets Data Pools:** +

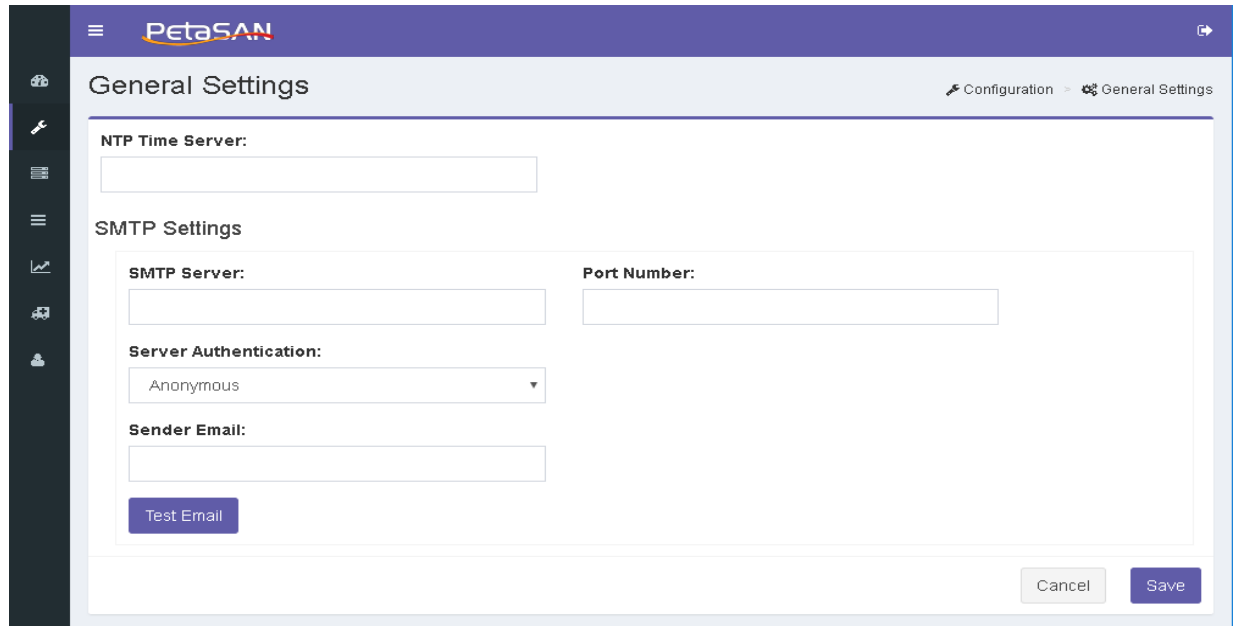
Storage Class	Pool Name	PGs Autoscale	Rule Name	Size	Action
STANDARD	<input type="text" value="Zone2.rgw.buckets.data"/>	<input type="text" value="on"/>	<input type="text" value="replicated_rule"/>	<input type="text" value="3"/>	

Cancel
Save

## 4.9. General Settings

### ➤ Update General Settings

You should use the General settings form to update the NTP Time Server and SMTP server settings.



The form contains the following fields:

- **NTP Time Server**

You should enter the time server that you want to synch your cluster time with. If not specified, PetaSAN will sync the cluster to the local machine clocks of the management nodes.

- **SMTP Settings**

You should enter the SMTP server settings if you want PetaSAN to send email notifications. You should enter the following fields:

- SMTP server name
- Port Number
- Server Authentication

You can select one of following the server authentication options:

- Anonymous
- Basic Authentication
- SSL
- TLS
- Sender Email
- Password of the sender email

### ➤ Notes:

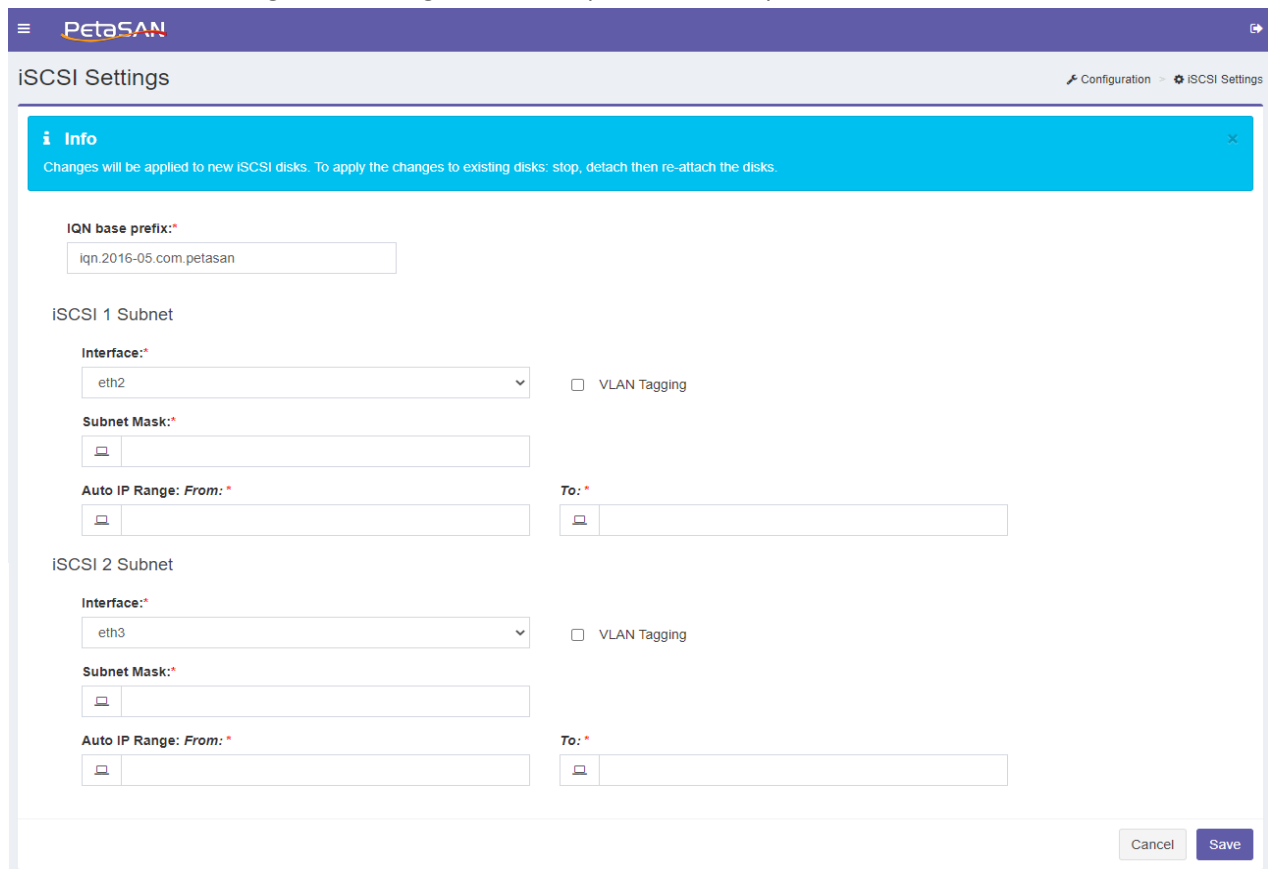
- *Update the profiles of the users who should receive email notification by adding their email addresses and checking the Receive Notifications option.*

- You can test the SMTP server settings by sending a test email to all users registered to receive Notifications.
- PetaSAN sends notifications in the following cases
  - Node failure.
  - OSD failure.
  - Raw storage approaching full capacity
  - Replication job failure

## 4.10. iSCSI Settings

### ➤ Update iSCSI Settings

The iSCSI settings form configures various parameters required for the creation of iSCSI disks.



The form contains the following fields:

- **IQN base/prefix**
  - The prefix entered will be used as the base for all disks IQNs that will be created. The disk id will be appended to the base prefix to form the unique IQN associated with the disk.
  - The IQN format takes the form `iqn.yyyy-mm.naming-authority:unique name`, where:
    - `yyyy-mm` is the year and month when the naming authority was established.

- Naming-authority is usually reverse syntax of the Internet domain name of the naming authority. For example, the `iscsi.vmware.com` naming authority would have the iSCSI qualified name form of `iqn.1998-01.com.vmware.iscsi`. The name indicates that the `vmware.com` domain name was registered in January of 1998, and iSCSI is a subdomain, maintained by `vmware.com`.
- Unique name is any name you want to use, and in PetaSAN we are using the disk id.

For example if the base IQN is `iqn.2016-05.com.petasan` then the IQN for disk 00001 will be `iqn.2016-05.com.petasan:00001`

## iSCSI Subnets

You should specify the settings for the iSCSI subnet1 and iSCSI subnet 2 by entering the following information:

- **Interface**

The interface card assigned to the iSCSI subnet, it can be updated at any time as long as all existing nodes with the iSCSI role have the selected interface.

- **Subnet Mask**

Enter the subnet mask of the iSCSI subnet

- **VLAN Tagging**

Enable VLAN tagging for the iSCSI subnet, if enabled then you must enter the subnet VLAN Id

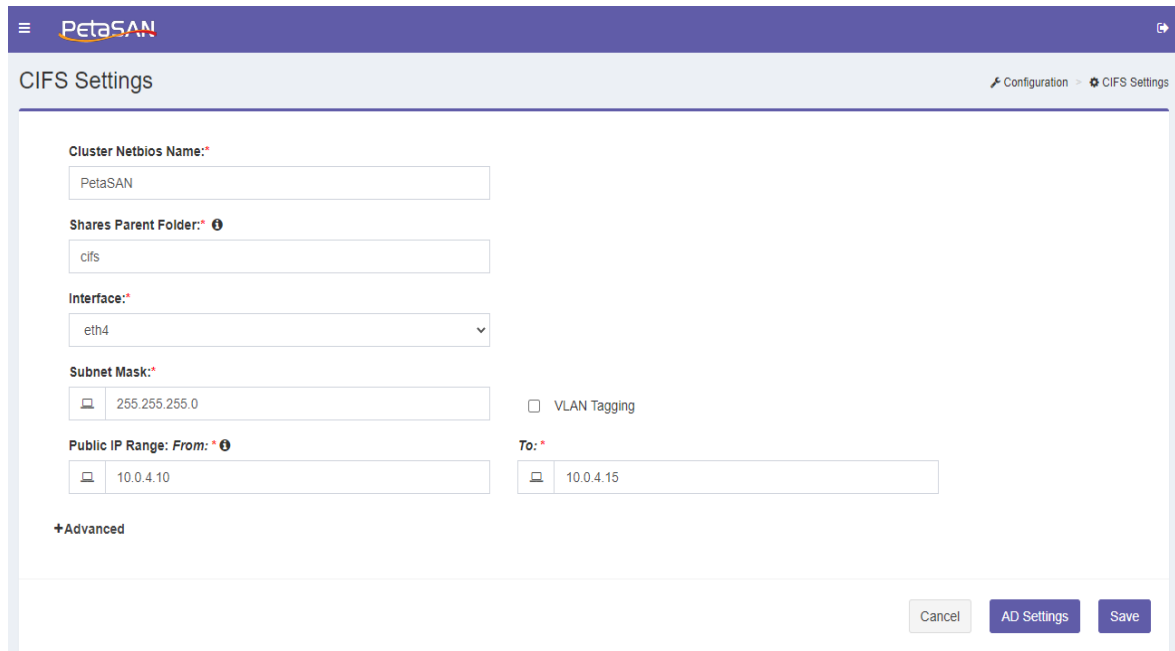
- **Auto IP Range**

Enter the start and end auto IP Range which define a pool of IPs that are available for assignment when creating iSCSI disks with the Auto assign IP option.

## 4.11. CIFS Settings

### ➤ Update CIFS Settings

The CIFS Settings include configuration required to set up CIFS services, networking as well as Active Directory security.



The CIFS settings form allows specifying the following fields

- **Cluster Netbios Name**  
The Netbios name for the CIFS cluster which is set to PetaSAN by default
- **Shares parent folder**  
The top level folder within a pool or layout in which all CIFS shares will be created.
- **Interface**  
Interface for the CIFS subnet ,it can be changed any time as long as all nodes hosting the CIFS role have the selected interface card.
- **Subnet Mask**  
The subnet mask of the CIFS public subnet used by clients.
- **VLAN Tagging**  
Enable VLAN tagging for the CIFS subnet, if enabled then you must enter the subnet VLAN Id
- **Public IP Range**  
Enter the start and end public IP which define a range of IPs allocated on the CIFS servers which clients connect to. These are dynamic IPs that can be re-assigned to support high availability.
- **Gateway**  
By default use the default gateway on the management network else specify the gateway IP on public subnet.
- **Advanced**

**smb.conf**

- Global Public

- Samba daemon global section for public shares.
- Global Secure
  - Samba daemon global section for secure shares using AD.
- Share Public
  - Samba daemon public shares section.
- Share Secure:
  - Samba daemon protected shares section

## krb5.conf

- krb5
  - Configuration for kerberos support libraries.
- **Actions**
  - You can select one of the following actions
  - Save
    - Saves the current settings
  - AD Settings
    - Opens the Active Directory Settings Form

## ➤ Active Directory Settings

You can add the active directory settings

AD Settings

**i Info** Please make sure time gap between server and client nodes does not exceed 5 minutes

**AD Domain Name:\***  
company.local

**AD IP:\***  
192.168.123.7

Close Join

You will need to enter the Active Directory settings by entering the following fields:

- **AD Domain Name**
  - The name of the active directory domain you want to join
- **AD DNS IP**

The Active Directory DNS server IP address

- **Actions**

- Join

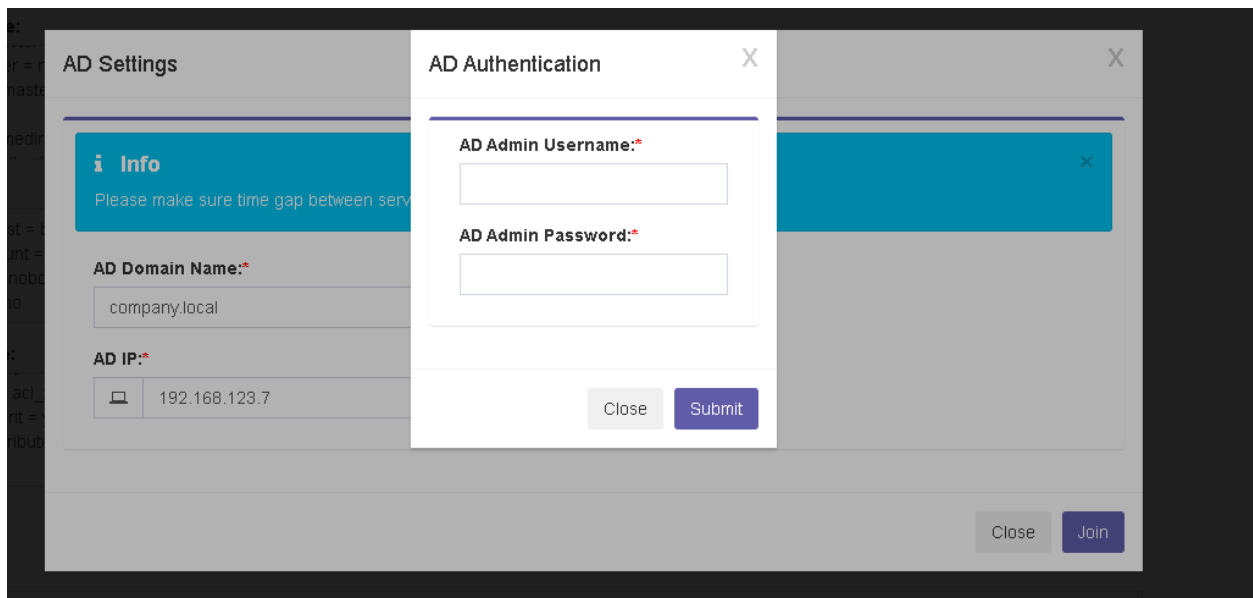
- After entering the AD information, you should select to join; the system will ask you to enter the AD user name and password

- Unjoin

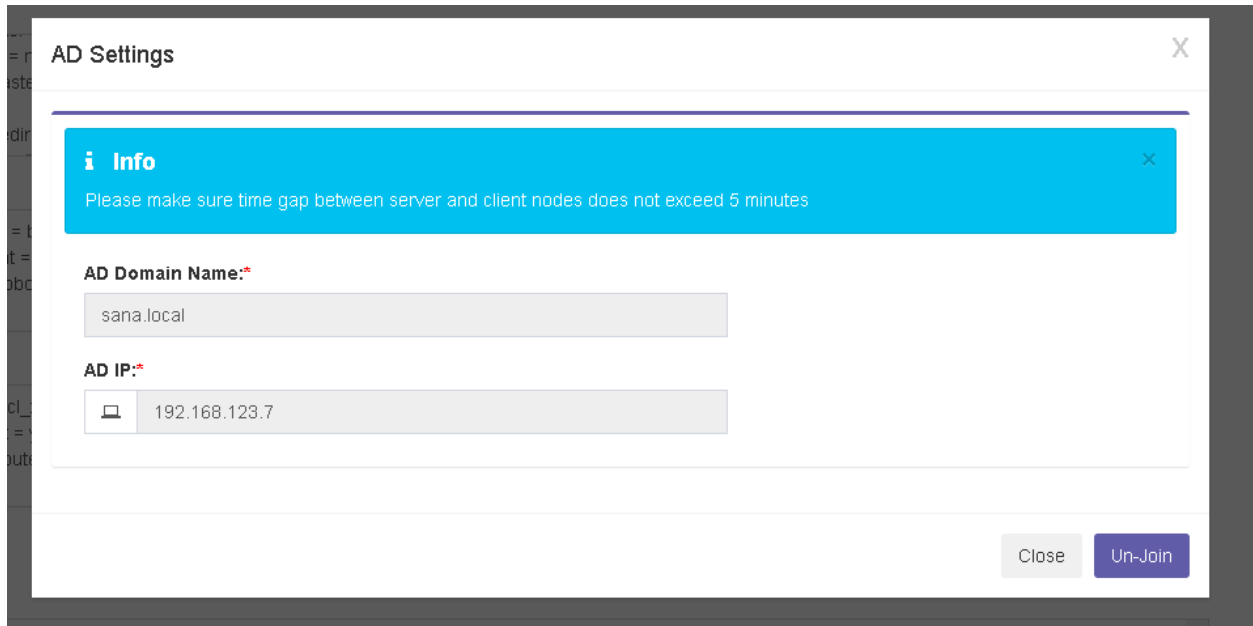
- If you already joined an active directory, you can unjoin it, this allows you to join another active directory server

➤ **Notes:**

*The system will not save the active directory password; it will be only used to during the join operation.*



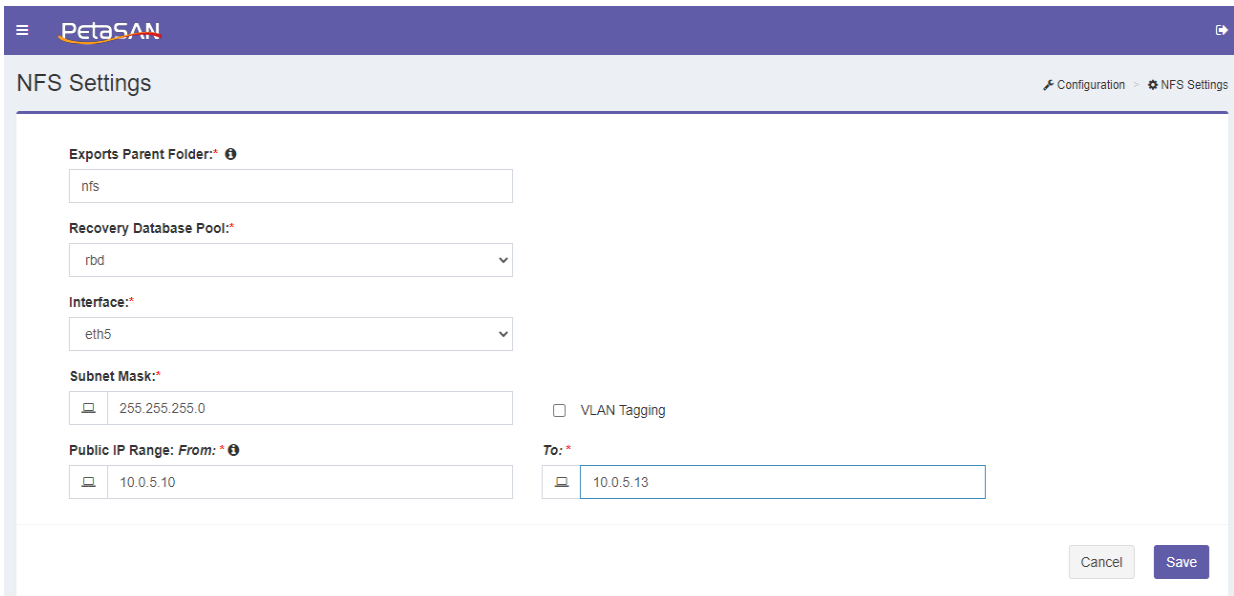




## 4.12. NFS Settings

### ➤ Update NFS Settings

The NFS Settings include configuration required to set up NFS services.



The NFS settings form allows specifying the following fields

- **Exports parent folder**

The top level folder within a pool or layout in which all NFS exports will be created.

- **Interface**

Interface for the NFS subnet ,it can be changes any time as long as all nodes hosting the NFS role has the selected Ethernet card.

- **Subnet Mask**

The subnet mask of the NFS public subnet used by clients.

- **VLAN Tagging**

Enable VLAN tagging for the NFS subnet, if enabled then you must enter the subnet VLAN Id

- **Public IP Range**

Enter the start and end public IP which define a range of IPs allocated on the NFS servers which clients connect to. These are dynamic IPs that can be re-assigned to support high availability.

- **Gateway**

- By default uses the default gateway on the management network else specify the gateway IP on public subnet.

## 4.13.S3 Settings

### ➤ Update S3 Settings

The S3 Settings include configuration required to set up S3 server.

The screenshot shows the PetaSAN S3 Settings configuration page. The page has a purple header with the PetaSAN logo and a navigation menu. The main content area is titled "S3 Settings" and contains several configuration sections:

- RadosGW http port:** A text input field containing "7480".
- Load Balancer:**
  - Port:** A text input field containing "8000".
  - HTTPS:** A checked checkbox.
  - Private Key:** A large text area for the private key, with a "Generate" button to its right.
  - Certificate:** A large text area for the certificate.
- Interface:** A dropdown menu showing "eth6".
- Subnet Mask:** A text input field containing "255.255.255.0".
- VLAN Tagging:** An unchecked checkbox.
- Public IP Range:** Two text input fields: "From:" containing "10.0.7.100" and "To:" containing "10.0.7.110".
- Gateway:** Two radio buttons: "Default" (selected) and "Custom".

At the bottom right of the page, there are "Cancel" and "Save" buttons.

The S3 settings form allows specifying the following fields

- **RadosGW http port**

The backend radosgw port, default is 7480. **Load Balancer Port**

The load balancer port

- **HTTPS**

Select if you want to enable https support.

If https is selected you will need either enter a key and certificate or generate them by clicking on generate button

System will ask you about the common name to use for certificate.

Using the common name system will generate the certificate

- **Interface**

Interface for the S3 Load Balancer subnet ,it can be changed any time as long as all nodes hosting the S3 role have the selected Ethernet card.

- **Subnet Mask**

The subnet mask of the S3 public subnet used by clients.

- **VLAN Tagging**

Enable VLAN tagging for the S3 subnet, if enabled then you must enter the subnet VLAN Id

- **Public IP Range**

Enter the start and end public IP which define a range of IPs allocated on the S3 servers which clients connect to. These are dynamic IPs that can be re-assigned to support high availability.

- **Gateway**

By default uses the default gateway on the management network else specify the gateway IP on public subnet.

Once you saved the settings, you will be able to download the certificate.

PetaSAN
⌵

## S3 Settings

Configuration > S3 Settings

**RadosGW http port:**

**Load Balancer**

**Port:**

**HTTPS**

**Private Key:**

```
-----BEGIN PRIVATE KEY-----
MIIJRAIBADANBgkqhkiG9w0BAQEFAASCCS4wggkqAgEAAoICAQC1dVSq0/RTvYa1
qX/SASvKTP1f2tG0RFr3bgDuxKpbNKh8a3Z1WnioaN/ZgPQFCcnvXe3TogXHObJi
+of6KFqRT7V7HQG+qtid6Jrpxg+idPzqQp95uROCOgI3V1dX4o8xMiRNLb2+i1s
GMOEARhYz17B4LEVH8qb5rrdQg3bmK85ozpRzIsLHQQ/Xga13qMUx08H/Ko8Q1G7
Jwy51tS8+P1yhkhKFn7FjhvBt/dNLU8c7zcQGDaH0Rrzc7Lc1D7pqfCtnw/7Xo1V
.....
```

**Certificate:**

```
-----BEGIN CERTIFICATE-----
MIIEVzCCAqegAwIBAgIUHoZJfcN1LK68x8FqNlqMfU5n1BcwDQYJKoZIhvcNAQEL
BQAwDzENMAsGA1UEAwwEZGVtZzAeFw0yMTA2MDIuNTE3MjNaFw00MTA1MjcNTE3
HjNaMA8xDALBglVBAMMBGR1bW8wgI1HA0GCSqGSIb3DQEBAQUAA4ICDwAwggIK
AoICAQC1dVSq0/RTvYa1qX/SASvKTP1f2tG0RFr3bgDuxKpbNKh8a3Z1WnioaN/Z
gPQFCcnvXe3TogXHObJi+of6KFqRT7V7HQG+qtid6Jrpxg+idPzqQp95uROCOgI
3V1dX4o8xMiRNLb2+i1sGMOEARhYz17B4LEVH8qb5rrdQg3bmK85ozpRzIsLHQQ/
.....
```

**Interface:**

**Subnet Mask:**

**VLAN Tagging**

**Public IP Range: From:**

**To:**

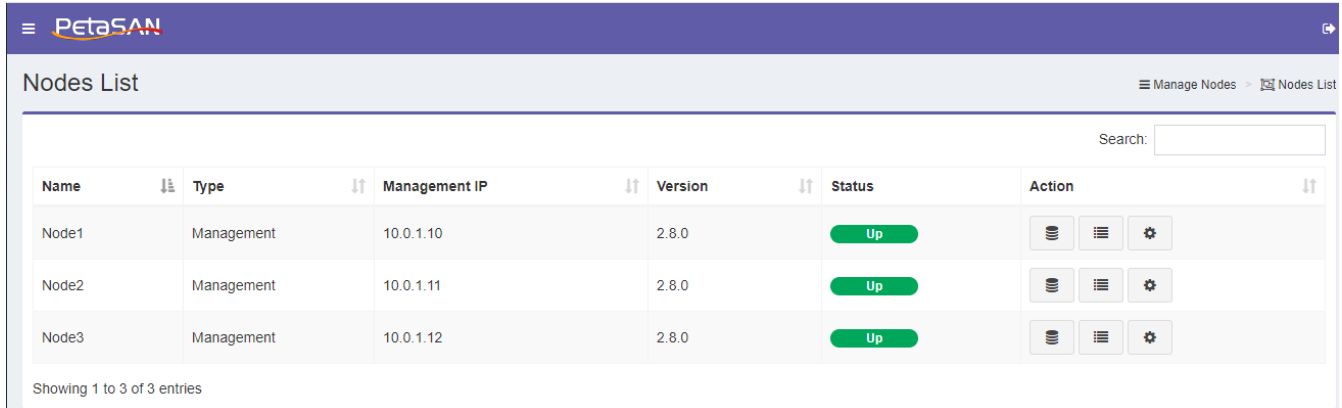
**Gateway:**

Default     Custom

## 5. Manage Nodes

### ➤ Node List

The system views the list of cluster nodes and enables you to manage them.



Name	Type	Management IP	Version	Status	Action
Node1	Management	10.0.1.10	2.8.0	Up	View Local disks, View Log, Manage Roles
Node2	Management	10.0.1.11	2.8.0	Up	View Local disks, View Log, Manage Roles
Node3	Management	10.0.1.12	2.8.0	Up	View Local disks, View Log, Manage Roles

Showing 1 to 3 of 3 entries

The node list views the following information about each node:

- **Name**  
The node host name
- **Type**  
The node type which can be either “Management” or “Storage”
- **Management IP**  
The node management IP addresses
- **Version**  
The version number of the node
- **Status**  
The node status whether it is up or down
- **Action**  
The system displays the actions that can be done on the node according to its type and status which are:
  - View Local disks  
Shows the list of local disk for the node and enables you to manage them refer to “Node Local Disks”.  
  
The button will be available for running nodes only
  - View Log  
Shows the system log for the node  
The button will be available for running nodes only
  - Manage Roles  
Opens the Manage Roles form, refer to “Manage Roles”.  
The button will be available for running nodes only
    - Delete

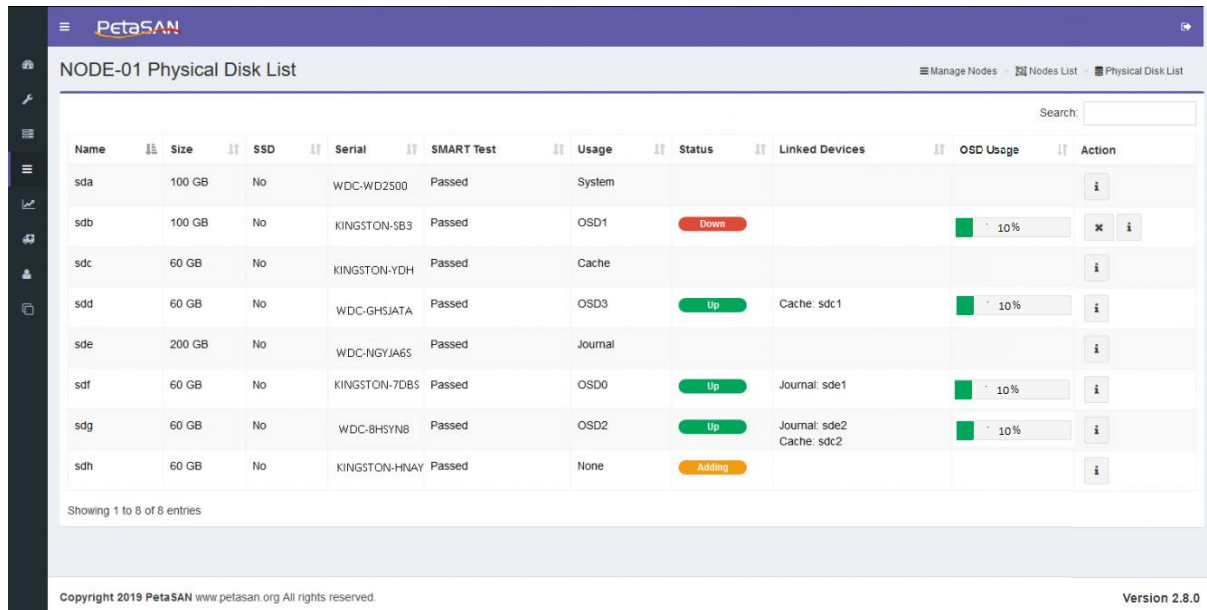
Enables you to delete a down storage node which will remove the node from the cluster and all its OSDs.

The button will be available for nodes of type “Storage” and only when they are down.

- *Note: If you have a down storage node, you can still re-use its OSDs by moving them to other nodes before node deletion.*

## ➤ Node Physical Disk List

The system displays a list of local disks for a specific node.



Name	Size	SSD	Serial	SMART Test	Usage	Status	Linked Devices	OSD Usage	Action
sda	100 GB	No	WDC-WD2500	Passed	System				<i>i</i>
sdb	100 GB	No	KINGSTON-SB3	Passed	OSD1	Down		10%	<i>x</i> <i>i</i>
sdc	60 GB	No	KINGSTON-YDH	Passed	Cache				<i>i</i>
sdd	60 GB	No	WDC-GHSJATA	Passed	OSD3	Up	Cache: sdc1	10%	<i>i</i>
sde	200 GB	No	WDC-NGYJA6S	Passed	Journal				<i>i</i>
sdf	60 GB	No	KINGSTON-7DBS	Passed	OSD0	Up	Journal: sde1	10%	<i>i</i>
sdg	60 GB	No	WDC-BHSYNB	Passed	OSD2	Up	Journal: sde2 Cache: sdc2	10%	<i>i</i>
sdh	60 GB	No	KINGSTON-HNAY	Passed	None	Adding			<i>i</i>

The node local disks list views the following information about each disk:

- **Name**  
The local disk name
- **Size**  
The disk size
- **SSD**  
This column specifies if the disk is SSD or not
- **Serial**  
The disk serial number
- **SMART Test**  
The result of overall SMART health test for the disk, test results are Passed, Failed or Unknown
- **Usage**  
This column shows if the disk is used in the cluster and what its usage is.  
The disk usage can be one of the following values:
  - System

This means the disk is used as a system disk.

- Journal

This means the disk is used as a Journal WAL/DB

- OSD

This means the disk is used as a Ceph OSD storage disk. The system displays its OSD number.

- No

This means the disk is not used in the cluster.

- Mounted

This means the disk is not used in the cluster but it is mounted.

- **Linked Devices**

This column shows the journal disk or cache disk used by the OSD if any.

- **OSD Usage**

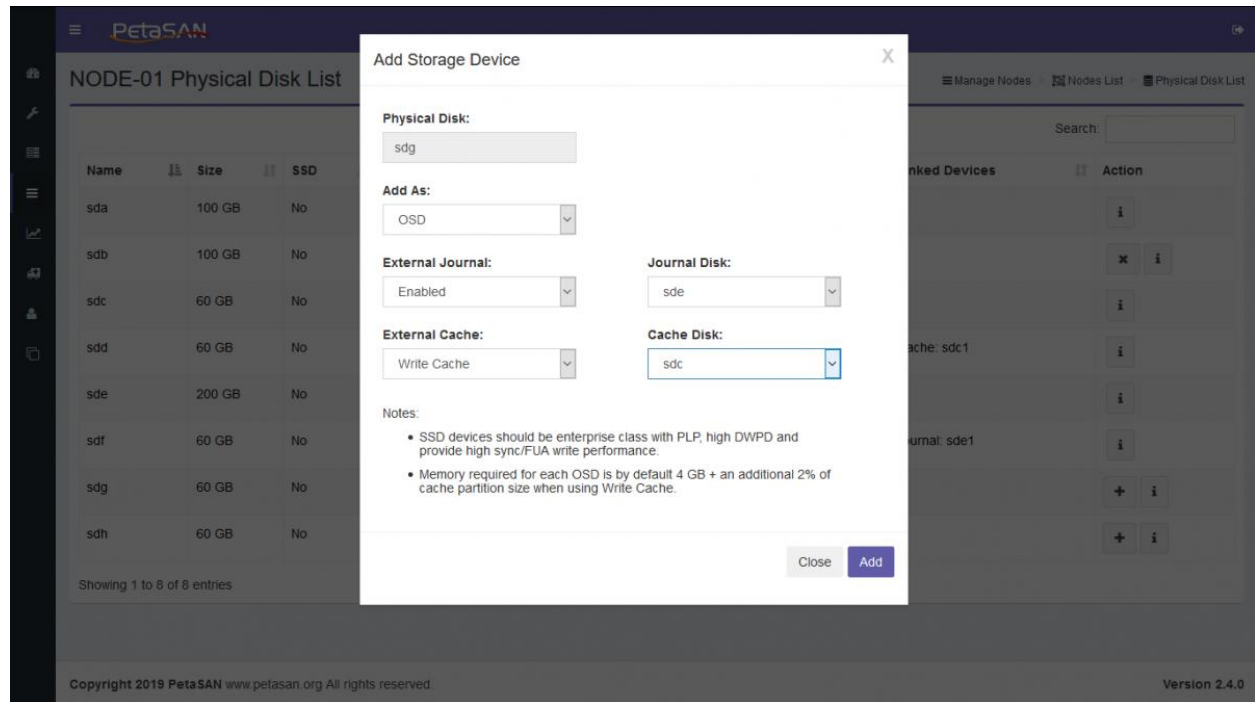
Percentage of used space in the OSD.

- **Action**

The system displays the actions that can be done on the disk according to its type and Status which are:

- Add Storage Device

Opens that Add Storage Device form .



The system enables you to add the disk as

- OSD

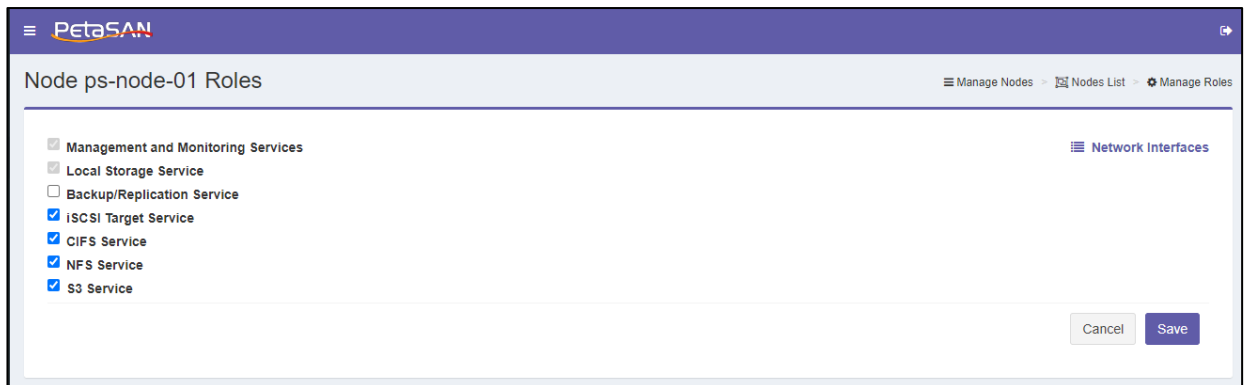
Adds the local disk as an OSD, it is available only if the disk is not used or mounted

You can specify an external journal disk or Cache disk or both. You can manually



select the specific journal and cache devices or let the system select the available device by selecting Auto.

- Journal
  - Adds the local disk as Journal, it is available only if the disk is not used or mounted  
Cache  
Adds the local disk as Cache, it is available only if the disk is not used or mounted
  - Delete  
Deletes the local disk from the cluster, it is available only if the disk is not a system disk and is either a down OSD or an unused cache or journal..
  - Info  
Views the storage device full information
- **Notes:**
- You will not be able to add disk as either OSD, cache or Journal if its node does not have Local Storage Service assigned to its roles, you will first need to update the node roles using the Manage Roles form.
- **Manage Roles**  
The system shows the roles currently assigned to the node and enables you to add more roles or remove existing roles using the Manage Roles Form.



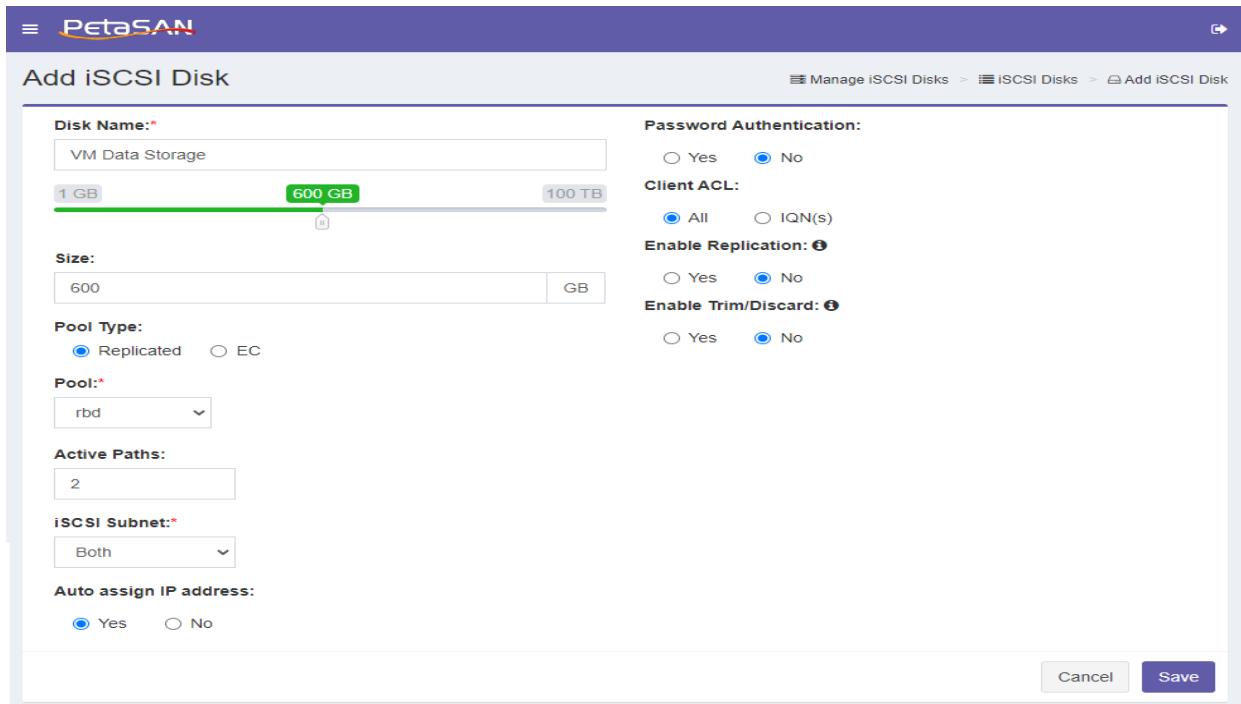


## 6. Manage iSCSI Disks

### 6.1. iSCSI Disks

#### ➤ Add iSCSI Disk

By clicking on iSCSI Disk List->Add iSCSI Disk the system will open a form that you must fill in all its required fields then click Save



The Add iSCSI disk form contains the following fields:

- **Disk Name**  
The disk name will make it easier to identify the disk either in the view list or while performing search. Disk name must be unique
- **Disk Size**  
To enter the disk size you can use the slider or manually enter the size
- **Pool Type**  
Select the desired pool type, replicated or EC
- **Pool**  
The system will show the list of pools which have their usage set to rbd.
  - In case replicated pool type is selected ,then you should select desired pool from the list of active replicated pools
  - In case of EC pool type is selected ,then you should select a replicate pool to store disk meta data and an EC pool to store the actual disk data
- **Active paths**  
You should select the number of active paths for the iSCSI disk.

➤ *Note: selecting more than one path load balances the io traffic across several nodes which helps improve performance.*

- **iSCSI subnet**

You should select whether you want to allocate the active paths on iSCSI Subnet1, iSCSI Subnet2 or on both.

- **Automatically assign IP address**

- By default the system will assign the IP address automatically using the Auto IP range already configured for the iSCSI subnets.
- If you select No then you will be able to assign the IP address(s) manually.
- Assigning IP addresses manually is only enabled for 1 or 2 active paths.
- You will not be able to use the “Auto Assign IP Address” feature in case there are any inactive pools in the cluster to avoid re-using already assigned IPs

- **Password Authentication**

- By default your disk will not require password authentication ,you can select yes to enable it
- If enabled you will have to enter user name and password

- **Client ACL**

You can allow any client to connect to the disk by selecting “All” or you can restrict access to some client IQNs

You should write the IQNs separated by comma “,”.

- **Enable Replication**

You can enable the disk to be a target of a replication job

- **Enable Trim/Discard**

Enables or Disables Trim/Discard functionality.

➤ **iSCSI Disk List**

You can view list of all disks using the iSCSI Disk List page

Disk Id	Size	Name	Created	Pool	IQN	Active Paths	Status	Action
00001	60GB	Backup	2023-06-21	rbd	iqn.2016-05.com.petasan:00001	2	Started	[Stop] [Edit] [Snapshot]
00003	50GB	Database	2023-08-18	Fast-VM_Pool	iqn.2016-05.com.petasan:00003	2	Started	[Stop] [Edit] [Snapshot]
00004	1GB	Test	2023-08-18	Backup_Pool	iqn.2016-05.com.petasan:00004	2	Stopped	[Play] [Edit] [X] [Refresh] [Snapshot]
00005	1GB	Test1	2023-08-18	Fast-VM_Pool		N/A	Detached	[Refresh] [X]
00006	10GB	Local-VMs	2023-08-18	Fast-VM_Pool	iqn.2016-05.com.petasan:00006	2	Starting	

The iSCSI Disk List displays the following information for each disk:

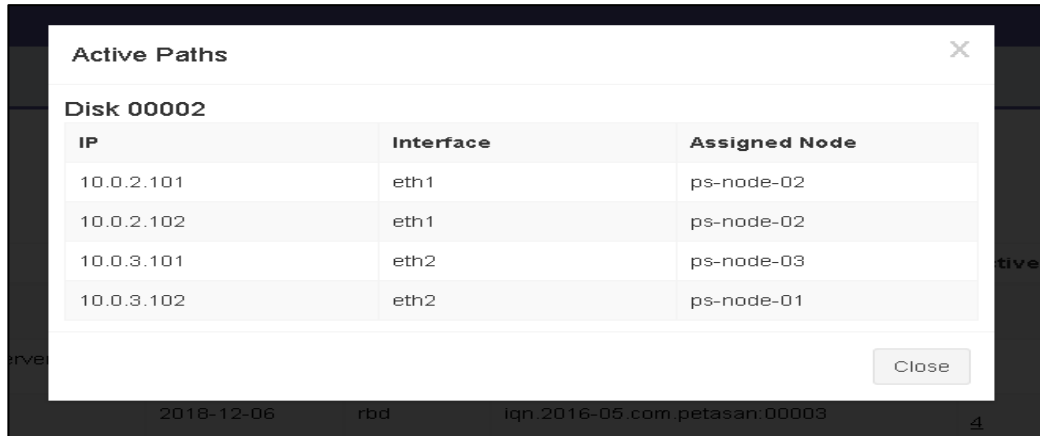
- **Replication information**

In case the disk is source or destination of a replication job, a special icon will appear which allows you to mouse over it and view information of the replication job.

Name	Created
Eman3	2019-04-15
Eman5	2019-04-15
ahmed_test	2019-04-15

- **Disk Id**  
The unique disk id assigned to the disk by the system
- **Size**  
The disk size
- **Name**  
The disk name.
- **Created**  
The date when the disk was created

- **Pool**  
The storage pool storing the disk image
- **IQN**  
The IQN assigned to the disk
- **Active Paths**  
The number of active paths selected for the disk  
By clicking on the number of paths the system will open a new page displaying the node hosting each path and its interface



Active Paths		
Disk 00002		
IP	Interface	Assigned Node
10.0.2.101	eth1	ps-node-02
10.0.2.102	eth1	ps-node-02
10.0.3.101	eth2	ps-node-03
10.0.3.102	eth2	ps-node-01

- **Status**  
Display the current status of the disk which can be one of the following:
  - Started  
At least one of the paths has been assigned and clients can connect to the disk.
  - Stopped  
All disk paths have been stopped and clients cannot connect to the disk.
  - Detached  
Disk image is available on Ceph but the iSCSI information has been removed from the image metadata.
  - Starting  
The system is starting the disk
  - Stopping  
The system is stopping the disk
- **Action**  
The column action displays the actions that can be done on each disk according to its current status; each action has a tool tip that is displayed when the mouse is over the action button.
  - Stop  
Stops the disk, it is available only when the disk is started.  
In case the disk is destination of a replication job, it will always be stopped after running the replication job.
  - Start

Starts the disk, it is available only when the disk is stopped.

In case the disk is destination of a replication job, any data you write to it will be lost after the next replication job.

- Detach

Removes the iSCSI information assigned to the disk but keeps the disk image on Ceph, it is available only when the disk is stopped.

- Attach

Opens the attach disk form so you can enter the iSCSI information associated with the disk, which is similar to information required when adding a new disk, it is available only when the disk is detached.

The screenshot shows the 'Attach iSCSI Disk' configuration page in the PetaSAN interface. The page has a dark sidebar on the left with navigation icons. The main content area is white with a purple header. The form contains the following fields and options:

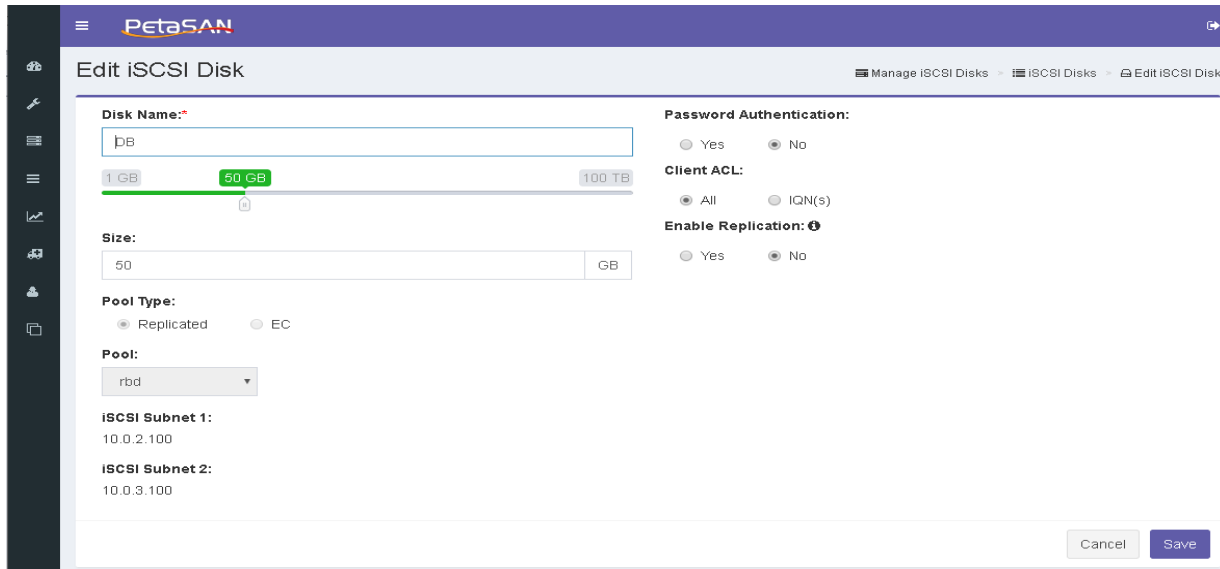
- Disk Name:** Text input field containing 'pb'.
- Size:** A slider bar ranging from 1 GB to 100 TB, with a green bar indicating the current size of 50 GB.
- Size:** Text input field containing '50' and a unit dropdown set to 'GB'.
- Pool Type:** Radio buttons for 'Replicated' (selected) and 'EC'.
- Pool:** A dropdown menu showing 'rbd'.
- Active Paths:** Text input field containing '2'.
- iSCSI Subnet:** A dropdown menu showing 'Both'.
- Auto assign IP address:** Radio buttons for 'Yes' (selected) and 'No'.
- Password Authentication:** Radio buttons for 'Yes' and 'No' (selected).
- Client ACL:** Radio buttons for 'All' (selected) and 'IQN(s)'.
- Enable Replication:** Radio buttons for 'Yes' and 'No' (selected).

At the bottom right of the form, there are 'Cancel' and 'Save' buttons.

- Edit

Opens the Edit Disk form which enables updating the disk name and/or size; it is available only when the disk is stopped.

➤ *Note: You will not be able to change the disk's pool type, pool or assigned IP addresses.*



- **Delete**  
Deletes the disk and all its data from the system and it is available only when the disk is stopped or detached.
- **View Snapshots**  
Views list of snapshots that were previously taken for the current iSCSI disk, snapshots may be taken manually or at a selected schedule (daily, weekly or monthly).

You can select number of scheduled snapshots the system should keep.

Disk 00001 Snapshots

X

Take Snapshot
Schedule Snapshots
View Schedule
Delete Schedule

X

Show 10 entries Search:

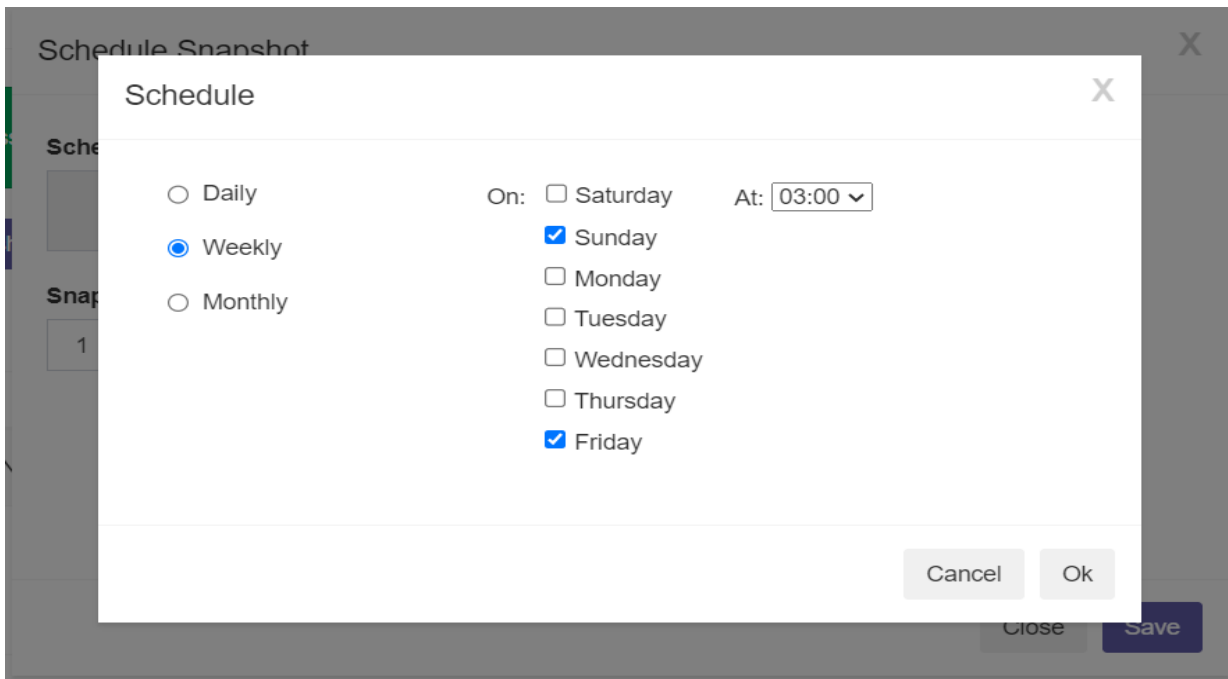
Time	Name	Type	Schedule	Exported	Status	Action
2023-08-18 04:00:01	PETASAN-S-20230818040001	Scheduled	Daily every 1 Hour	No		<span style="border: 1px solid #ccc; padding: 2px 5px;">✕</span> <span style="border: 1px solid #ccc; padding: 2px 5px;">↺</span> <span style="border: 1px solid #ccc; padding: 2px 5px;">↻</span>
2023-08-18 02:26:17	PETASAN-U-20230818022617	User		No		<span style="border: 1px solid #ccc; padding: 2px 5px;">✕</span> <span style="border: 1px solid #ccc; padding: 2px 5px;">↺</span> <span style="border: 1px solid #ccc; padding: 2px 5px;">↻</span>

Showing 1 to 2 of 2 entries 
Previous
1
Next

Close

You can have only one scheduled snapshot, when added, you will be able to view the saved schedule or delete schedule if needed.





For each snapshot taken either scheduled or taken by the user you can use one of the following actions:

-Delete

Deletes the selected snapshot.

-Rollback

Rolls back the current disk data to the selected snapshot.

-Export /Disable Export

Exports the selected snapshot so it can be accessed outside the system, system uses the last IP in the iSCSI 1 subnet IP range, if it was already used you can increase the range of IPs (Note, this will require stopping all running iSCSI disks first, update the IP range then start them again).

When a snapshot is exported you will be able to view the export information.

## Export Information

**Disk Name:**

00001-20230822220250

**Ethernet Card:**

eth2

**IQN:**

iqn.2016-05.com.petasan:00001-20230822220250

**IP:**

10.0.3.35

Close

Some actions are common on all iSCSI disks in the list which are:

- Start all  
Starts all stopped disks except the replication target disks.
- Stop all  
Stops all started disks
- Snapshots  
Views number of snapshots taken for each disk with the ability to view list of snapshots for a selected disk or delete all snapshots of the selected disk.  
You can also delete all snapshots for all disks using the Delete All button.

### Snapshots

X

Delete All

Show  entries

Search:

Disk ID	Number of Snapshots	Actions
00001	1	<input type="button" value="x"/>
00003	1	<input type="button" value="x"/>

Showing 1 to 2 of 2 entries

Previous **1** Next

Close

➤ **Note:**

- You can use the arrows in each column for sorting in descending or ascending order.
- You can search for a disk using any part of the displayed information.

## 6.2. Path Assignment

You can view all running nodes and all paths assigned to each node and the network interface used; you can reassign paths by clicking on the “Reassign Paths” button.

The screenshot shows the 'Path Assignment List' interface. At the top, there is a navigation bar with the PetaSAN logo and a 'Reassign Paths' button. Below the navigation bar, the page title 'Path Assignment List' is displayed. The main content area shows three nodes, each with a table of paths. The nodes are 'ps-node-01', 'ps-node-02', and 'ps-node-03'. Each node has a dropdown arrow and a 'Reassign Paths' button. The tables show IP Address, Interface, and Disk Name for each path.

IP Address	Interface	Disk Name
10.0.3.103	eth2	Quorum Disk
10.0.3.105	eth2	Hyper-V Datastore
10.0.3.106	eth2	Hyper-V Datastore

IP Address	Interface	Disk Name
10.0.2.105	eth1	Hyper-V Datastore
10.0.2.106	eth1	Hyper-V Datastore
10.0.3.104	eth2	Quorum Disk
10.0.3.107	eth2	Hyper-V Datastore

IP Address	Interface	Disk Name
10.0.2.103	eth1	Quorum Disk
10.0.2.104	eth1	Quorum Disk

You can expand, collapse nodes or use the expand all or collapse all buttons

After clicking on the “Reassign Paths” the system will enable you to re-assign paths manually or automatically

The screenshot shows the 'Path Assignment List' interface with search and reassignment options. At the top, there is a navigation bar with the PetaSAN logo and a 'Reassign Paths' button. Below the navigation bar, the page title 'Path Assignment List' is displayed. The main content area shows a 'Path Assignment:' section with radio buttons for 'Manual' and 'Automatic'. Below this is a 'Search by:' section with radio buttons for 'Disk Name' and 'Path IP', and a search button. The main content area also shows three nodes, each with a table of paths. The nodes are 'ps-node-01', 'ps-node-02', and 'ps-node-03'. Each node has a dropdown arrow and a 'Reassign Paths' button. The tables show IP Address, Interface, and Disk Name for each path.

**Path Assignment:**  
 Manual  Automatic

**Search by:**  
 Disk Name:   Path IP:

IP Address	Interface	Disk Name
10.0.3.103	eth2	Quorum Disk
10.0.3.105	eth2	Hyper-V Datastore
10.0.3.106	eth2	Hyper-V Datastore

IP Address	Interface	Disk Name
10.0.2.105	eth1	Hyper-V Datastore
10.0.2.106	eth1	Hyper-V Datastore
10.0.3.104	eth2	Quorum Disk

You can search for specific paths using part of the disk name or using specific path IP

**Path Assignment List**

Path Assignment :  
 Manual  Automatic

Search by :  
 Disk Name :   Path IP :

>> <<

**ps-node-01** ▾ 2 Paths

<input type="checkbox"/>	IP Address	Interface	Disk Name
<input type="checkbox"/>	10.0.3.105	eth2	Hyper-V Datastore
<input type="checkbox"/>	10.0.3.106	eth2	Hyper-V Datastore

**ps-node-02** > 3 Paths

**ps-node-03** ▾ 3 Paths

<input type="checkbox"/>	IP Address	Interface	Disk Name
<input type="checkbox"/>	10.0.2.107	eth1	Hyper-V Datastore
<input type="checkbox"/>	10.0.2.108	eth1	Hyper-V Datastore
<input type="checkbox"/>	10.0.3.108	eth2	Hyper-V Datastore

You can select one or more paths and select the destination node. You can select Auto which will assign the selected paths to the node(s) containing the least number of paths.

10.0.3.106 eth2 Hyper-V Datastore

**ps-node-02** ▾ 4 Paths

<input type="checkbox"/>	IP Address	Interface	Disk Name
<input type="checkbox"/>	10.0.2.105	eth1	Hyper-V Datastore
<input checked="" type="checkbox"/>	10.0.2.106	eth1	Hyper-V Datastore
<input type="checkbox"/>	10.0.3.104	eth2	Quorum Disk
<input type="checkbox"/>	10.0.3.107	eth2	Hyper-V Datastore

**ps-node-03** ▾ 5 Paths

<input type="checkbox"/>	IP Address	Interface	Disk Name
<input type="checkbox"/>	10.0.2.103	eth1	Quorum Disk
<input checked="" type="checkbox"/>	10.0.2.104	eth1	Quorum Disk
<input checked="" type="checkbox"/>	10.0.2.107	eth1	Hyper-V Datastore
<input type="checkbox"/>	10.0.2.108	eth1	Hyper-V Datastore
<input type="checkbox"/>	10.0.3.108	eth2	Hyper-V Datastore

Assign To:

You can select one or more paths then select the destination node  
 After clicking on the “Assign ” button , the system will start moving the paths , while the current assignments is in progress and you will be able to see the summary of the current path assignments as well as the current moving status of each path,

**Path Assignment List**

**Info**  
 Paths re-assignment in progress...  
 Total: 3      Moved: 1      Failed: 0      Moving: 1      Pending: 1

Node01 ▾ 5 Paths

IP Address	Disk Name	Status
10.0.2.114	Quorum	
10.0.3.111	VMWare	
10.0.2.118	Hyper-V	
10.0.2.126	Windows 2012	Moving
10.0.2.133	SQL2014	Moved

Node02 ▾ 3 Paths

IP Address	Disk Name	Status
10.0.3.115	Quorum	
10.0.3.117	Quorum	Pending
10.0.2.130	SQL2014	

The status can be one of the following:

- Pending  
 Path still waiting its turn to be moved
- Moving  
 Path is being moved from source to destination node.
- Moved  
 Path has been assigned to the destination node

You can move the paths automatically; currently the system will distribute the paths according to the number of paths on each node, trying to assign paths equally on while trying to avoid assigning too many paths that belong to the same disk to a single node.

**Path Assignment List**

**Path Assignment :**  
 Manual     Automatic

**Distribute by:**  
 Path Count     Resource Load

Assign

*Note: In case you are using VLANs for the iSCSI subnets, system will view the VLAN id next to the interface name ex: eth1.2.*

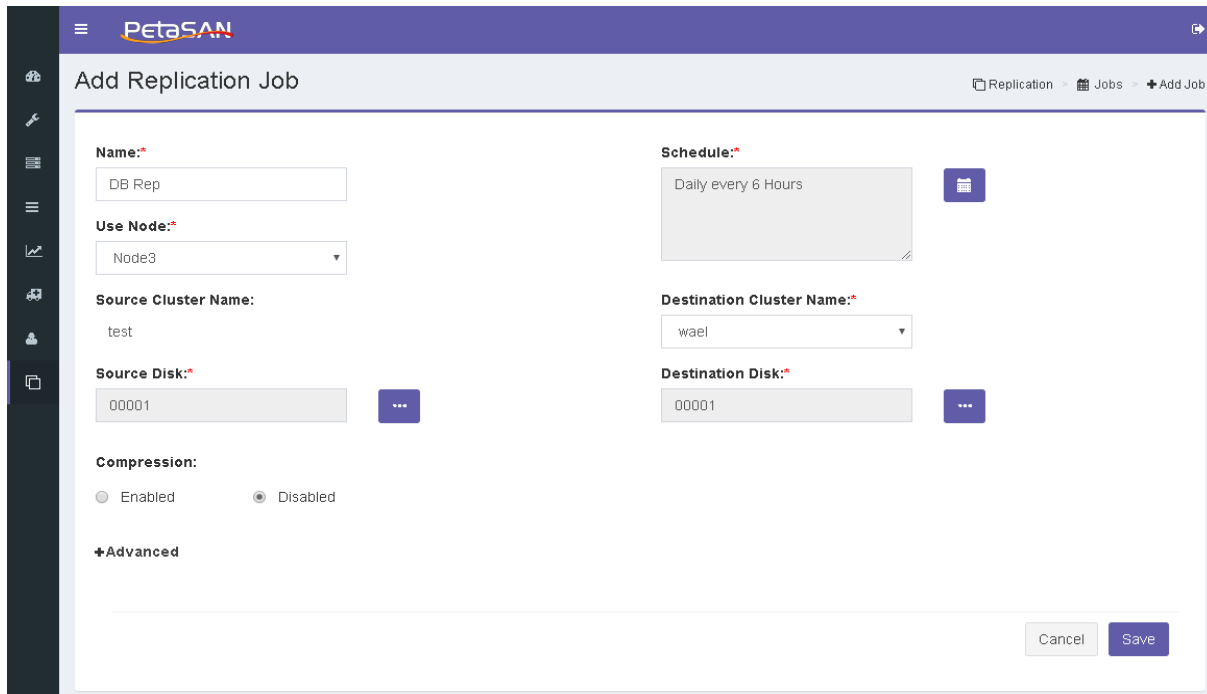
### 6.3. Replication

Replication enables you to replicate data from a source disk to a destination disk in a different cluster; this is activated by creating a replication job on the source cluster.

Replication is enabled by creating a replication job on the source cluster. The first time a replication job runs, it will transfer all data from the source disk to the destination disk, later jobs instances will only transfer the changes/differences done since the previous replication instance.


#### ➤ Add Replication Job

Adds a replication job on the source cluster



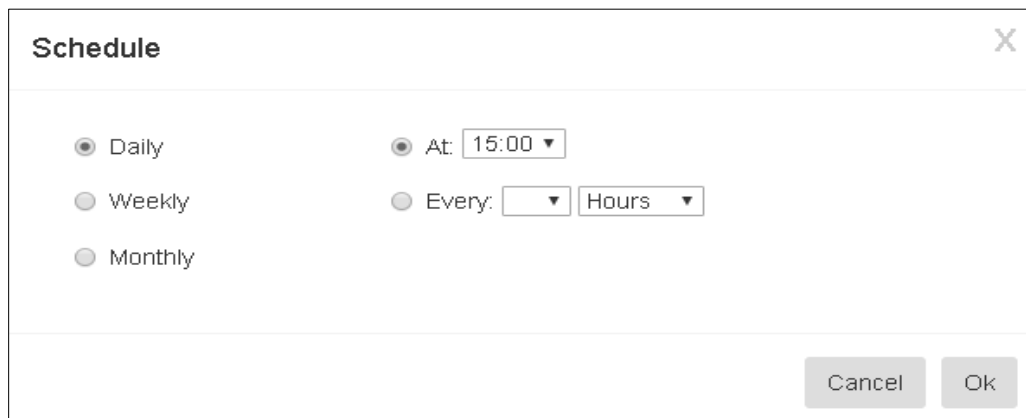
The form requires the following information:

- **Name**  
The job name
- **Use Node**  
The node that will be used to run the replication job on the source cluster, the system allows you to select a node from list of Backup/Replication nodes.
- **Schedule**  
Allows you to select how frequent the job will be executed.



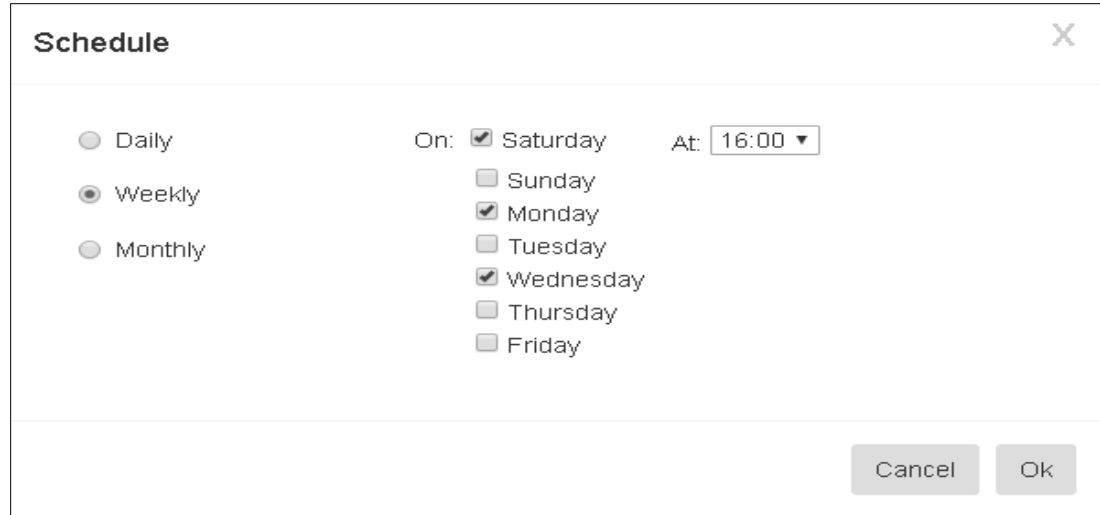
First you should select the job frequency from one of the following options:

○ **Daily**



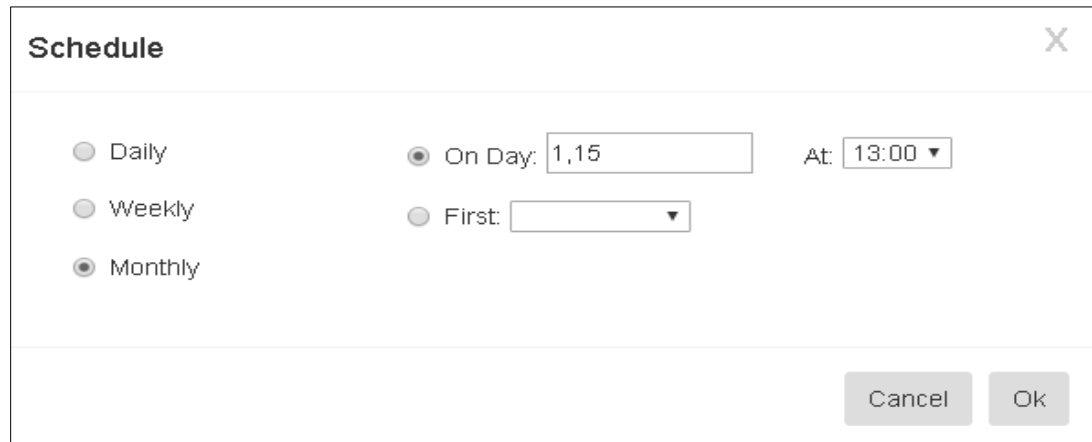
- **At**  
Job will run daily at a specific hour, you select the hour in 24 hours format
- **Every**  
Job will run every x minutes or x hours.

○ **Weekly**



- **Week Days**  
You can select the weeks days the job will run on.
- **At**  
Job will run daily at a specific hour, you will select the hour in 24 hours format.

- **Monthly**



- **On Day**  
You can select the day(s) to run the replication job on.
- **First**  
Job can run on first Sunday, Monday.., of the month.
- **At**  
Job will run at specific hour, you will select the hour in 24 hours format.

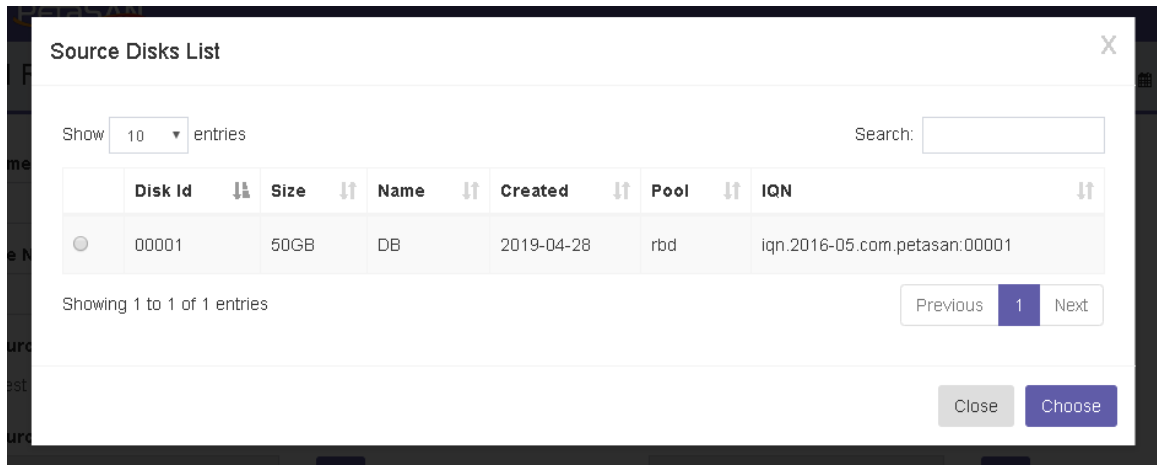
- **Source Cluster Name**

The source cluster name (name of the current cluster).

- **Source Disk**

The system enables you to select an iSCSI disk from the list of disks on the source cluster.

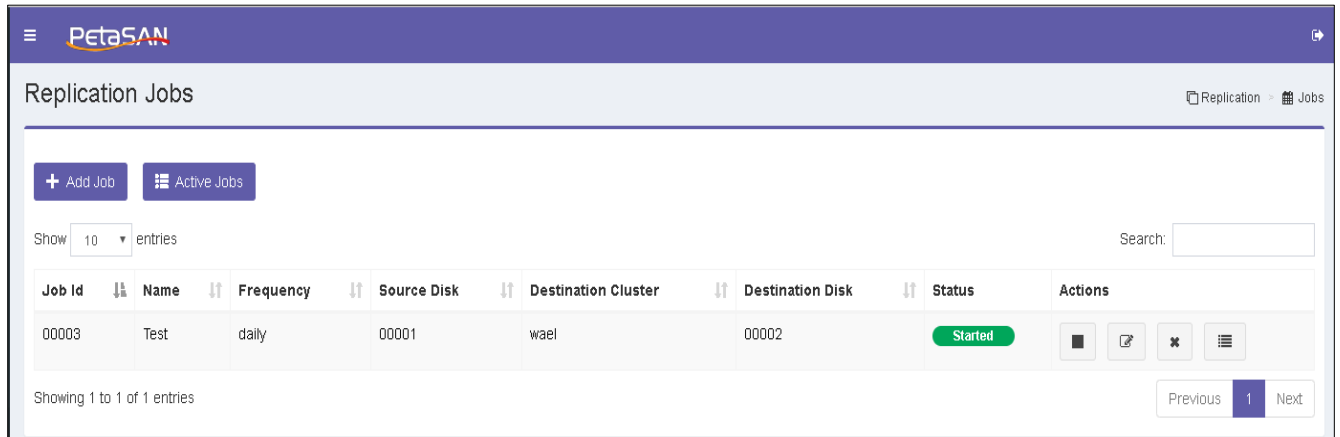




- Destination Cluster Name**  
 The destination cluster name which you can select from the destination clusters you defined; refer to Manage Destination Clusters for more information
- Destination Disk**  
 The system enables you to select an iSCSI disk from the list of replication target disks on the destination cluster.
- Compression**  
 Select this if you want to enable compression during data transfer.
- Pre Snapshot Script URL**  
 Optional URL of a custom script you want to run before the replication job takes a snapshot from the source disk.
- Post Snapshot Script URL**  
 Optional URL of a custom script you want to run after taking the snapshot from the source disk.

## ➤ Jobs List

The system displays the list of defined jobs.



The Jobs List contains the following information for each Job:

- **Job Id**  
The job serial number
- **Name**  
The Job Name
- **Frequency**  
How often the job will run (daily, weekly or monthly)
- **Source Disk**  
The source disk id.
- **Destination Cluster**  
The name of the destination cluster.
- **Destination Disk**  
The destination disk id.
- **Status**  
Displays the current status of the job, which can be one of the following:
  - Started  
Job will run according to its schedule.
  - Stopped  
Job is suspended and will not run.
- **Actions**  
The Actions column displays the actions that can be done on each job according to its current status; each action has a tool tip that is displayed when the mouse is over the action button.
  - Stop  
Stops the job, it is available only when the job is started.
  - Start  
Starts the job, it is available only when the job is stopped.

- Edit  
Enables editing the job information except the source and destination disks.
- Delete  
Deletes the selected job
- Log  
Shows the job log information

### ➤ Active Jobs List

The system views the list of jobs currently in progress.

Active Replication Jobs X

Show  entries Search:

Id	Name	Start Time	Elapsed Time	Transfer Rate	Transferred	Compression	Progress	Actions
00001	DB Rep	2019-04-28 14:20:08	0:02 hh:mm	1.6 MB/s	0.05 GB	1.25	4%	✖

Showing 1 to 1 of 1 entries Previous **1** Next

[Close](#)

The Active Jobs List displays the following information for each Job:

- **Id**  
The job id
- **Name**  
The Job Name
- **Current Occurrence**  
The current job instance time occurrence
- **Elapsed Time**  
Time elapsed since the job started
- **Transfer Rate**  
Transfer speed in bytes per second
- **Transferred Data**  
Number of bytes transferred till now
- **Compression Ratio**  
Ratio of original byte count / count after compression
- **Progress**  
Percentage of the transferred data to the total data to be transferred
- **Actions**  
The Actions column displays the actions that can be done on each active job
  - Cancel  
Cancel the running job.

## ➤ Add Replication User

This enables the administrator of the destination cluster to create a new user for usage by the system during job execution

The replication user will have limited access to enable him to run the replication job.

The screenshot shows the 'Add Replication User' interface in the PetaSAN web console. The breadcrumb trail indicates the path: Replication > Users > +Add User. The form fields are as follows:

- User Name:\***: A text input field containing the value 'Replicator'.
- Authorized Pools:\***: A dropdown menu with 'rbd' selected.
- User's private key:**: A large, empty text area for pasting a private key.

At the bottom right of the form, there are two buttons: 'Cancel' and 'Save'.

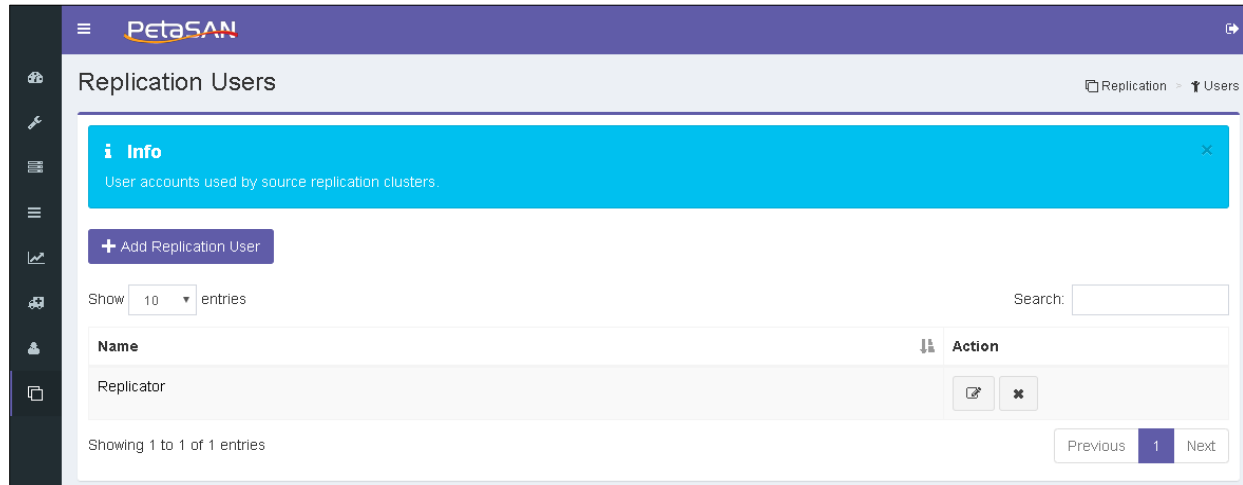
The form requires the following information:

- **User Name**  
The user name, the system will create an OS system user as well as a Ceph user.
- **Authorized Pools**  
The pools the user is granted access to.
- **User's Private Key**  
The user's private key generated by the system after saving the user information.

➤ *Note: The user will be created only on the Backup/Replication Nodes of the destination cluster.*

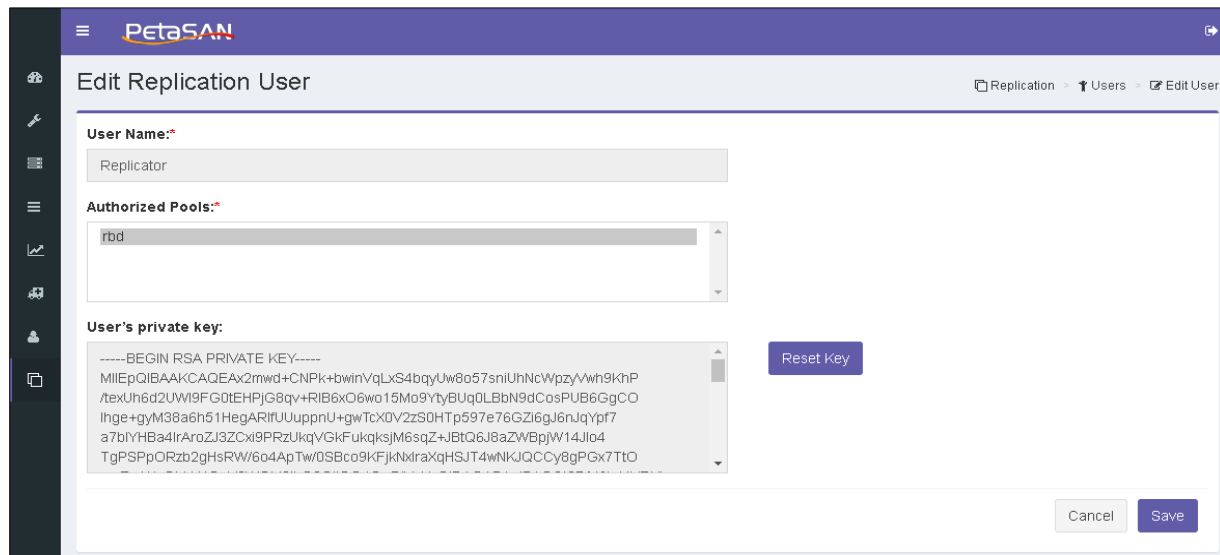
## ➤ Replication Users List

The system lists the current replication users.



The Replication Users list shows the following information for each user:

- **Name**  
User Name
- **Actions**  
The system displays the actions that can be done on the user
  - Edit  
The system enables the administrator to update the authorized pools or reset the user's private key.

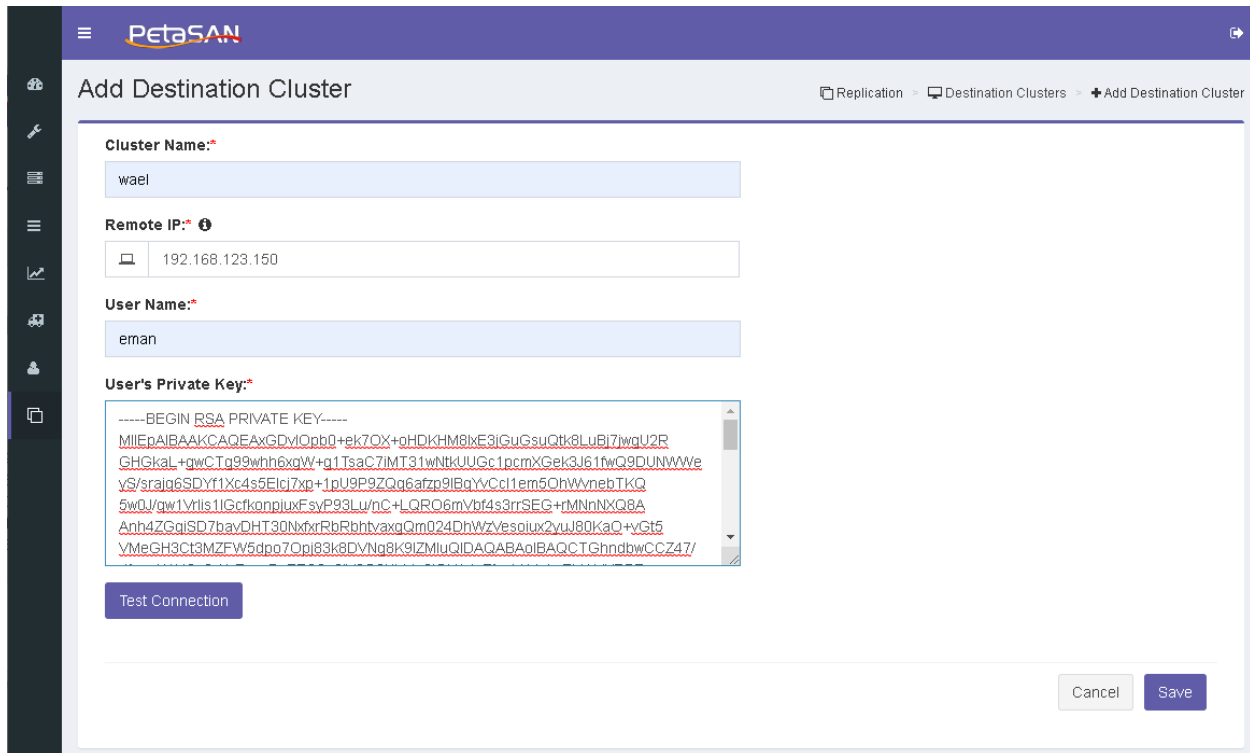


- **Delete**  
The system deletes the selected user from the system OS and from Ceph users.

➤ *Note: Any updates done to the user will be applied to all backup/replication nodes.*

## ➤ Add Destination Cluster

The system enables the administrator of the source cluster to define a target destination cluster



The screenshot shows the 'Add Destination Cluster' configuration interface. The fields are as follows:

- Cluster Name:** wael
- Remote IP:** 192.168.123.150
- User Name:** eman
- User's Private Key:** -----BEGIN RSA PRIVATE KEY-----  
MIIEpAIBAAKCAQEAxGDvI0pb0+ek7OX+oHDkHM8kE3iGuGsuQtk8LuBj7wgU2R  
GHGkaL+qwCTq99whh6xgW+g1TsaC7iMT31wNtkUUGc1pcmXGek3J61fwQ8DUNWWe  
vS/sraja6SDYf1Xc4s5Eicj7xp+1pU9P9ZQq6afz9IBqYVccl1em5OhWwnebtKQ  
5wQJgw1Vrlis1IGcftkonpiuxFsyP93Lu/nC+LQRO6mVbf4s3rrSEG+rMNNXQ8A  
Anh4ZGqiSD7bavDHT30NMrRbRbhtvaxgQm024DhwZvesoiux2yuj80KaQ+yGt5  
VMeGH3Ct3MZFW5dpg7Opj83k8DVNg8k9IZMluQIDAQABAQCTGhndbwCCZ47/

Buttons: Test Connection, Cancel, Save

The Destination Cluster from contains the following fields:

- **Cluster Name**  
Name of the destination cluster.
- **Remote IP**  
IP of the Backup/Replication node at the destination cluster.
- **User Name**  
Name of the replication user created at the destination cluster.
- **User's Private Key**  
The private key of the user at the destination cluster.

## ➤ Destination Clusters List

The system views list of destination clusters

The screenshot shows the PetaSAN web interface for managing destination clusters. At the top, there is a navigation bar with the PetaSAN logo and a breadcrumb trail: Replication > Destination Clusters. Below the navigation bar, there is a header section with the title "Destination Clusters" and a search input field. A button labeled "+ Add Destination Cluster" is located on the left. Below the button, there is a "Show 10 entries" dropdown menu. The main content area displays a table with the following columns: Cluster Name, Remote IP, and Actions. The table contains one entry with the cluster name "wael" and the remote IP "192.168.123.151". The Actions column for this entry contains two icons: a pencil (edit) and a trash can (delete). At the bottom of the table, there is a pagination control showing "Showing 1 to 1 of 1 entries" and buttons for "Previous", "1", and "Next".

The Destination Clusters list views the following information about each user:

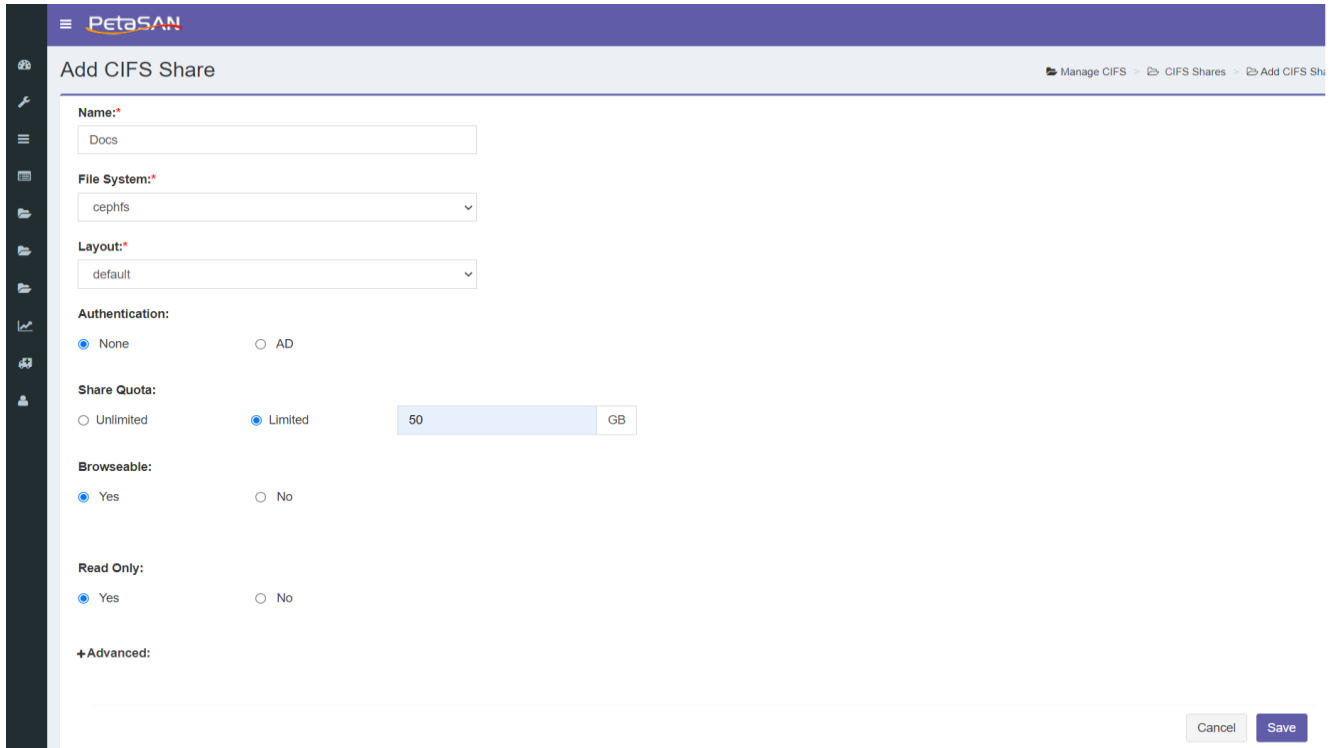
- **Cluster Name**  
The destination cluster name.
- **Remote IP**  
IP of the target Replication/Backup Node at the destination cluster.
- **Actions**  
The system displays the actions that can be done on the destination cluster.
  - Edit  
The system enables the administrator of the source cluster to update the destination cluster information.
  - Delete  
The system deletes the selected destination cluster.

## 7. Manage CIFS

### 7.1. CIFS Shares

#### ➤ Add CIFS Share

The system opens the CIFS Share form



The screenshot shows the 'Add CIFS Share' form in the PetaSAN interface. The form is titled 'Add CIFS Share' and has a breadcrumb trail: 'Manage CIFS > CIFS Shares > Add CIFS Share'. The form contains the following fields and options:

- Name:** A text input field containing 'Docs'.
- File System:** A dropdown menu with 'cephfs' selected.
- Layout:** A dropdown menu with 'default' selected.
- Authentication:** Radio buttons for 'None' (selected) and 'AD'.
- Share Quota:** Radio buttons for 'Unlimited' and 'Limited' (selected). A text input field next to 'Limited' contains '50', and a dropdown menu next to it is set to 'GB'.
- Browseable:** Radio buttons for 'Yes' (selected) and 'No'.
- Read Only:** Radio buttons for 'Yes' (selected) and 'No'.
- +Advanced:** A section that is currently collapsed.

At the bottom right of the form, there are 'Cancel' and 'Save' buttons.

The CIFS Share form contains the following fields:

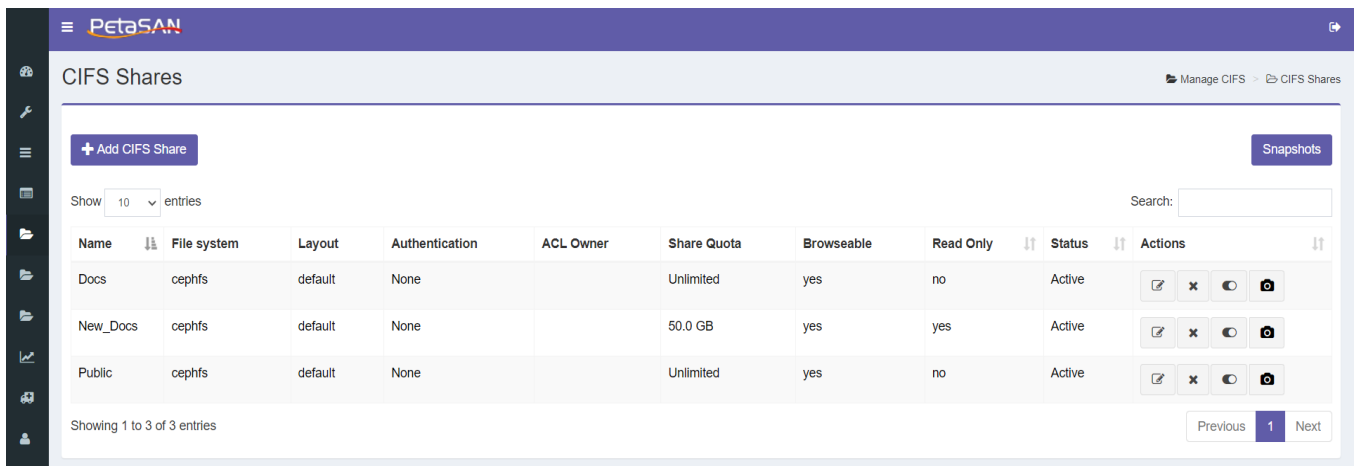
- **Name**  
The share name
- **File System**  
The file system on which the share will be created.
- **Layout**  
The layout in which the share will be created.
- **Authentication**  
You can select no authentication or AD; in case of you select AD you will be required to enter the ACL owner.
- **Share Quota**  
You can select to have unlimited quota, or specify limited quota in gigabytes
- **Browsable**  
Indicates if the share is browsable by clients.
- **Read only**



You can select to set the share to be read only or not.

- **smb.conf override**  
Configuration parameters to add/modify over existing common CIFS Settings.
- **Actions**
  - Save  
Saves the added share
  - Cancel  
Cancels any changes were done

## ➤ View CIFS Shares



Name	File system	Layout	Authentication	ACL Owner	Share Quota	Browseable	Read Only	Status	Actions
Docs	cephfs	default	None		Unlimited	yes	no	Active	[Edit] [Delete] [Toggle] [Snapshot]
New_Docs	cephfs	default	None		50.0 GB	yes	yes	Active	[Edit] [Delete] [Toggle] [Snapshot]
Public	cephfs	default	None		Unlimited	yes	no	Active	[Edit] [Delete] [Toggle] [Snapshot]

The system views list of existing shares showing the following columns:

- **Name**  
Share name
- **File System**  
The file system on which the share is created.
- **Layout**  
The layout in which the share is created.
- **Authentication**  
Shows if authentication is required to access the share.
- **ACL owner**  
The owner of the ACL, this is the user who can assign others permission on the share.
- **Share Quota**  
Shows sharequota, either unlimited or limited by number of GBs.
- **Browsable**  
Indicates if the share is browsable by clients.
- **Read Only**  
Shows if the share is read only or not.
- **Status**

Shows if the share is active so it can be accessed by the users or inactive.

- **Actions**

- Add  
Opens the new share form
- Edit  
Opens the edit share form, enables updating some of shares information such as the quota size.
- Delete  
Deletes the current share including all stored data.
- Activate/ Deactivate  
If the share is active, the system deactivates the share so it will not be accessible and vice versa.
- View Snapshots  
Views list of snapshots that were previously taken for the current CIFS share, snapshots may be taken manually or at a selected schedule (daily, weekly or monthly).

You can select number of scheduled snapshots the system should keep. You can have one scheduled

Share Docs Snapshots

X

Take Snapshot
Schedule Snapshots
View Schedule
Delete Schedule

Show 10 entries

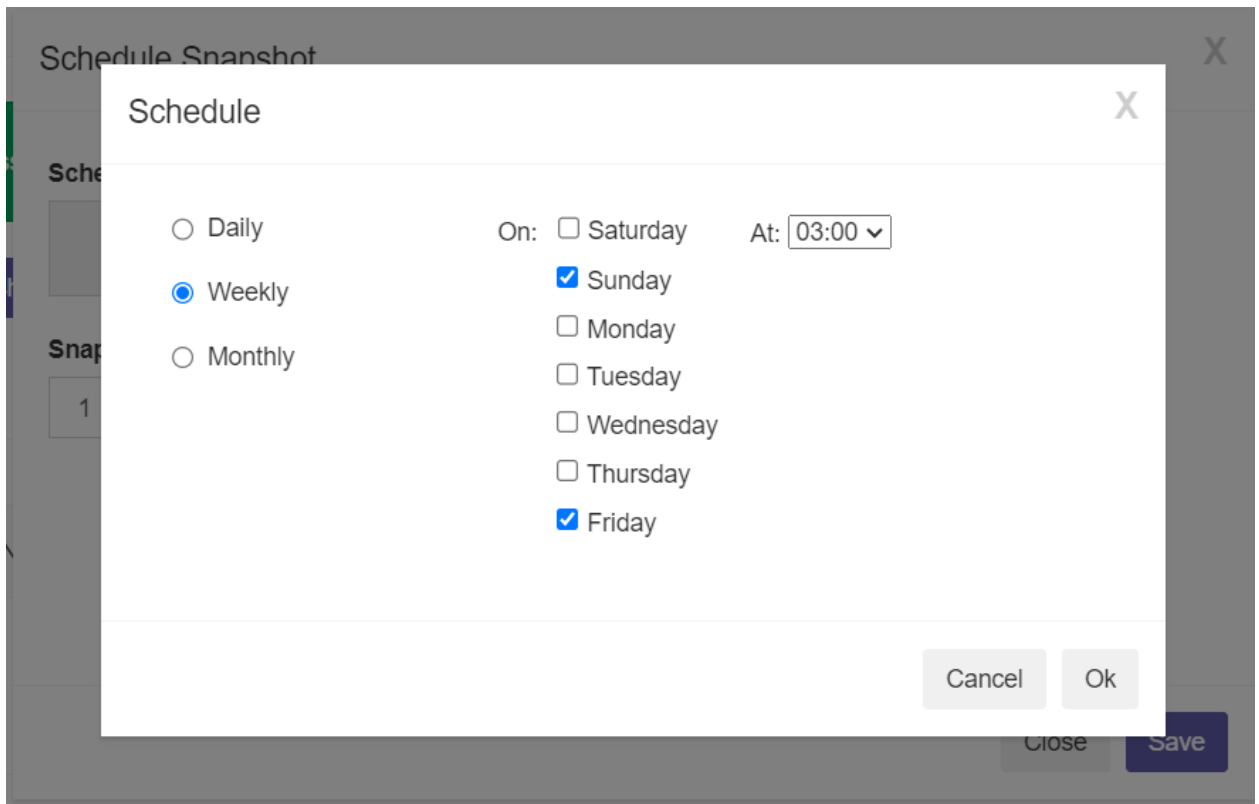
Time	Name	Type	Schedule	Exported	Status	Action
2023-08-20 17:00:02	PETASAN-S-20230820170002	Scheduled	Daily every 1 Hour	No		<span style="border: 1px solid #ccc; padding: 2px 5px;">✕</span> <span style="border: 1px solid #ccc; padding: 2px 5px;">↺</span> <span style="border: 1px solid #ccc; padding: 2px 5px;">↻</span>
2023-08-20 16:47:01	PETASAN-U-20230820164701	User		Yes		<span style="border: 1px solid #ccc; padding: 2px 5px;">✕</span> <span style="border: 1px solid #ccc; padding: 2px 5px;">↺</span> <span style="border: 1px solid #ccc; padding: 2px 5px;">↻</span> <span style="border: 1px solid #ccc; padding: 2px 5px;">📄</span> <span style="border: 1px solid #ccc; padding: 2px 5px;">i</span>
2023-08-20 16:00:02	PETASAN-S-20230820160002	Scheduled	Daily every 1 Hour	No		<span style="border: 1px solid #ccc; padding: 2px 5px;">✕</span> <span style="border: 1px solid #ccc; padding: 2px 5px;">↺</span> <span style="border: 1px solid #ccc; padding: 2px 5px;">↻</span>
2023-08-20 15:00:02	PETASAN-S-20230820150002	Scheduled	Daily every 1 Hour	No		<span style="border: 1px solid #ccc; padding: 2px 5px;">✕</span> <span style="border: 1px solid #ccc; padding: 2px 5px;">↺</span> <span style="border: 1px solid #ccc; padding: 2px 5px;">↻</span>
2023-08-20 14:48:58	PETASAN-U-20230820144858	User		No		<span style="border: 1px solid #ccc; padding: 2px 5px;">✕</span> <span style="border: 1px solid #ccc; padding: 2px 5px;">↺</span> <span style="border: 1px solid #ccc; padding: 2px 5px;">↻</span>

Showing 1 to 5 of 5 entries

Previous
1
Next

Close

snapshot at a time, when added, you will be able to view the saved schedule or delete the schedule if you want.



For each snapshot taken either scheduled or manually by the user you can use one of the following actions:

-Delete

Deletes the selected snapshot

-Rollback

Rolls back the current share data to the selected snapshot.

-Export/Disable export

Exports the selected snapshot so it can be accessed outside the system, system can export one snapshot at a time, when the share is exported the Information action appears

-Export Information

The system view the exported share information, the exported snapshot name and range of IPs can be used to connect

## Export Information

**Share Name:**

Docs-20230820164701

**IP Range:**

**From:**

10.0.3.10

**To:**

10.0.3.40

Close

Some actions are common on all CIF Shares in the list which are:

- Snapshots

Views number of snapshots taken for each share with the ability to view list of snapshots for a selected share, delete all snapshots of a selected share.

You can delete all snapshots for all shares using Delete All button.

Snapshots

X

Delete All

Show  entries

Search:

Share Name	Number of Snapshots	Actions
Docs	5	<input type="button" value="x"/>
New_Docs	1	<input type="button" value="x"/>
Public	1	<input type="button" value="x"/>

Showing 1 to 3 of 3 entries

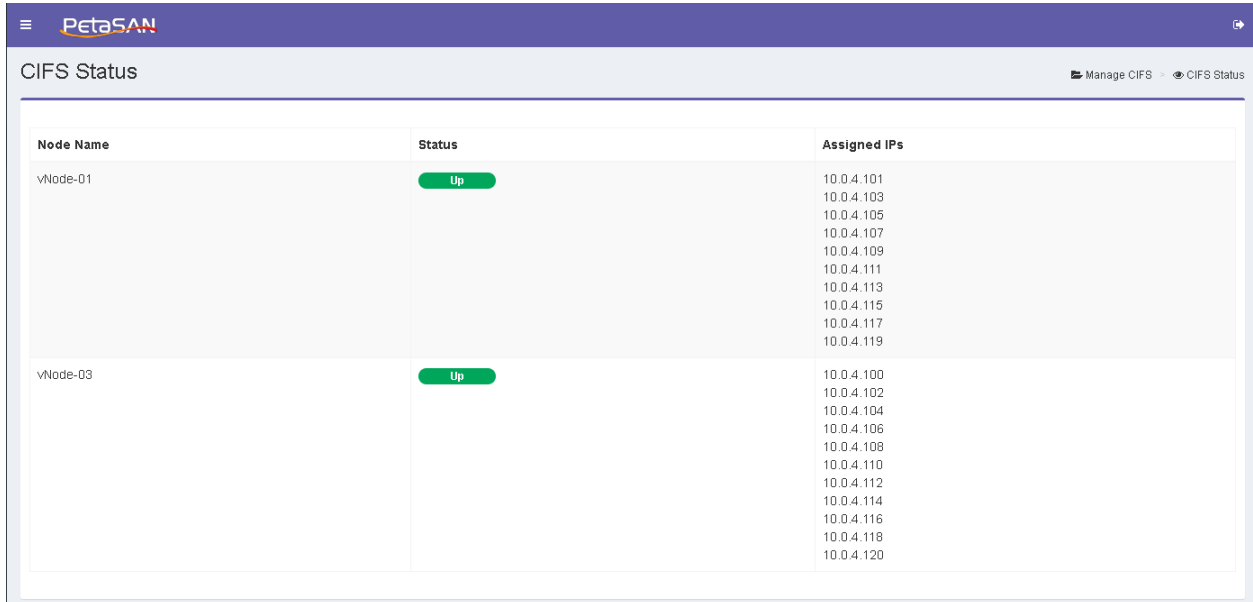
Previous **1** Next

Close

## 7.2. CIFS Status

### ➤ View CIFS Status

You can view the current CIFS status



Node Name	Status	Assigned IPs
vNode-01	Up	10.0.4.101 10.0.4.103 10.0.4.105 10.0.4.107 10.0.4.109 10.0.4.111 10.0.4.113 10.0.4.115 10.0.4.117 10.0.4.119
vNode-03	Up	10.0.4.100 10.0.4.102 10.0.4.104 10.0.4.106 10.0.4.108 10.0.4.110 10.0.4.112 10.0.4.114 10.0.4.116 10.0.4.118 10.0.4.120

The system will view the following columns:

- **Node Name**  
Name of the CIFS Server Node
- **Status**  
Node status, up or down
- **Assigned IPs**  
IPs assigned to each node
- **Number of Connections**  
Number of active client connections

### ➤ View Connections

You can view connections of a selected server node

Screen

The system will view the following columns for each active connection:

- **Username**  
The user name of the client that is connected
- **Group Name**  
The group name of the client that is connected.
- **Client IP**  
The IP of the connected client
- **Share Name**

The shares which the client is connected to

## 7.3. Reassign IPs

Opens the Reassign IPs form which enables you to move selected IPs manually to another node or distribute them automatically

The screenshot shows the 'IP Reassignment List' interface. At the top, there is a breadcrumb trail: 'Manage CIFS > IP Reassignment List'. Below this, the 'IP Reassignment:' section has two radio buttons: 'Manual' (selected) and 'Auto'. The main content area is a table with two columns: 'From Node' and 'Assigned IPs'. The table lists two nodes: 'Node2' and 'Node3'. Node2 has 4 IPs listed: 10.0.3.10 (checked), 10.0.3.12 (checked), 10.0.3.14 (unchecked), and 10.0.3.16 (unchecked). Node3 has 3 IPs listed: 10.0.3.11 (unchecked), 10.0.3.13 (unchecked), and 10.0.3.15 (unchecked). Below the table, there is an 'Assign To:' dropdown menu currently set to 'Node3'. A 'Reassign' button is located at the bottom right of the interface.

From Node	Assigned IPs
Node2	4 IPs <input type="checkbox"/> IP <input checked="" type="checkbox"/> 10.0.3.10 <input checked="" type="checkbox"/> 10.0.3.12 <input type="checkbox"/> 10.0.3.14 <input type="checkbox"/> 10.0.3.16
Node3	3 IPs <input type="checkbox"/> IP <input type="checkbox"/> 10.0.3.11 <input type="checkbox"/> 10.0.3.13 <input type="checkbox"/> 10.0.3.15

Assign To:

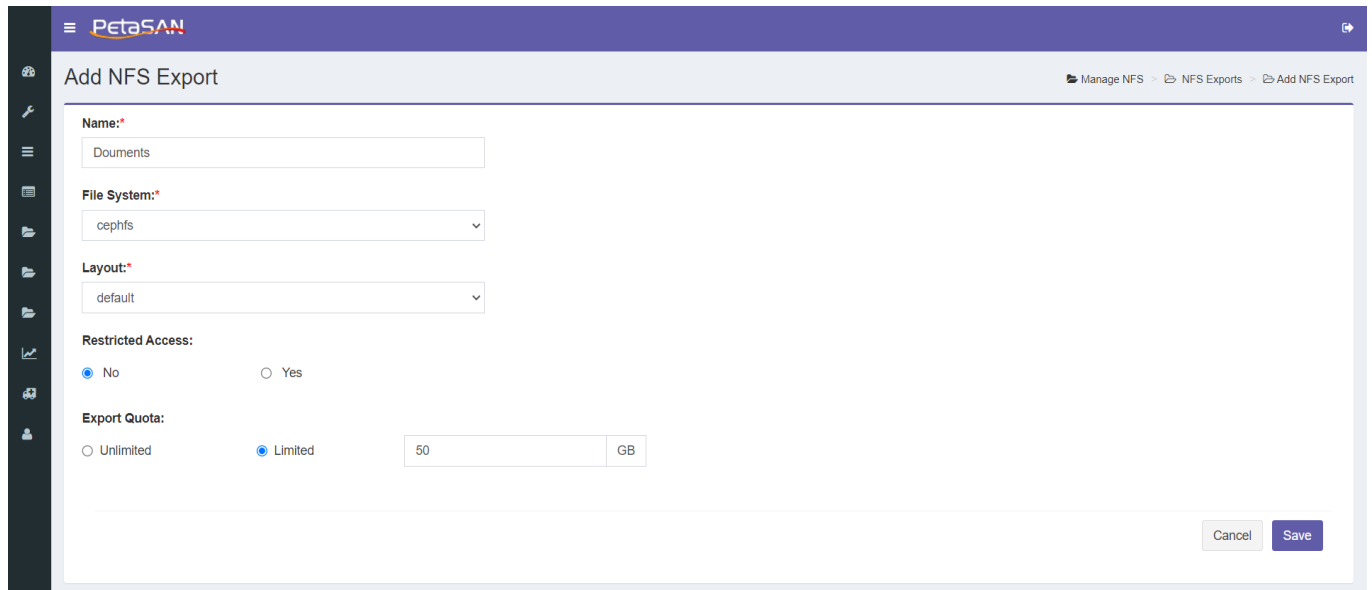
Reassign

## 8. Manage NFS

### 8.1. NFS Exports

#### ➤ Add NFS Export

The system opens the NFS export form



The screenshot shows the 'Add NFS Export' form in the PetaSAN interface. The form is titled 'Add NFS Export' and has a breadcrumb trail: 'Manage NFS > NFS Exports > Add NFS Export'. The form contains the following fields and options:

- Name:** A text input field containing 'Documents'.
- File System:** A dropdown menu with 'cephfs' selected.
- Layout:** A dropdown menu with 'default' selected.
- Restricted Access:** Two radio buttons: 'No' (selected) and 'Yes'.
- Export Quota:** Two radio buttons: 'Unlimited' and 'Limited' (selected). Next to 'Limited' is a text input field containing '50' and a 'GB' unit selector.

At the bottom right of the form, there are two buttons: 'Cancel' and 'Save'.

The NFS export form contains the following fields:

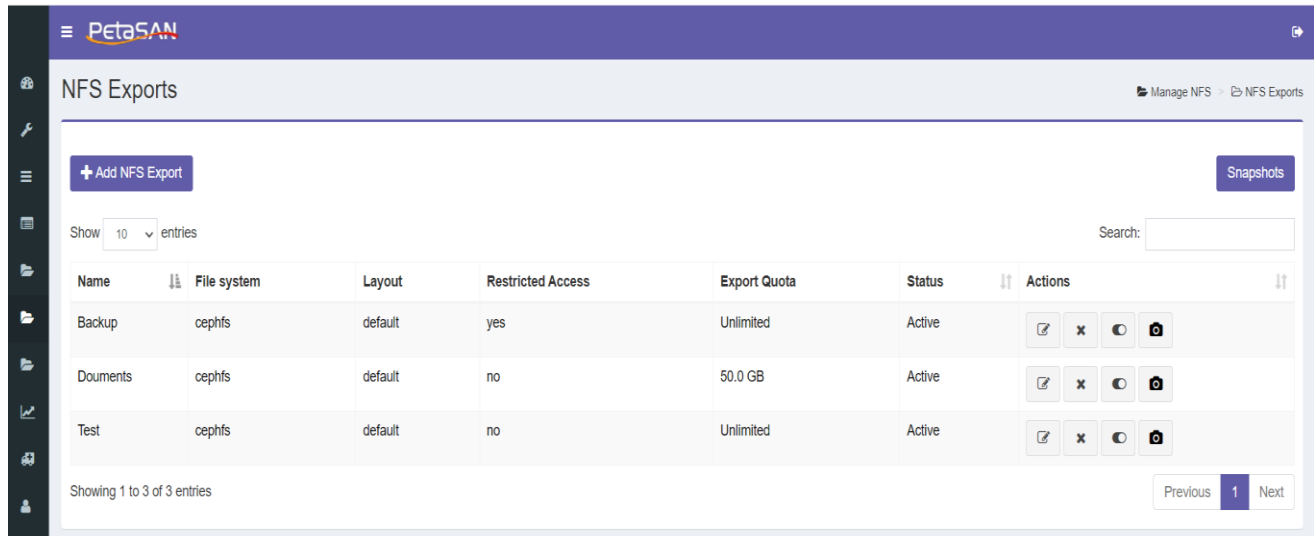
- **Name**  
The export name
- **File System**  
The file system on which the export will be created.
- **Layout**  
The layout which the export will be created.
- **Restricted Access**
- You can create the share without any access restrictions or specify the clients with read/write or read only access. **Export Quota:**  
You can specify that the share has unlimited quota, or you can specify quota in Gigabytes

**Actions**

- Save  
Saves the added export.
- Cancel  
Cancels any done changes.

## ➤ View NFS Exports

The system views the list of existing exports showing the following columns:



Name	File system	Layout	Restricted Access	Export Quota	Status	Actions
Backup	cephfs	default	yes	Unlimited	Active	[Edit] [Delete] [Toggle] [Snapshots]
Documents	cephfs	default	no	50.0 GB	Active	[Edit] [Delete] [Toggle] [Snapshots]
Test	cephfs	default	no	Unlimited	Active	[Edit] [Delete] [Toggle] [Snapshots]

- **Name**  
Export name
- **File System**  
The file system in which the export is created.
- **Layout**  
The layout in which the export is created.
- **Restricted Access**  
Views if access is restricted for the export or not.
- **Export Quota**  
Views if export has unlimited quota, or limited and showing the quota export in Gigabytes.
- **Status**  
Views if export is active or in active
- **Actions**
  - Add  
Opens the new export form
  - Edit  
Opens the edit export form, you can update the access writes or export quota.
  - Delete  
Deletes the current export including all stored data.
  - Activate/Deactivate share  
Deactivates the export if it was activated and vice verse.
  - Manage Snapshots



Manages snapshots that views list of snapshots that were previously taken for the current NFS Export, snapshots may be taken manually or at a selected schedule (daily, weekly or monthly).

Export Backup Snapshots



- Take Snapshot
- Schedule Snapshots
- View Schedule
- Delete Schedule

Show  entries

Search:

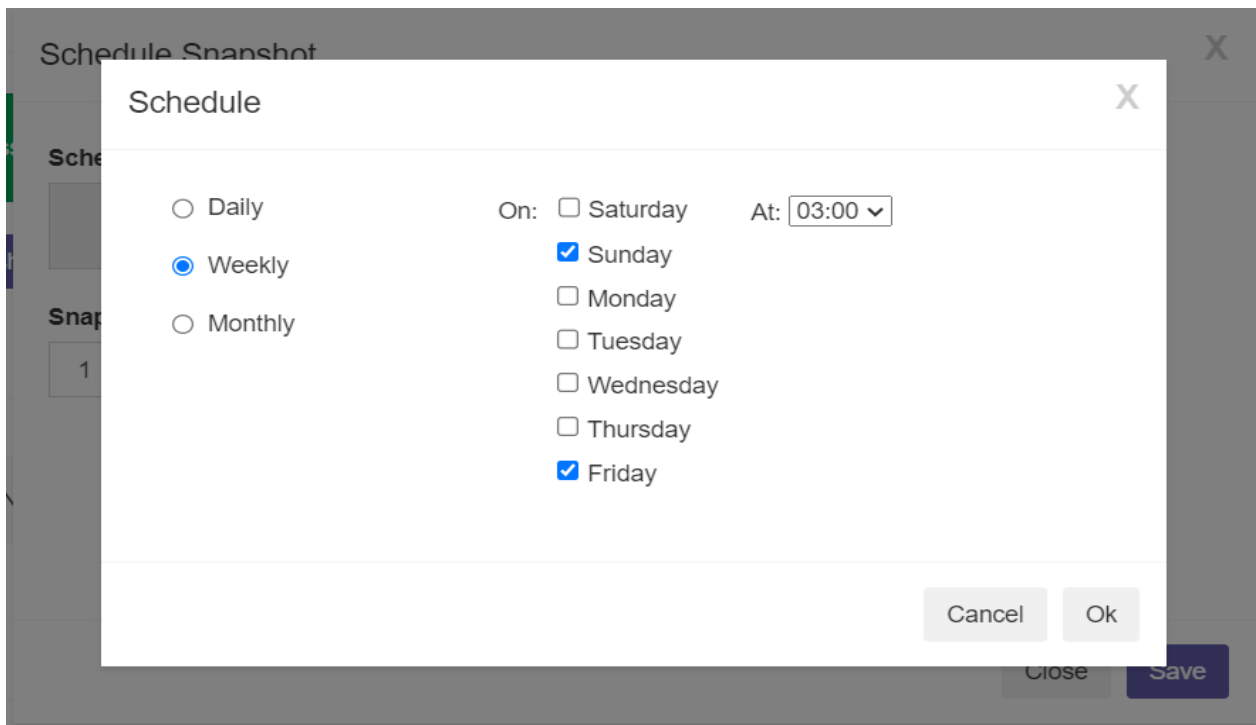
Time	Name	Type	Schedule	Exported	Status	Action
2023-08-20 22:00:02	PETASAN-S-20230820220002	Scheduled	Daily every 1 Hour	No		<input type="button" value="x"/> <input type="button" value="↺"/> <input type="button" value="↻"/>
2023-08-20 21:46:45	PETASAN-U-20230820214645	User		No		<input type="button" value="x"/> <input type="button" value="↺"/> <input type="button" value="↻"/>

Showing 1 to 2 of 2 entries

Previous **1** Next

Close

You can select number of scheduled snapshots the system should keep. You can have one scheduled snapshot at a time, when added, you will be able to view the saved schedule or delete the schedule if you want



For each snapshot taken either scheduled or taken by the user you can use one of the following actions:

-Delete

Deletes the selected snapshot

**-Rollback**

Rolls back the current export data to the selected snapshot.

**-Export/Disable export**

Exports the selected snapshot so it can be accessed outside the system, system can export one snapshot at a time, when the share is exported the system uses the last IP in the NFS subnet IP range. When an export is selected the export information button will appear.

**-Export Information**

The system view the exported export information, the exported snapshot name and the ip that should be used to connect

---

### Export Information

---

**NFS Export Name:**  
e1-20230822223322

**IP:**  
10.0.3.12

---

[Close](#)

Some actions are common on all NFS Exports in the list which are:

- Snapshots  
Views number of snapshots taken for each export with the ability to view list of snapshots for a selected export, delete all snapshots of a selected export.  
You can delete all snapshots for all exports using Delete All button.

Snapshots X

---

[Delete All](#)

Show  entries Search:

Export Name	Number of Snapshots	Actions
Backup	1	<a href="#">x</a>

Showing 1 to 1 of 1 entries [Previous](#) [1](#) [Next](#)

---

[Close](#)

## 8.2. NFS Status

### ➤ View NFS Status

You can view the current NFS status

Node Name	Status	Assigned IPs	Number Of Connections
node-01	Up	192.168.50.100 192.168.50.102	2
node-02	Up	192.168.50.101 192.168.50.103 192.168.50.104	0

The system will view the following columns:

- **Node Name**  
Name of the NFS Server Node
- **Status**  
Node status, up, down or in grace where in grace means node can't accept new client connections
- **Assigned IPs**  
IPs assigned to each node
- **Number of Connections**  
Number of active client connections

### ➤ View Connections

You can view connections of a selected server node

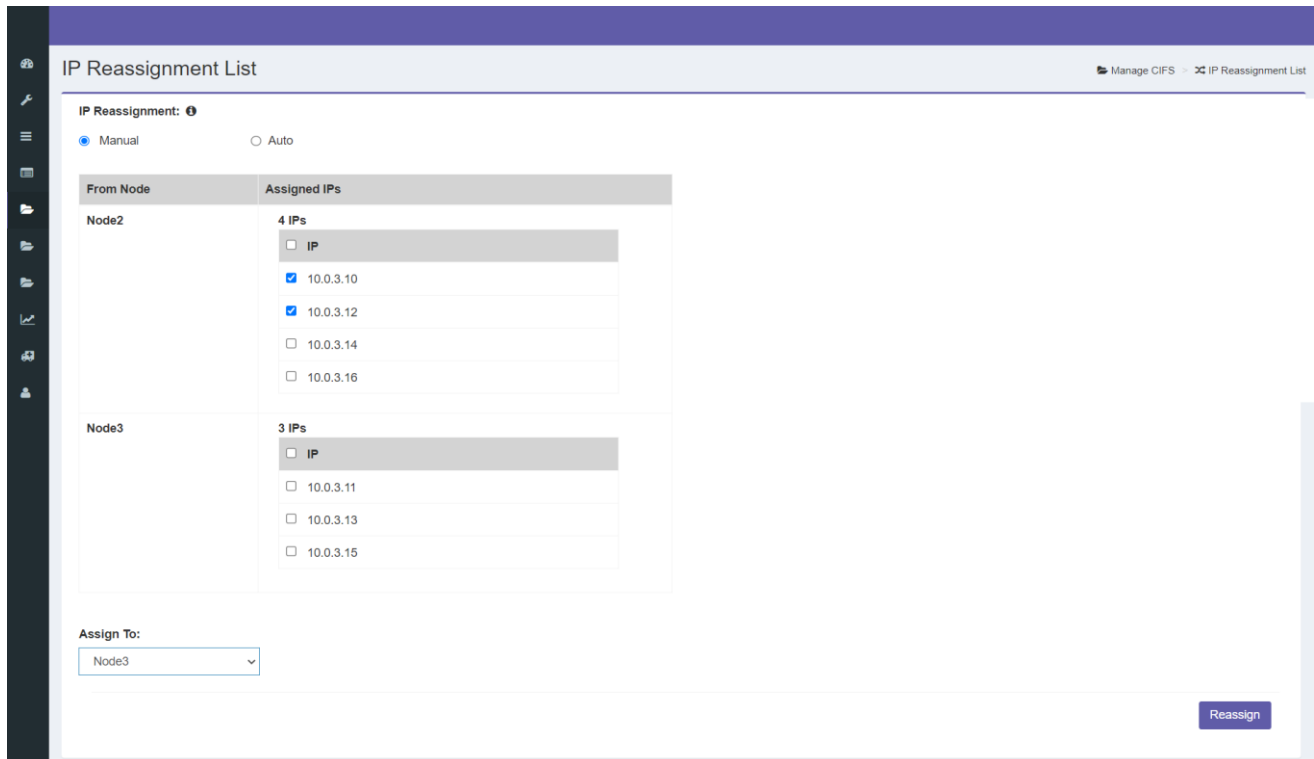
Client IP	Session ID
192.168.50.14	6847447354782515201 6847447320422776833

The system will view the following columns for each active connection:

- **Client IP**  
The IP of the client that is connected
- **Session ID**  
The session id of the connected client.

### 8.3. Reassign IPs

Opens the Reassign IPs form which enables you to move some selected IPs manually to another node or distribute them automatically

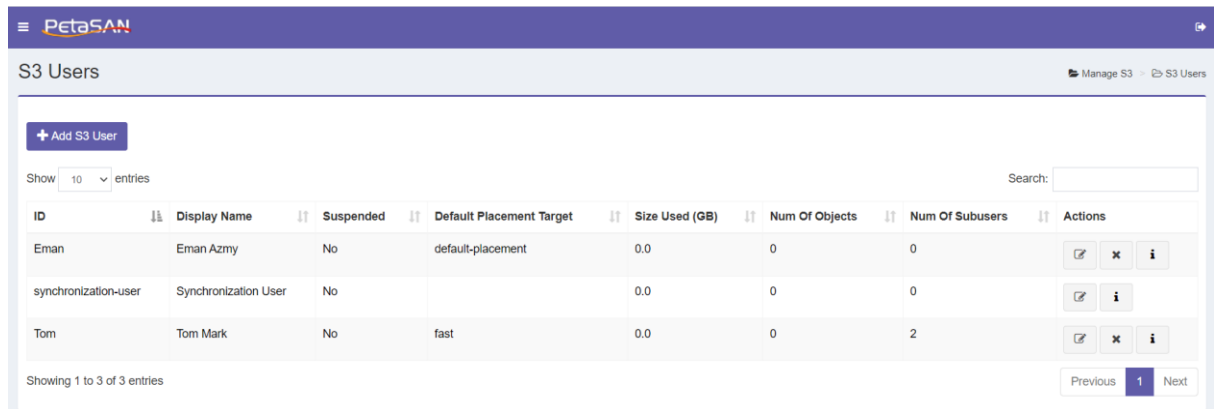


## 9. Manage S3

### 9.1. S3 Users

#### ➤ View S3 Users List

The system views the list of existing S3 Users showing the following columns:



ID	Display Name	Suspended	Default Placement Target	Size Used (GB)	Num Of Objects	Num Of Subusers	Actions
Eman	Eman Azmy	No	default-placement	0.0	0	0	[Edit] [Delete] [Info]
synchronization-user	Synchronization User	No		0.0	0	0	[Edit] [Info]
Tom	Tom Mark	No	fast	0.0	0	2	[Edit] [Delete] [Info]

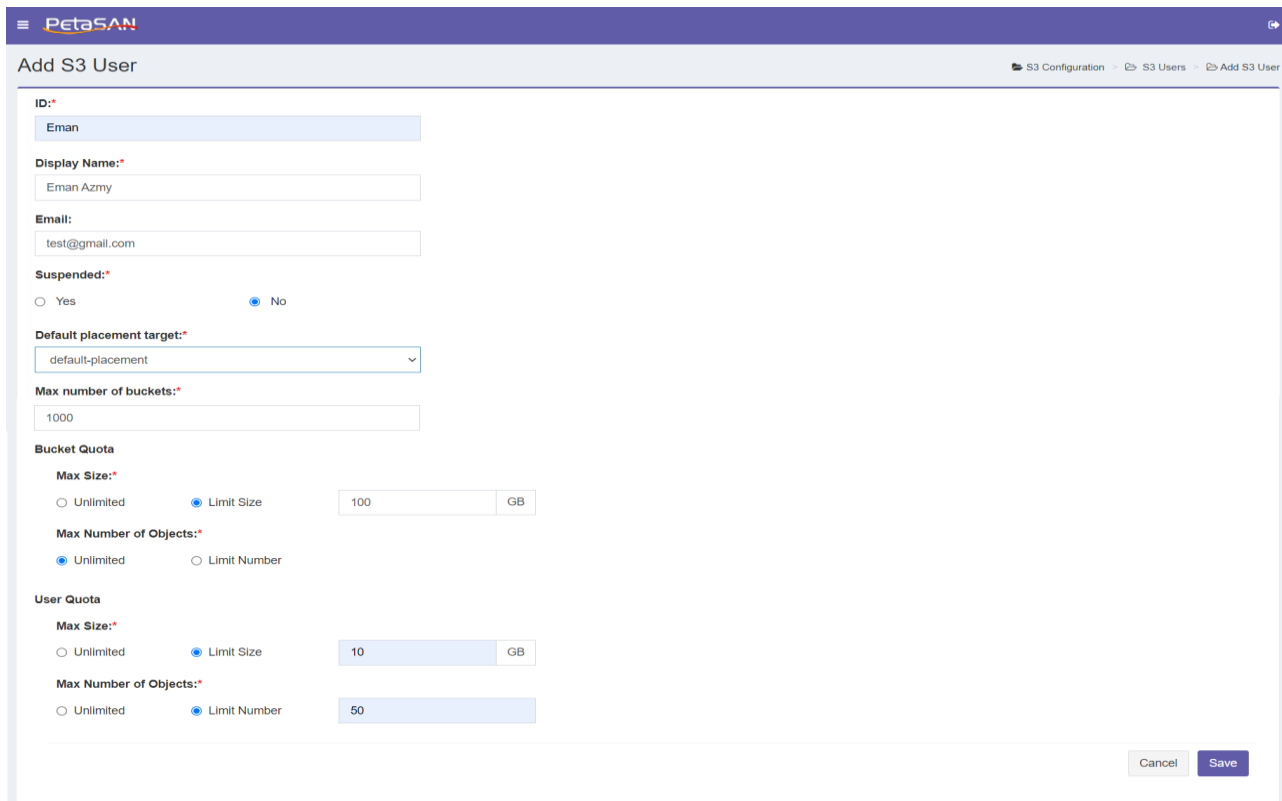
- **ID**  
The user ID
- **Display Name**  
The user name.
- **Suspended**  
The status of the user suspended or not.
- **Default Placement Target**  
The user's default placement target.
- **Size Used (GB)**  
The number of Gigabytes stored by the user.
- **Num of objects**  
The number of objects stored by the user.
- **Num of Subusers**  
The number of subusers defined under the user.
- **Actions**
  - Add  
Opens the S3 User form (this shown only in case you are in the master zone).
  - Edit  
Opens the Edit S3 User form (this shown only in case you are in the master zone).
  - Delete  
Deletes the selected user (this shown only in case you are in the master zone).
  - Info  
View the user information.

Note:

In case of multi zone, you will have Synchronization user that is created by the system and you can't be edited or deleted.

### ➤ Add S3 User

The system opens the S3 User form



The S3 User form contains the following fields:

- **ID**  
The user id, it must be unique.
- **Display Name**  
The user name.
- **Email**  
The user's email address.
- **Suspended**  
If the user is suspended or not.
- **Default Placement Target**  
The user's default placement target
- **Max number of buckets**  
The max number of buckets can be added by the user.
- **Bucket Quota**  
You can select the following quota for the user
  - Max Size  
Specify if the user can have unlimited or limited bucket size, in case of limited you must specify the size in GBs.
  - Max Number of Objects

Specify if the user can have unlimited or limited number of objects, in case of limited you must specify the number of objects.

- **User Quota**

You can select the following quota for the user

- Max Size

Specify if the user can have unlimited or limited bucket size, in case of limited you must specify the size in GBs.

- Max Number of Objects

Specify if the user can have unlimited or limited number of objects, in case of limited you must specify the number of objects.

- **Actions**

- Save

Saves the S3 users and redirects you to the Edit S3 User form where you can view the user's Access and secret keys.

- **Edit S3 User**

The system opens the Edit S3 User form, enable you to update all the user information except the user ID and can regenerate the keys if needed.

**Edit S3 User** S3 Configuration > S3 Users > Edit S3 User

**ID:**  
Eman

**Display Name:**  
Eman Azmy

**Email:**  
test@gmail.com

**Suspended:**  
 Yes  No

**Default placement target:**  
default-placement

**Max number of buckets:**  
1000

**Bucket Quota**

**Max Size:**  
 Unlimited  Limit Size

**Max Number of Objects:**  
 Unlimited  Limit Number

**Subusers:**  
[+ Add Subuser](#)

Show  entries

Subuser ID	Access Rights	Access Key	Secret Key	Actions
new	Read	2LZFUSIYEF4O85VFOYIU	BDOppzsyEyyZ0waRglghqGbLFG40daBsl5ZkduLC	<a href="#">✎</a> <a href="#">✕</a>
test	Full Control	1ZEPDYR83ZY4VL2FZ99C	3guk4s7NCdQWZYIDIOcgimZ8Ppt8t3YSnjtFBUCW	<a href="#">✎</a> <a href="#">✕</a>

Showing 1 to 2 of 2 entries

**User Quota**

**Max Size:**  
 Unlimited  Limit Size

**Max Number of Objects:**  
 Unlimited  Limit Number

**Access Key ID:**  
L43G5TL8D4XB4MLKNAAH

**Secret Access Key:**  
alpA8R9Wr03rcUWfVv1S7bXkib1qkFVO1KT3U4iY

**Size used:**  
0.0 GB

**Number of objects:**  
0

User can add subusers under the user and when added system will view the following columns for each subuser:

- **Subuser ID**  
The subuser ID
- **Access Rights**  
The subuser access right.
- **Access Key**  
The subuser access key.
- **Secret Key**  
The subuser secret key.

### ➤ Add Subuser

The system opens the Add Subuser form and enables the administrator to enter the following fields:



- **Parent User**  
The name of the parent user.
- **ID**  
The Subuser id.
- **Access Right**  
The subuser's access right.

The screenshot shows a web form titled "Add Subuser". It contains three main input sections: "Parent User:" with a text box containing "Eman"; "ID:\*" with a text box containing "New"; and "Access Right:\*" with a dropdown menu. The dropdown menu is open, showing four options: "Full Control" (highlighted in blue), "Read", "Read/Write", and "Write". At the bottom right of the form, there are two buttons: "Cancel" and "Save".

After saving the subuser system will redirect the user to the Edit subuser form showing the subuser's keys

## ➤ **Edit Subuser**

The system opens the Edit Subuser form and shows following fields:

- **Parent User**  
The name of the parent user and can't be modified.
- **ID**  
The Subuser id and can't be modified.
- **Access Right**  
The subuser's access right.
- **Access Key**  
The subuser's access key.
- **Secret Key**  
The subuser's secret key.

Edit Subuser
X

---

**Parent User:**

**ID:\***

**Access Right:\***

**Access Key:\***

**Secret Key:\***

Note:  
You can regenerate the subuser's access and secret keys

## 9.2. S3 Status

### ➤ View S3 Status

You can view the current S3 status

Node Name	Status	Assigned IPs
Node1	Up	10.0.3.11 10.0.3.14 10.0.3.15
Node2	Up	10.0.3.10 10.0.3.12 10.0.3.13

The system will view the following columns:

- **Node Name**  
Name of the S3 Server Node
- **Status**

Node status, up, down

- **Assigned IPs**

IPs assigned to each node

### 9.3. Reassign IPs

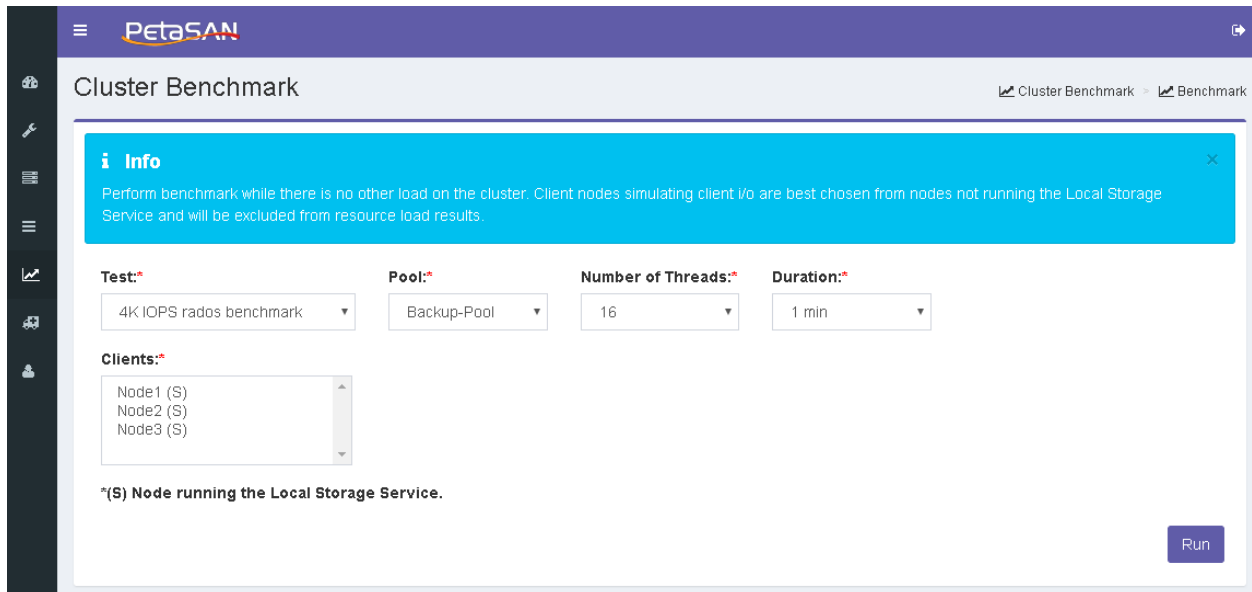
Opens the Reassign IPs form which enables you to move some selected IPs manually to another node or distribute them automatically

The screenshot displays the 'IP Reassignment List' interface. At the top, there are navigation links for 'Manage CIFS' and 'IP Reassignment List'. Below the title, there are radio buttons for 'Manual' (selected) and 'Auto'. The main content area is a table with two columns: 'From Node' and 'Assigned IPs'. The table lists two nodes: Node2 and Node3. Node2 is associated with 4 IPs: 10.0.3.10, 10.0.3.12, 10.0.3.14, and 10.0.3.16. Node3 is associated with 3 IPs: 10.0.3.11, 10.0.3.13, and 10.0.3.15. Below the table, there is a dropdown menu labeled 'Assign To:' with 'Node3' selected. A 'Reassign' button is located at the bottom right of the interface.

## 10. Cluster Benchmark

### ➤ Benchmark

You can use the cluster benchmark form to test the performance of your cluster. You should perform this benchmark while there is no other load on the cluster. Results will show the total cluster Throughput and IOPS as well as display the resource loads for all nodes that have Local Storage Role



You can customize your test using the following fields:

- **Test**

Select one of the following test types:

- 4K IOPS rados benchmark  
Runs the 4K rados test to benchmark total cluster IOPS.
- 4M Throughput rados benchmark  
Runs the 4M rados test to benchmark the total cluster Throughput

- **Pool**

Select the pool you want to perform the test on from the list of active pools.

- **Number of Threads**

Number of threads/queue depth that will be used in the test

- **Duration**

The test duration in minutes

- **Clients**

Select the nodes that will be acting as clients simulating client IO. Client nodes will be excluded from the resource load results.

➤ **Notes:**

- *For more accurate results, client nodes are best chosen from nodes not running the Local Storage Service.*
- *Nodes that have Local Storage Role will have (s) following their names.*  
The system will display the benchmarking results as follows

## Results

### Cluster IOPS

Write	Read
354	205

### Write Resource Load:

Node	Memory Util%	CPU Util%		Network Util%		Disks Util%		Actions
		Avg	Max	Avg	Max	Avg	Max	
Node1	51	59	59	0	0	3	3	Show Details

### Read Resource Load:

Node	Memory Util%	CPU Util%		Network Util%		Disks Util%		Actions
		Avg	Max	Avg	Max	Avg	Max	
Node1	49	99	99	0	0	100	100	Show Details

- **Cluster IOPS/Throughput**

Display IOPS/Throughput results depending on type of test chosen.

- Read
  - Total cluster Read IOPS/Throughput.
- Write
  - Total cluster Write IOPS/Throughput.

- **Write and Read Resource Load**

Displays resource load for storage nodes during both read and write operations.

- Node Name
  - The name of the storage node
- Memory Utilization
  - This is a measure of how much RAM/Memory was used.
- CPU Utilization
  - This is a measure of how much the CPU cores were busy.
    - Avg
      - Average utilization for all cores.
    - Max
      - Highest utilization used by an individual core.
- Network Utilization
  - This is a measure of how much the network interfaces were busy.
    - Avg
      - Average utilization for all network interface cards.
    - Max
      - Highest utilization used by an individual network interface cards.
- Disk Utilization
  - This is a measure of how much the disks were busy.

- Avg  
Average utilization for all disks.
- Max  
Highest utilization used by an individual disk.

You can view the resource load for any node in more details by clicking the Show Details button

#### Node1 Write Resource Load Details ×

##### Memory

kbmemfree	kbmemused	kbbuffers	kbcached	kbcommit	%commit	kbactive	kbinact	kbdirty	%util
745724	783532	1088	118024	4429960	289	553176	68416	560	51

##### CPUs

CPU	%user	%nice	%system	%iowait	%steal	%util
all	5	0	7	46	0	60
0	5	0	7	46	0	60

##### Network Interfaces

Interface	rxpck/s	txpck/s	rxkB/s	txkB/s	rxcmp/s	txcmp/s	rxmcast/s	%util
eth0	4	4	0	0	0	0	0	0
eth3	126	159	27	51	0	0	0	0
eth4	3	3	0	0	0	0	0	0

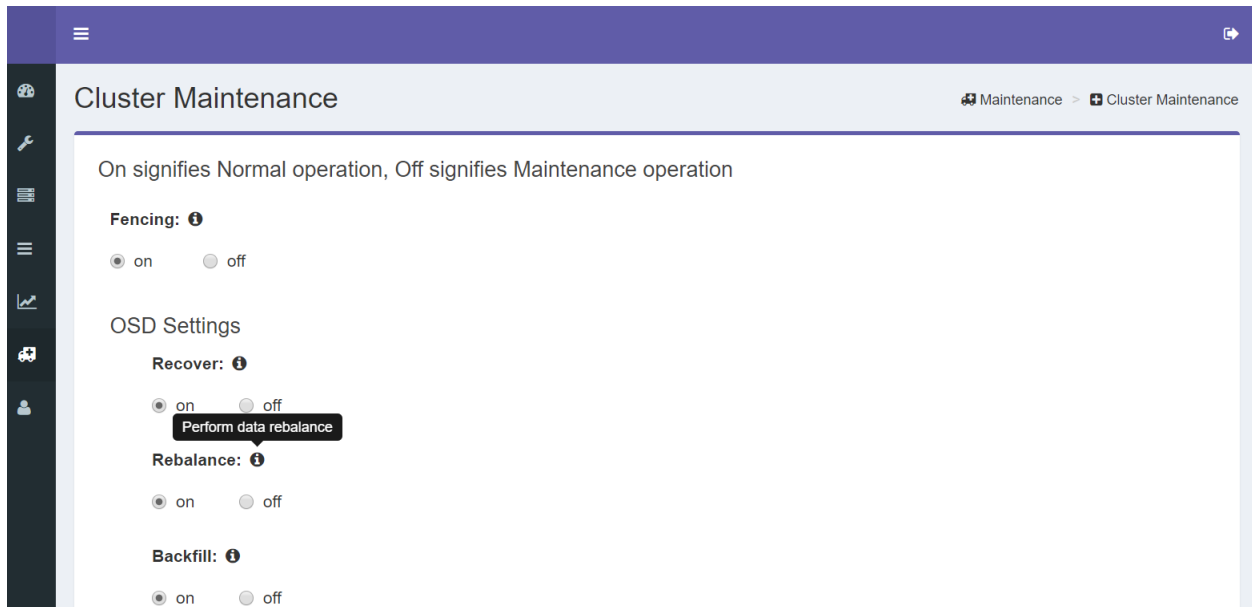
##### Disks

Disk	tps	rd_sec/s	wr_sec/s	avgrq-sz	avgqu-sz	await	svctm	%util
sdb	6	185	848	170	0	16	6	3

## 11. Maintenance

### 11.1. Cluster Maintenance

You can turn off one or more of the maintenance settings using the cluster maintenance form

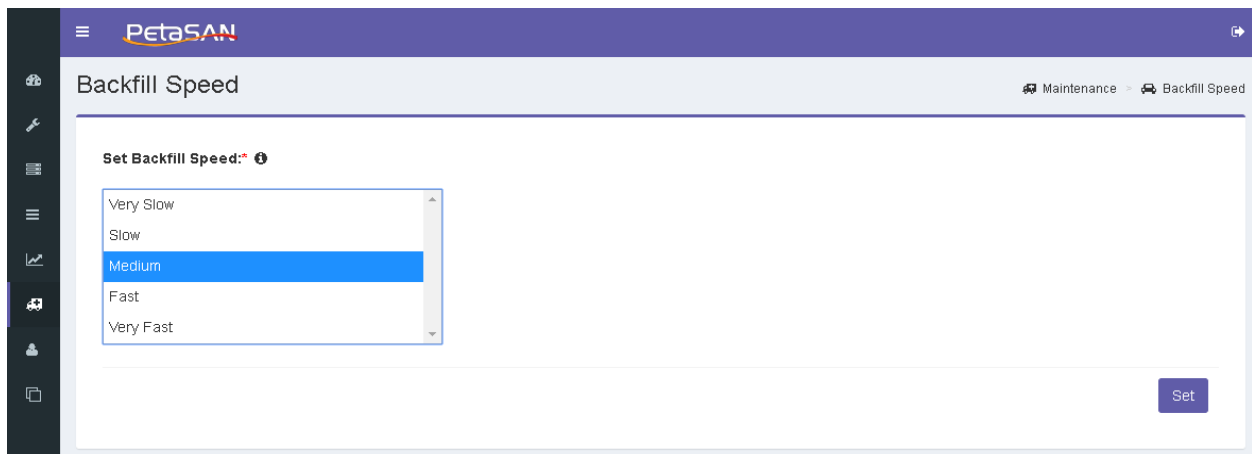


➤ **Notes:**

- *If any of the maintenance settings is turned off, the system will show that maintenance is on in the Dashboard.*

## 11.2. Backfill Speed

You can control the speed of background data recovery and rebalancing using the Backfill Speed form



## 11.3. Scrub Speed

You can control the speed of background scrub operations using the Scrub Speed form

### 11.4. Ceph Balancer

You can enable background fine tuning of PG assignments among OSDs to achieve better data distribution via the Ceph Balancer form.

### 11.5. OSD CRUSH Weight

You can adjust the individual weight of OSDs, measured in TB, via the OSD CRUSH Weight form. This can provide manual fine tuning of PG assignments among OSDs to achieve better data distribution and can also be used to gradually increase weight of newly added OSDs for better control of rebalance traffic.

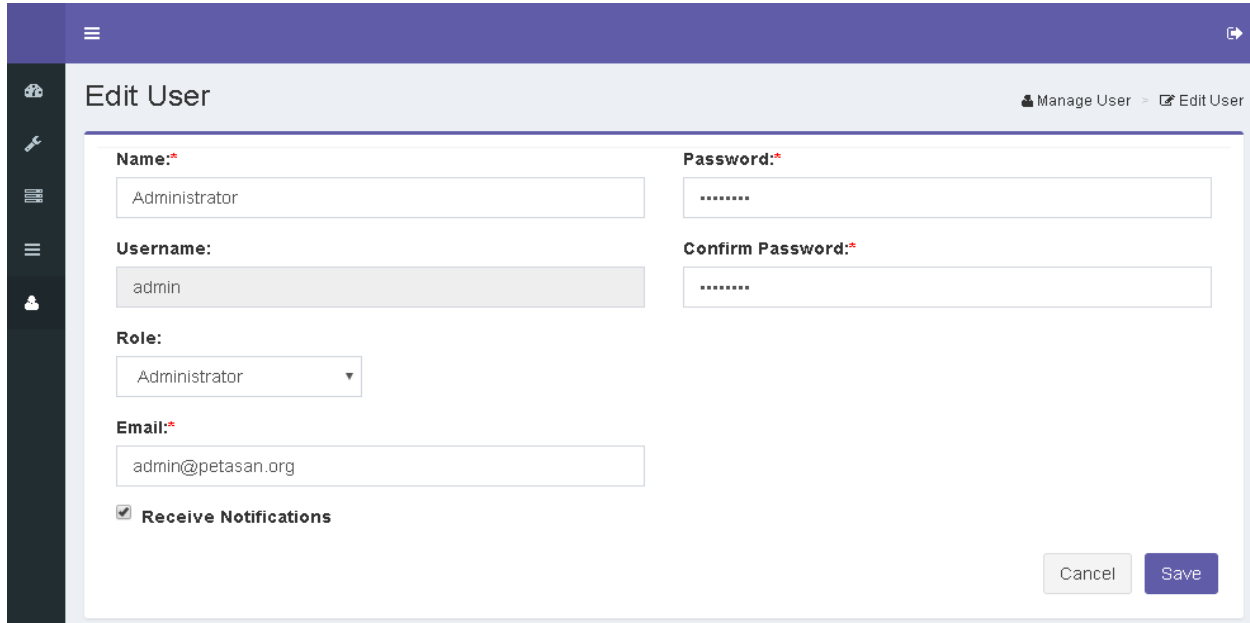
OSD Name	Node Name	% Usage	Weight
osd.0	Node3	0%	0.0977
osd.1	Node1	1%	0.0977
osd.2	Node2	1%	0.0977



## 12. Manage Users

### ➤ Add User

You can add new user by using the Add User Form



The screenshot shows the 'Edit User' form in the PetaSAN Administrator interface. The form is titled 'Edit User' and has a breadcrumb trail 'Manage User > Edit User'. The form contains the following fields and options:

- Name:** A text input field containing 'Administrator'.
- Username:** A text input field containing 'admin'.
- Password:** A password input field with masked characters '.....'.
- Confirm Password:** A password input field with masked characters '.....'.
- Role:** A dropdown menu with 'Administrator' selected.
- Email:** A text input field containing 'admin@petasan.org'.
- Receive Notifications:** A checked checkbox.
- Buttons:** 'Cancel' and 'Save' buttons at the bottom right.

The form requires the following information:

- **Name**  
The name of the user
- **Username**  
The username that will be used to login
- **Password/confirm password**  
The user password and its confirmation
- **Role**  
The user role which can be:

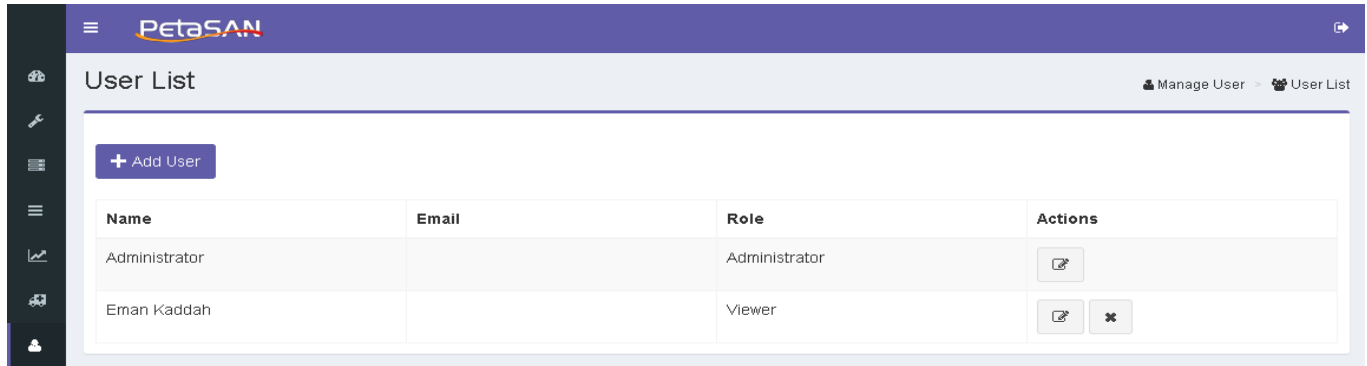
“Administrator” this enables the user to access all the system pages

“Viewer” this enables the user to monitor the cluster using the Dashboard.

- **Email**  
If you want the user to receive email notification then you should enter the user’s email and check the Receive notification option.

### ➤ User List

The system enables you to view all the users in the system using the User List Page

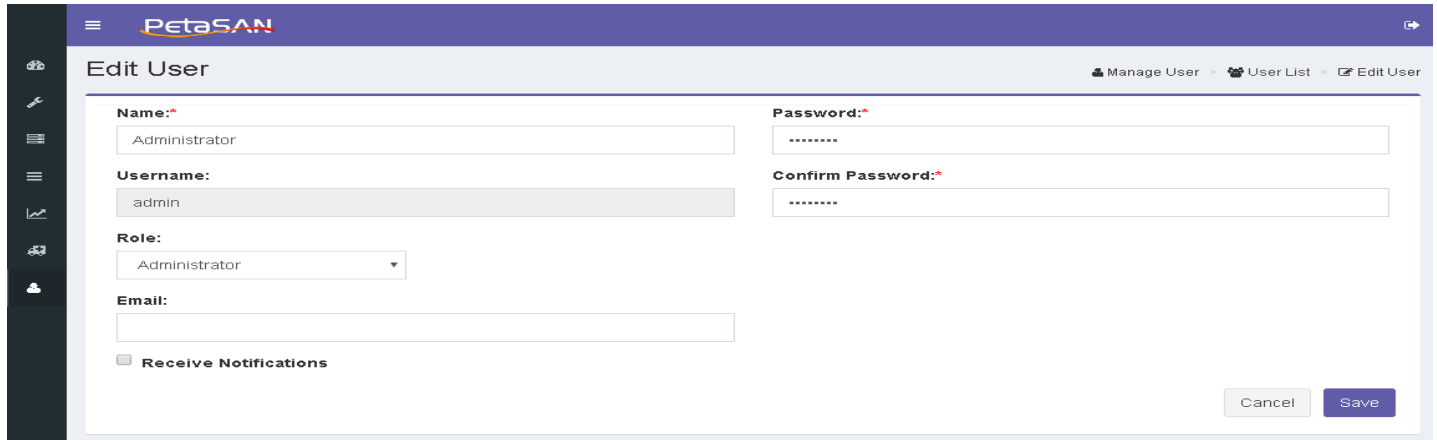


- **Actions**

The system enables you to do the following actions for each user:

- **Edit**

Enables updating the user information except the username.



- **Delete**

Enables you to delete any user except the default administrator

➤ **Change Password**

The system allows any user to change his password using the Change Password Form

