Administrator Guide

Version 1.8

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

Contents

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.

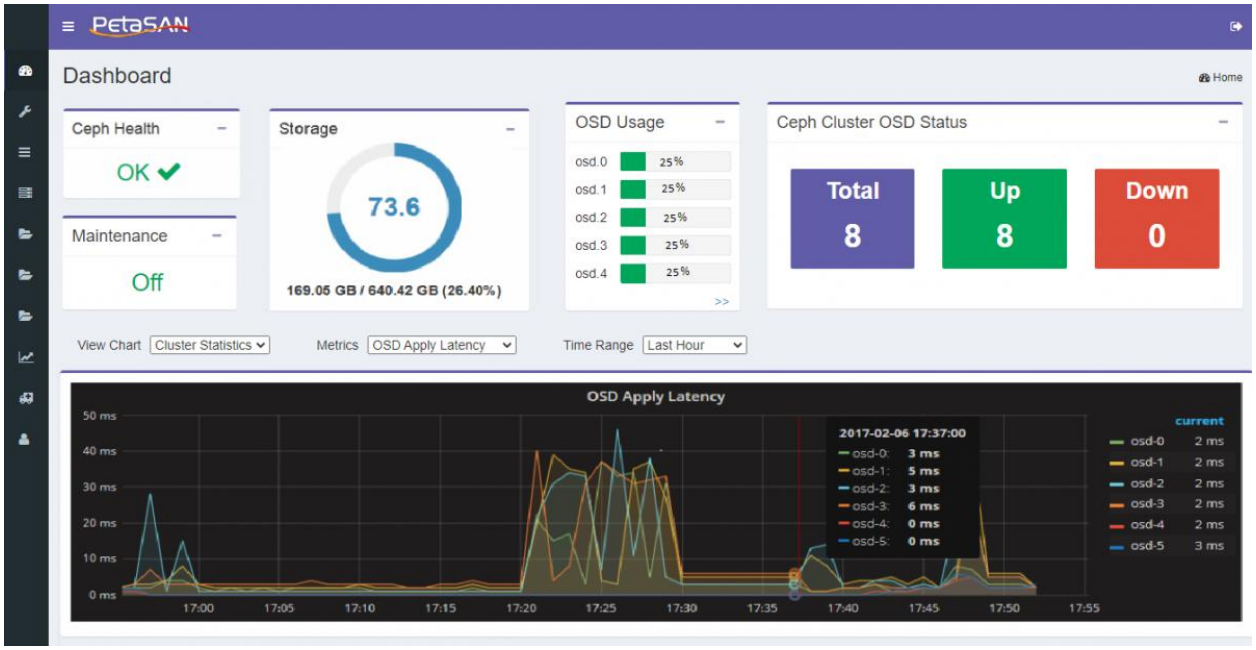

➢ Ceph Health

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

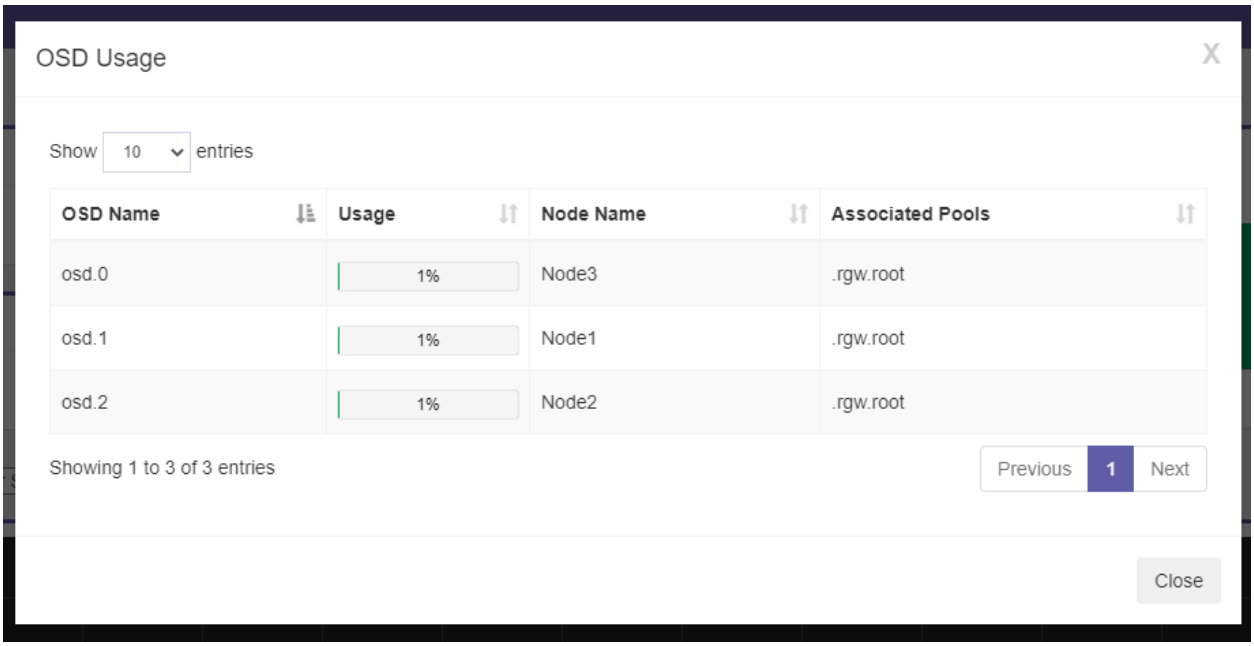

➢ Ceph OSD Status

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.

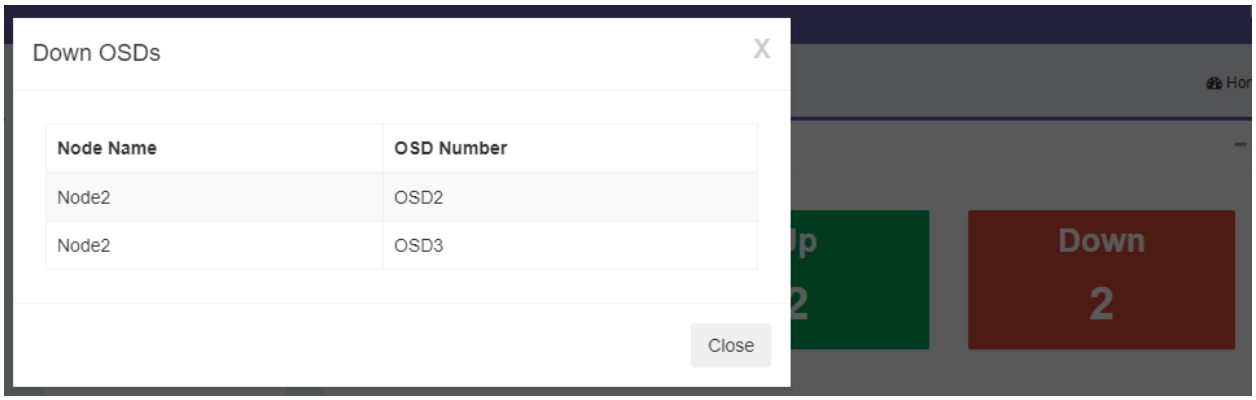

➢ View Chart

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:

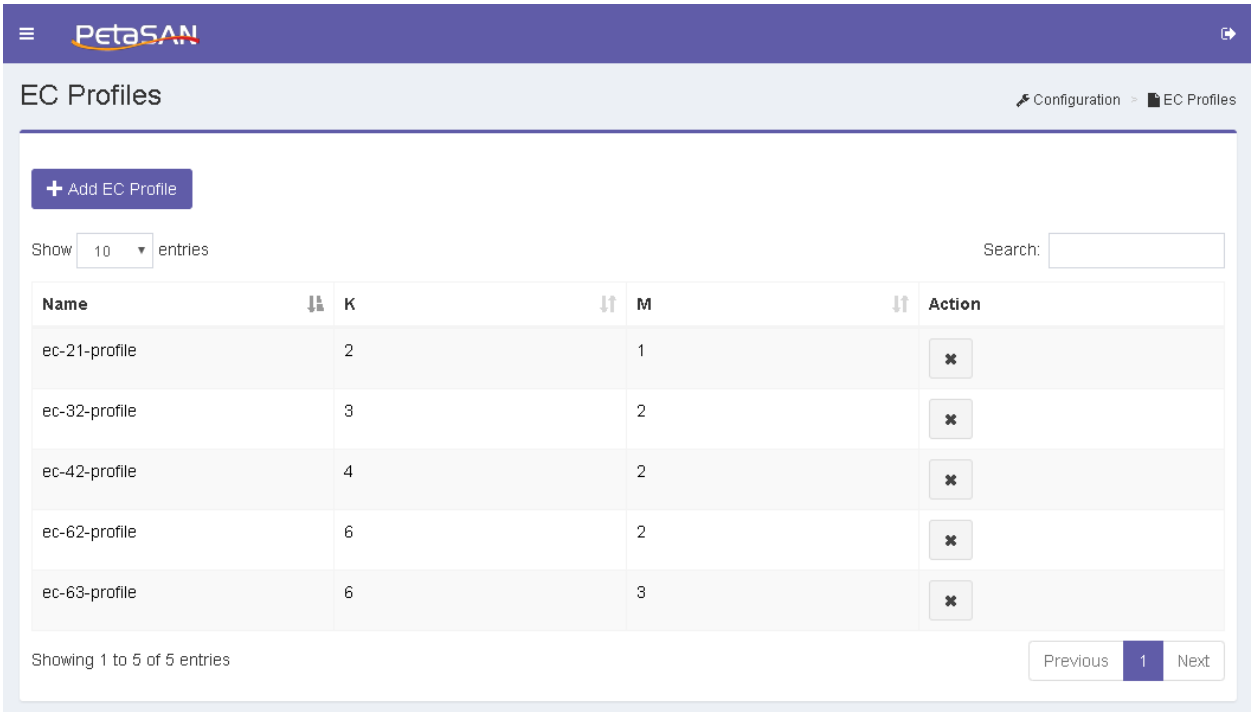

The EC Profiles List displays the following information for each Profile:

- **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

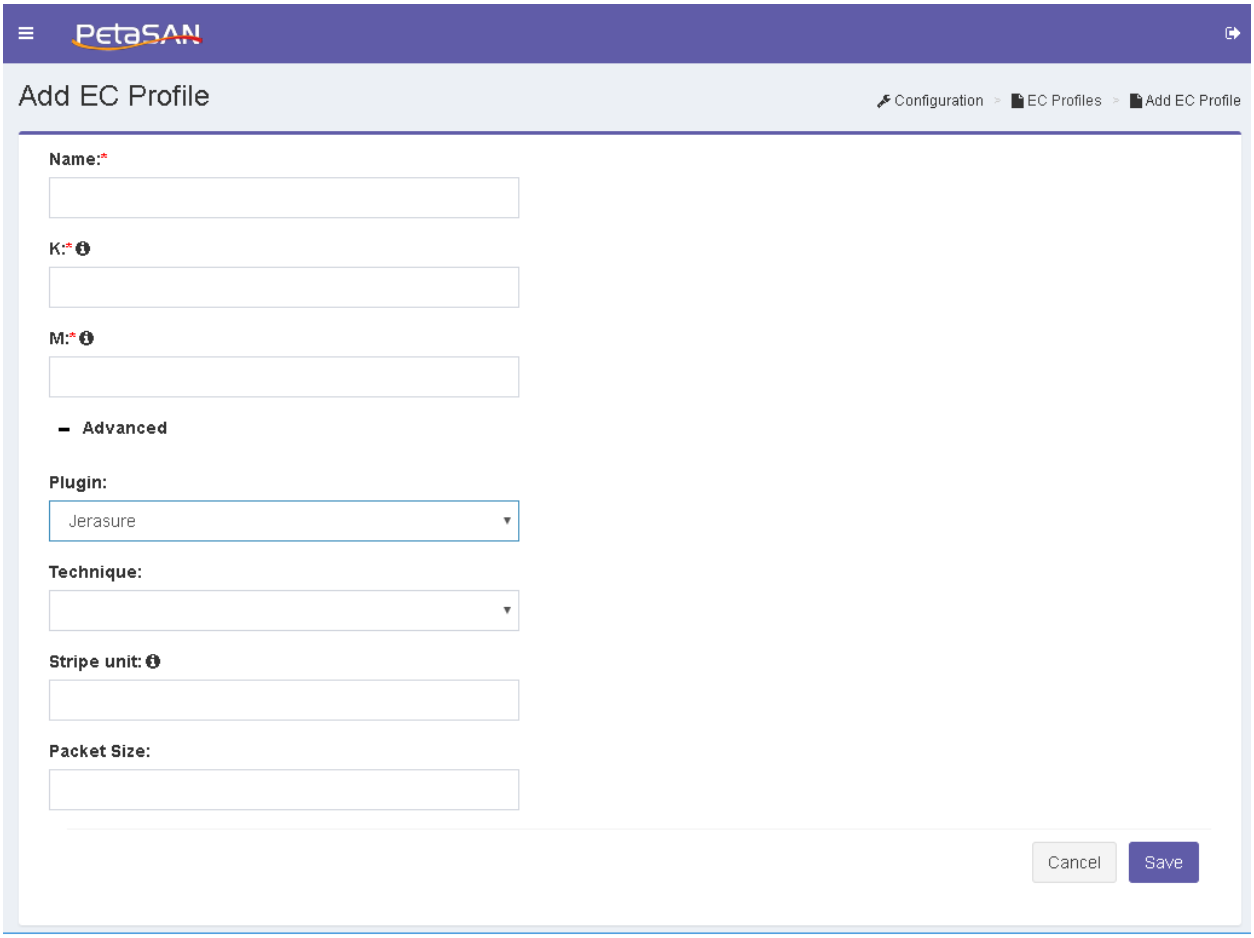

The Add EC Profile from contains the following fields:

- **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 respectivly.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.

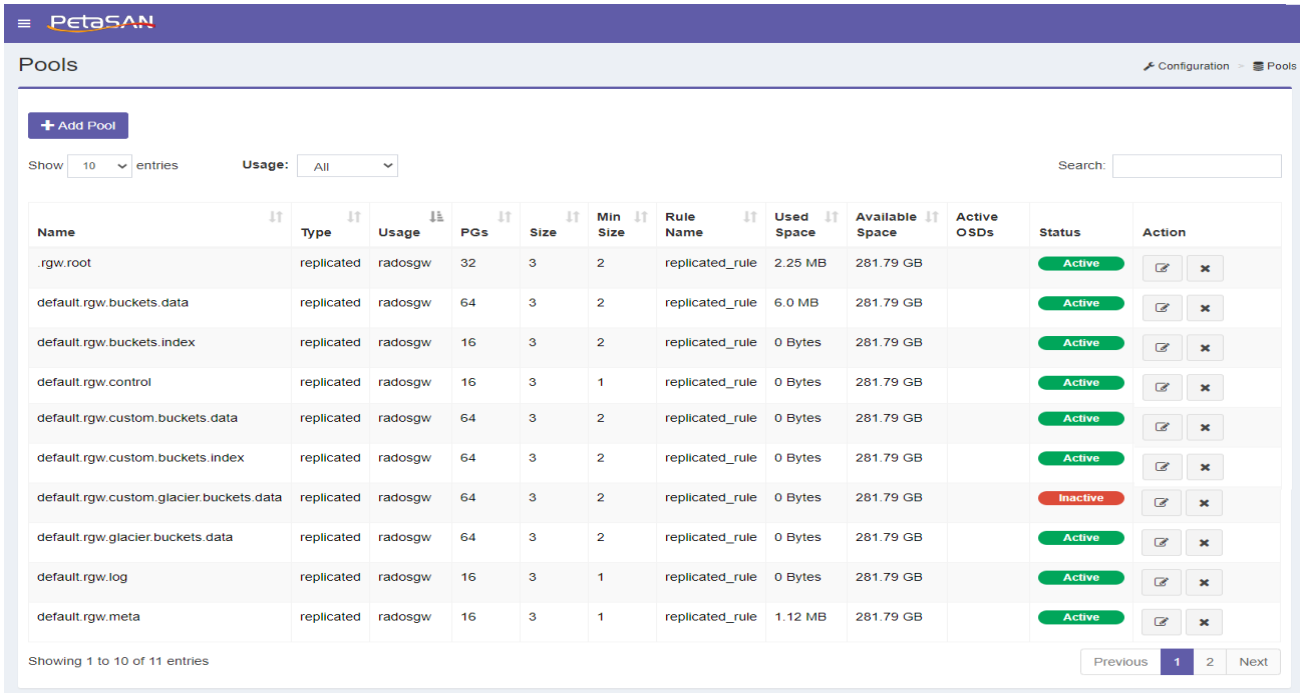

The Pools List displays the following information for each Pool:

- **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

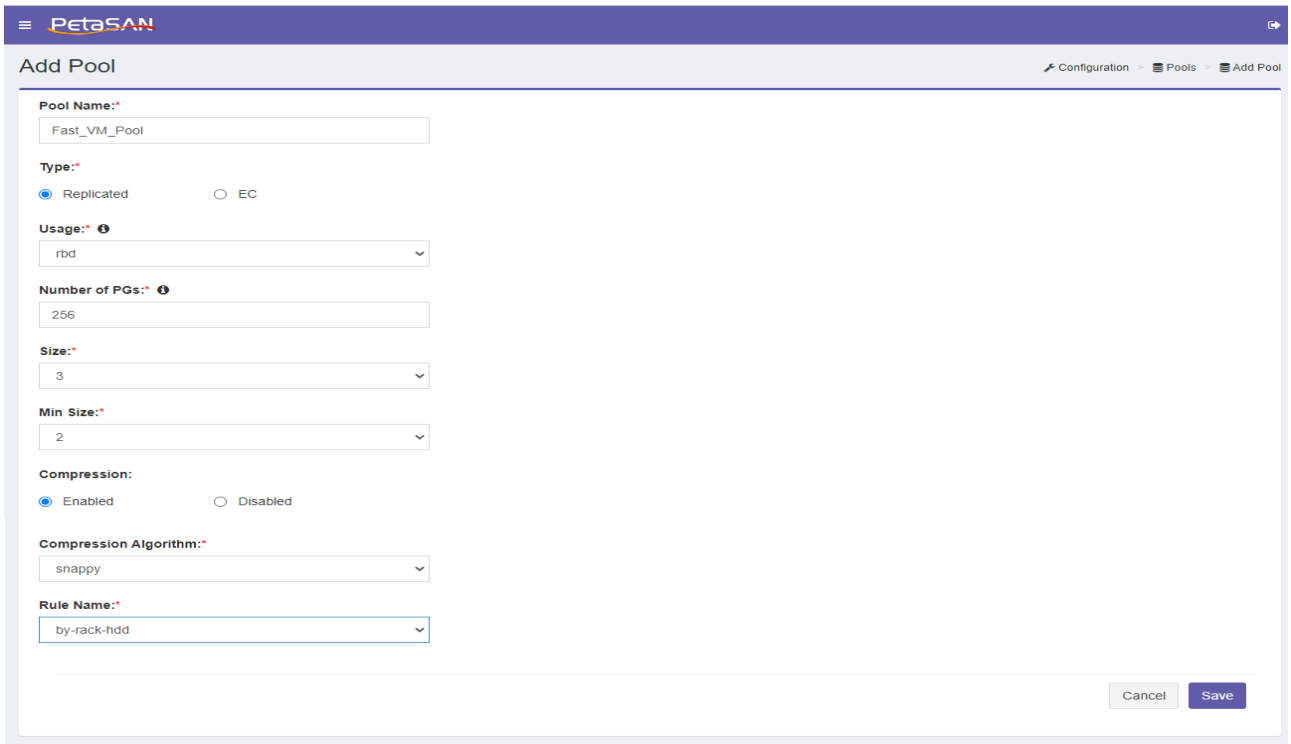

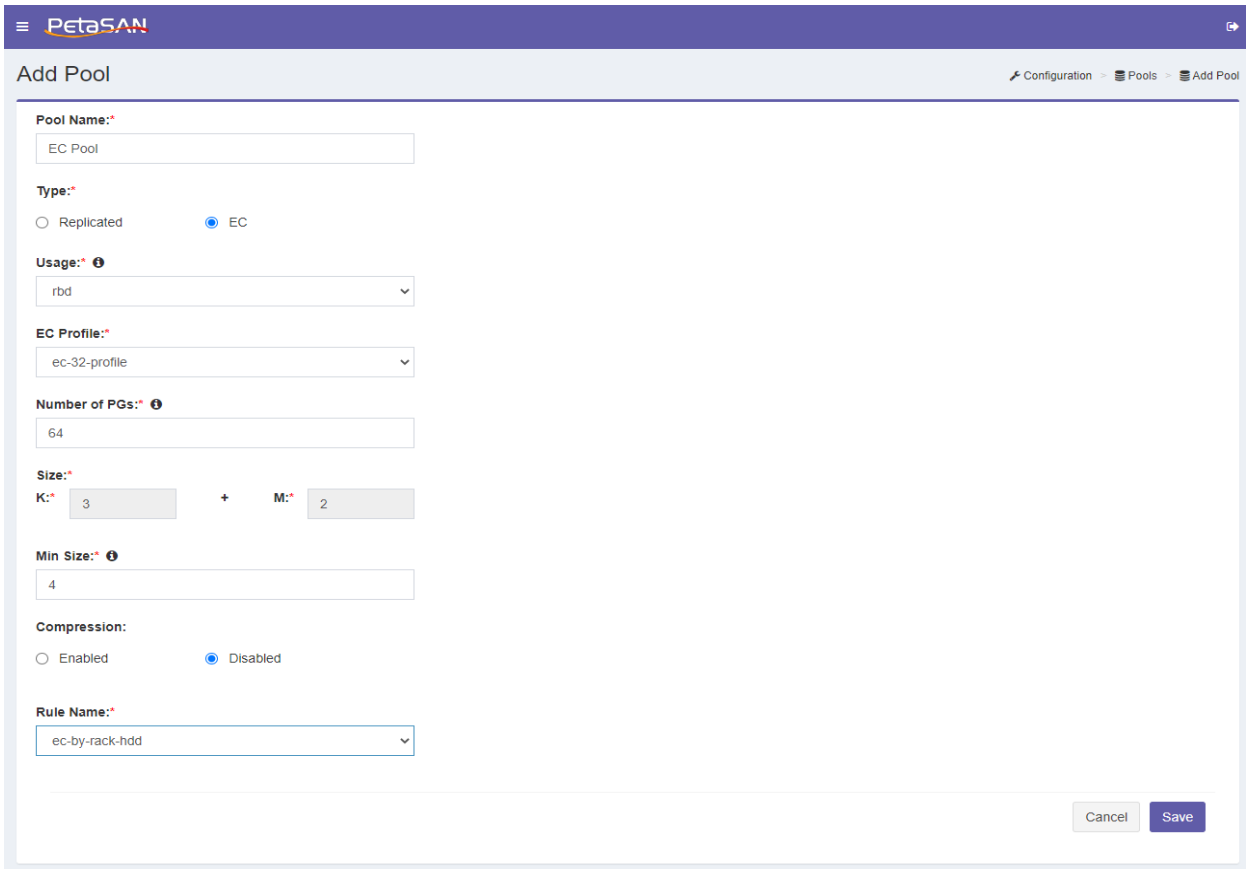

The Add Pool from contains the following fields:

- **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

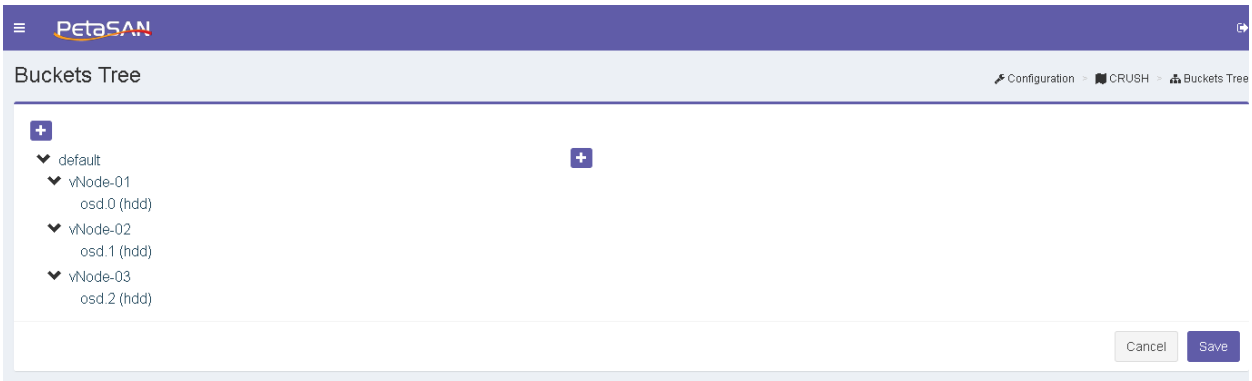

You will be able to manage buckets as following:

➢ 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:

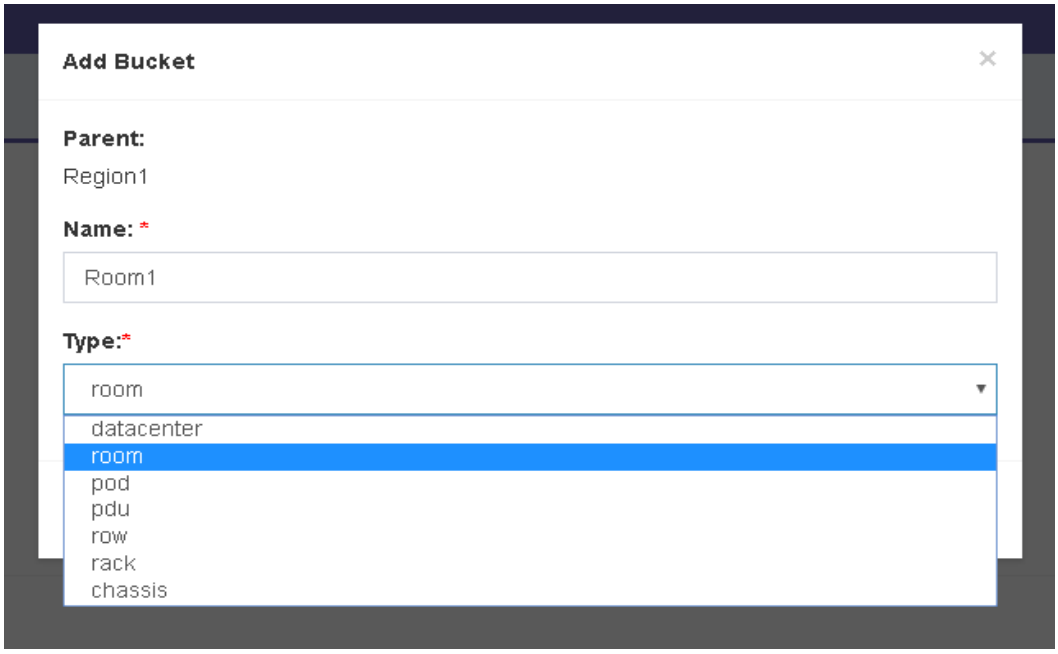

- **Parent**
  The name of the selected parent bucket.
- **Name**
  The new bucket name
- **Type**
  The new bucket type, you will be able to select types that are allowed Childs of the parent bucket, after entering the bucket information and clicking Add, the bucket will be added to the temporary buckets 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)

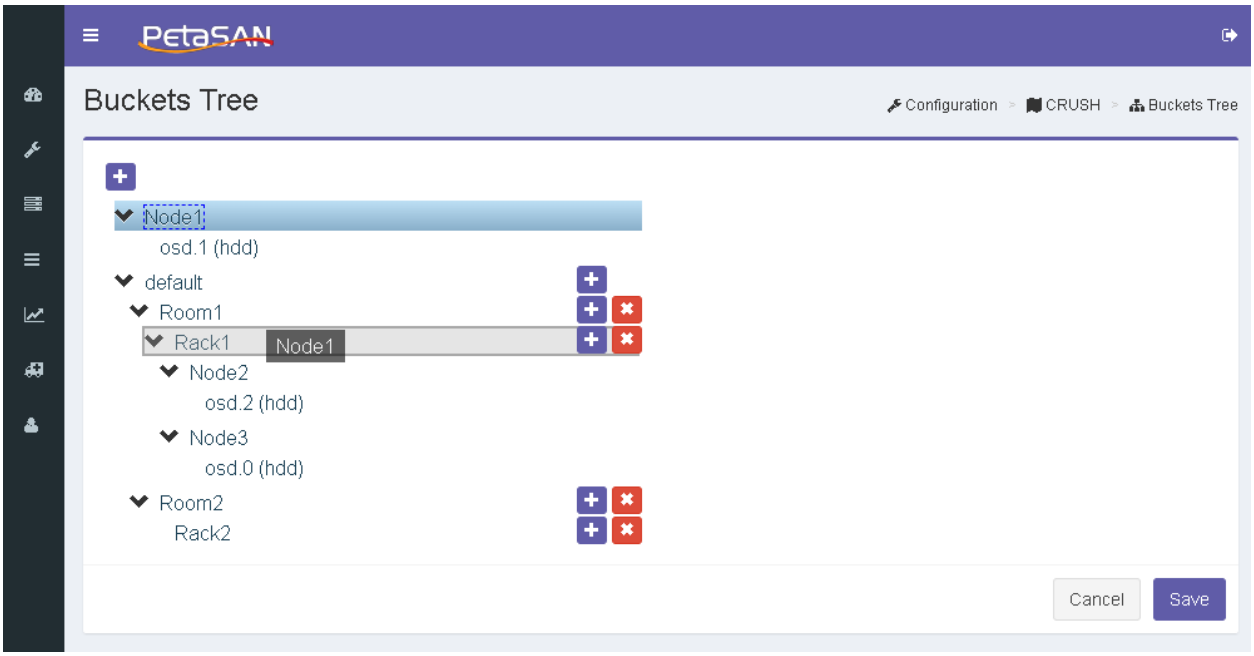

Any movements will not be saved to the CRUSH map until you click Save on the Buckets Tree form.

- **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

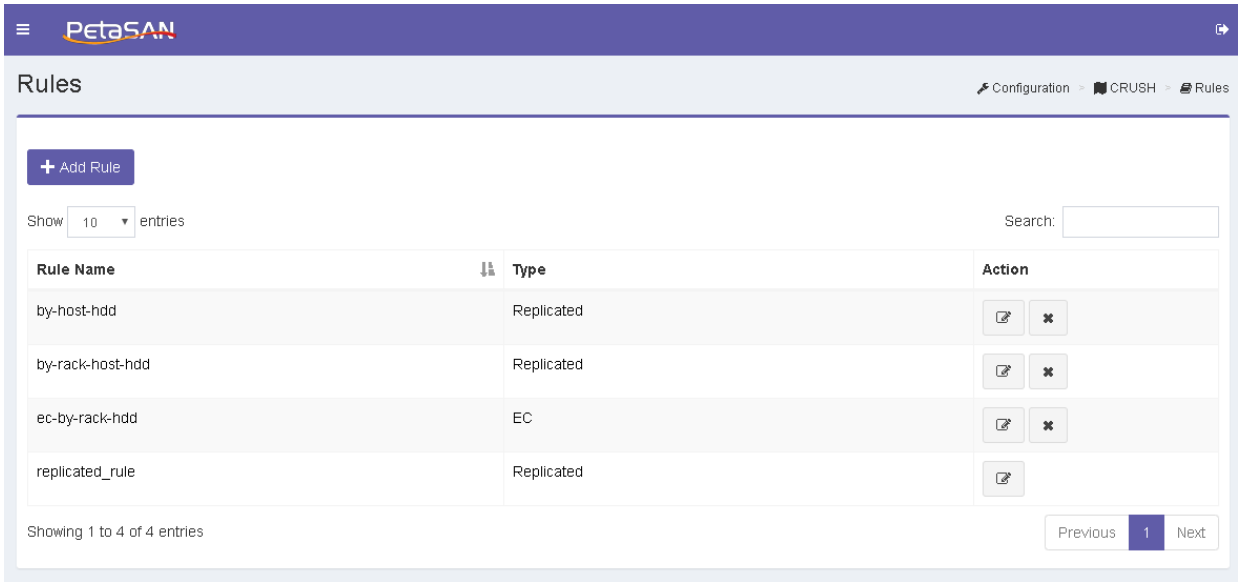

The Rules List displays the following information for each Rule:

- **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

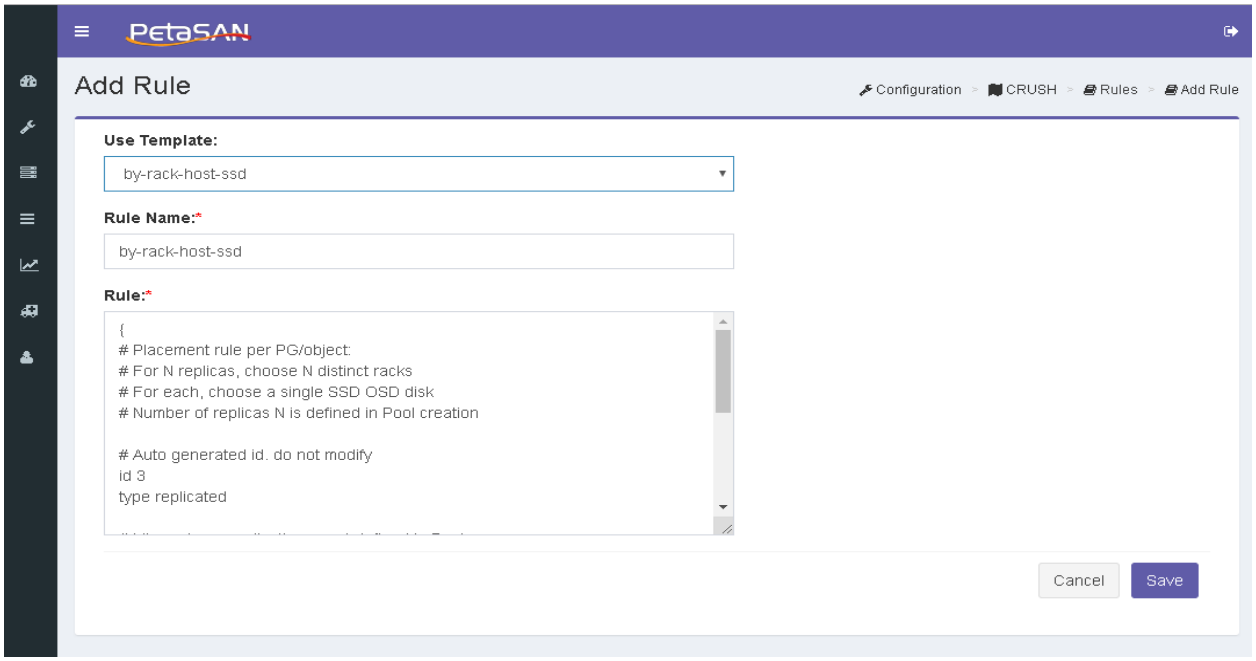

The Add Rule from contains the following fields:

- **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

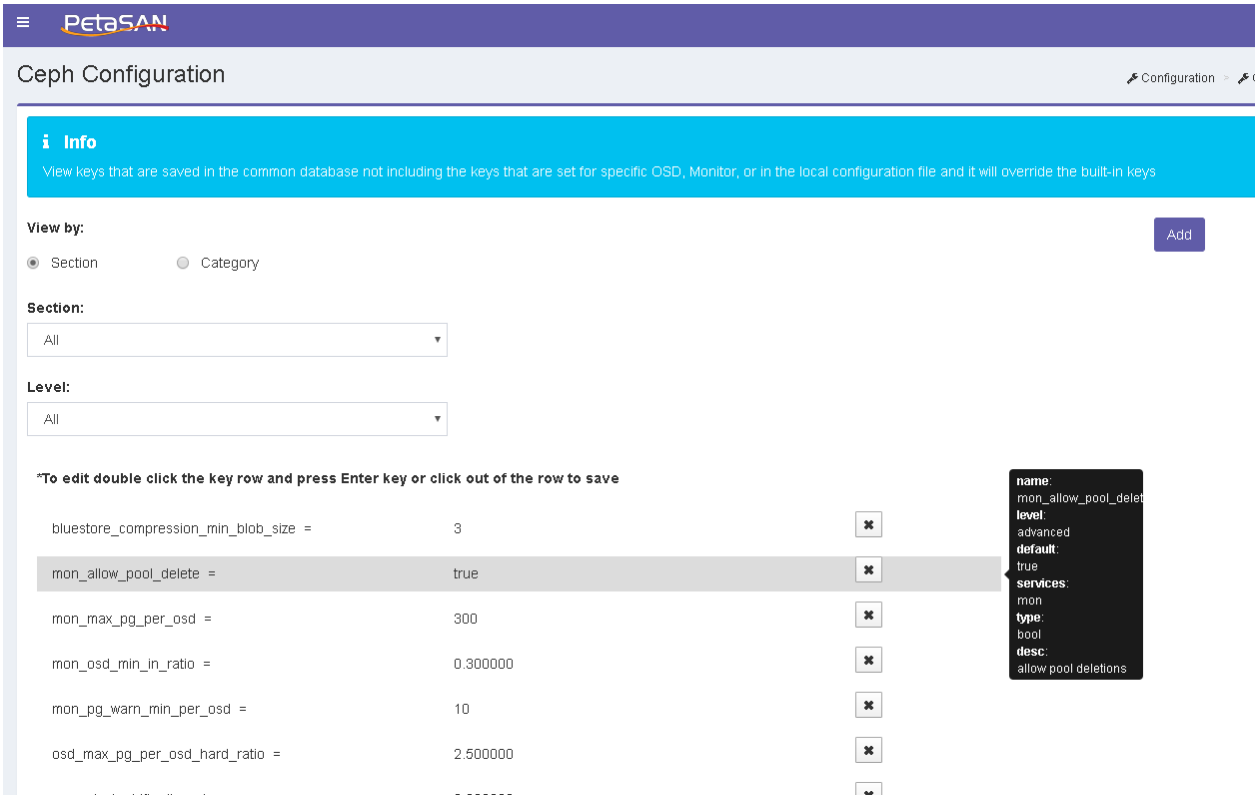

You can filter keys by:

- **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

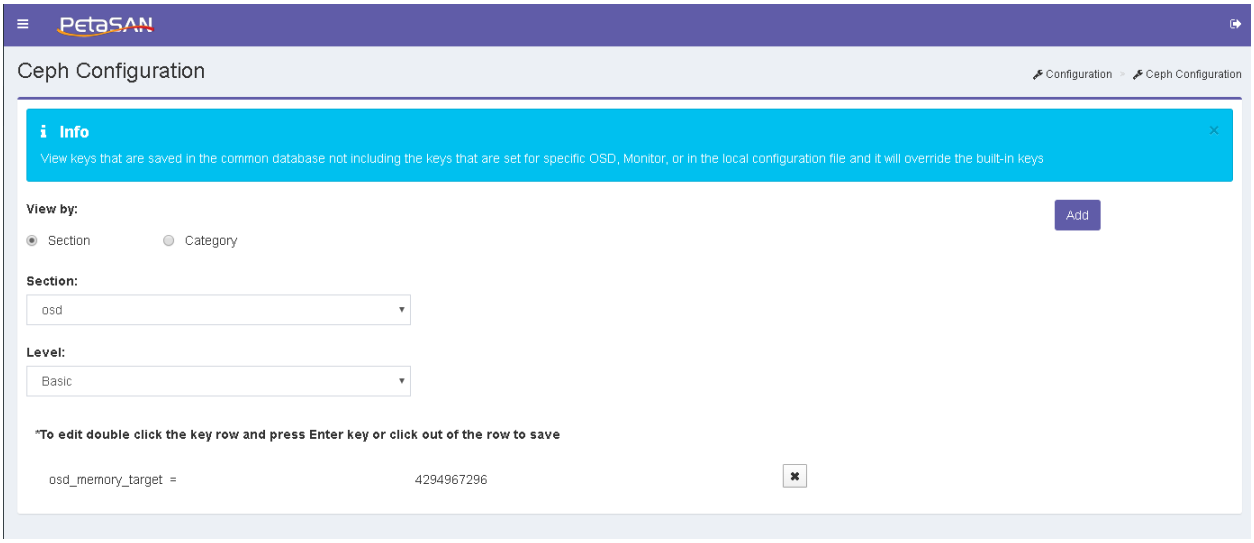

- **Category**

  The current categories options are
  - Recovery
  - Scrubbing

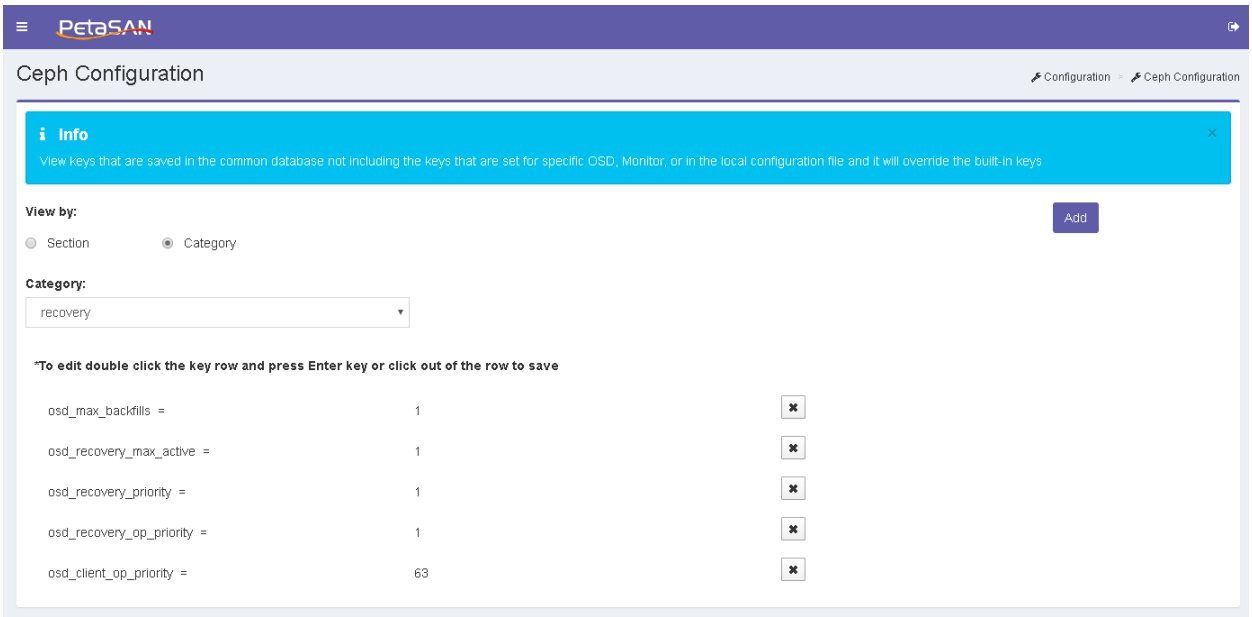

➢ Add Ceph Configuration Key

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

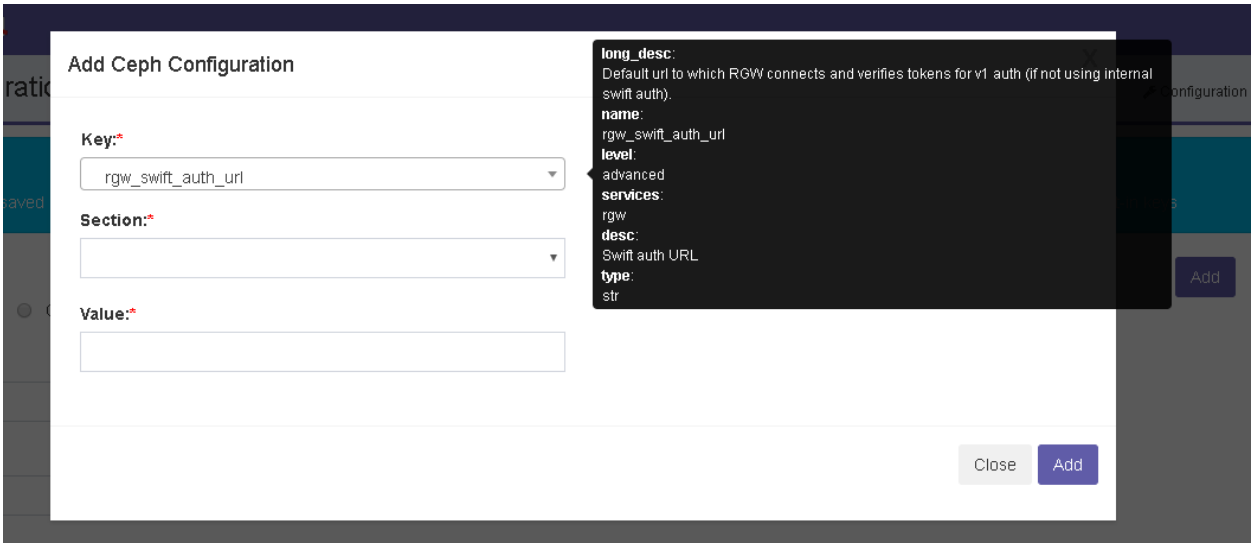

➢ Delete Ceph Configuration Key Value

You can delete one of the existing Ceph configuration keys

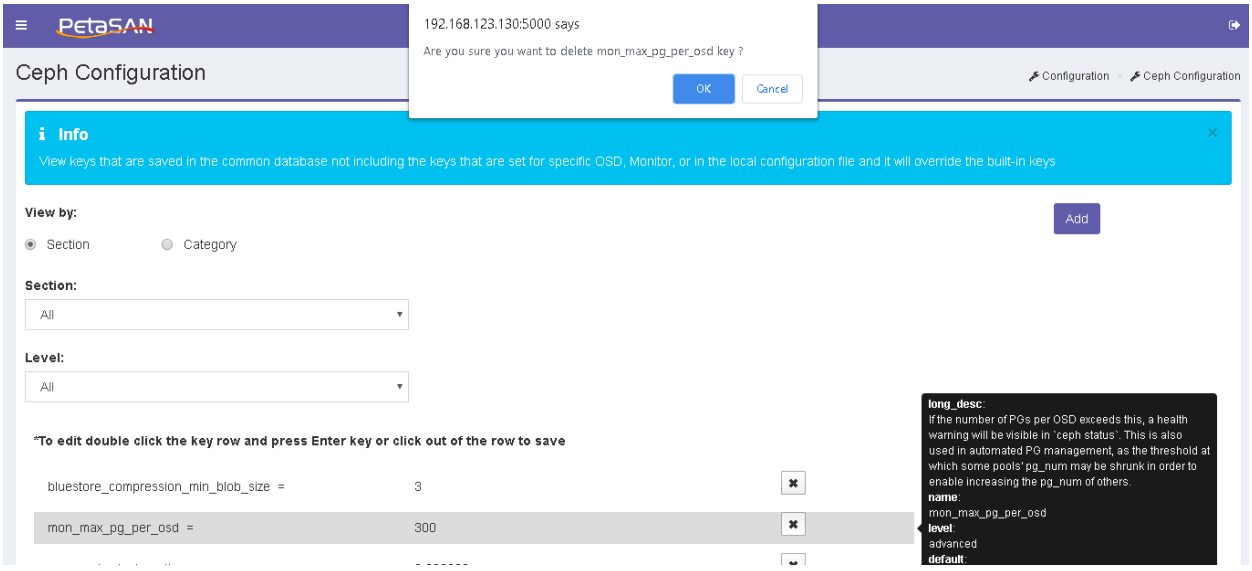

4.6. NUMA Pinning

➢ View NUMA Pinning

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

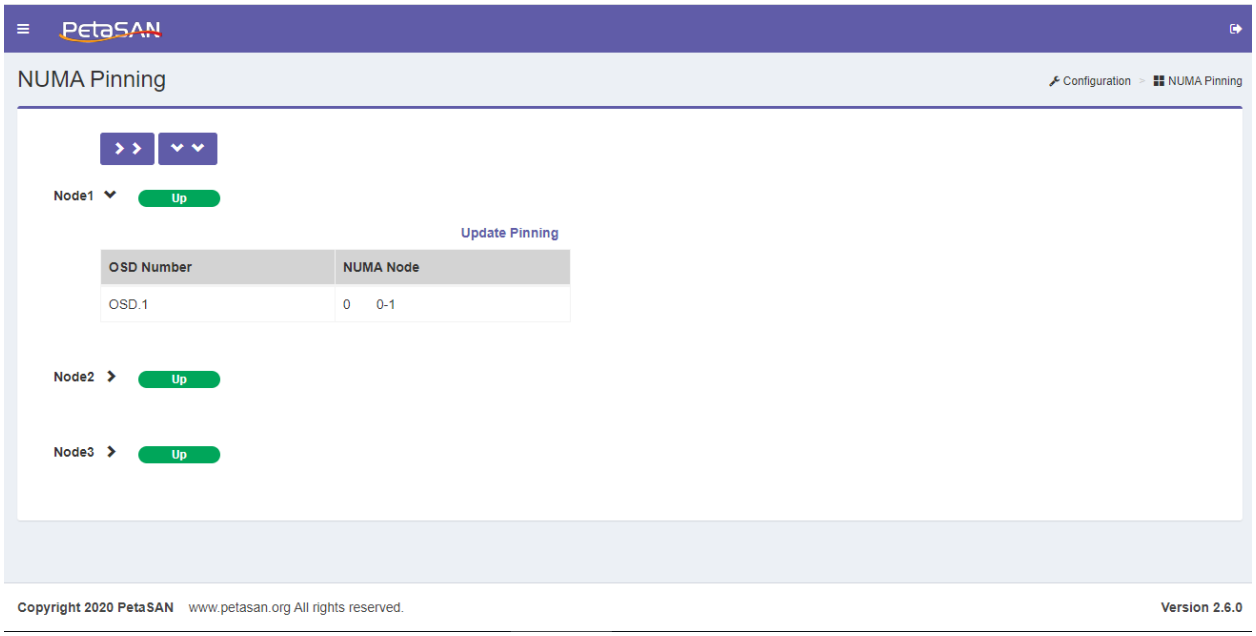

➢ Update Pinning

You can update the pinning for OSDs on specific NUMA nodes

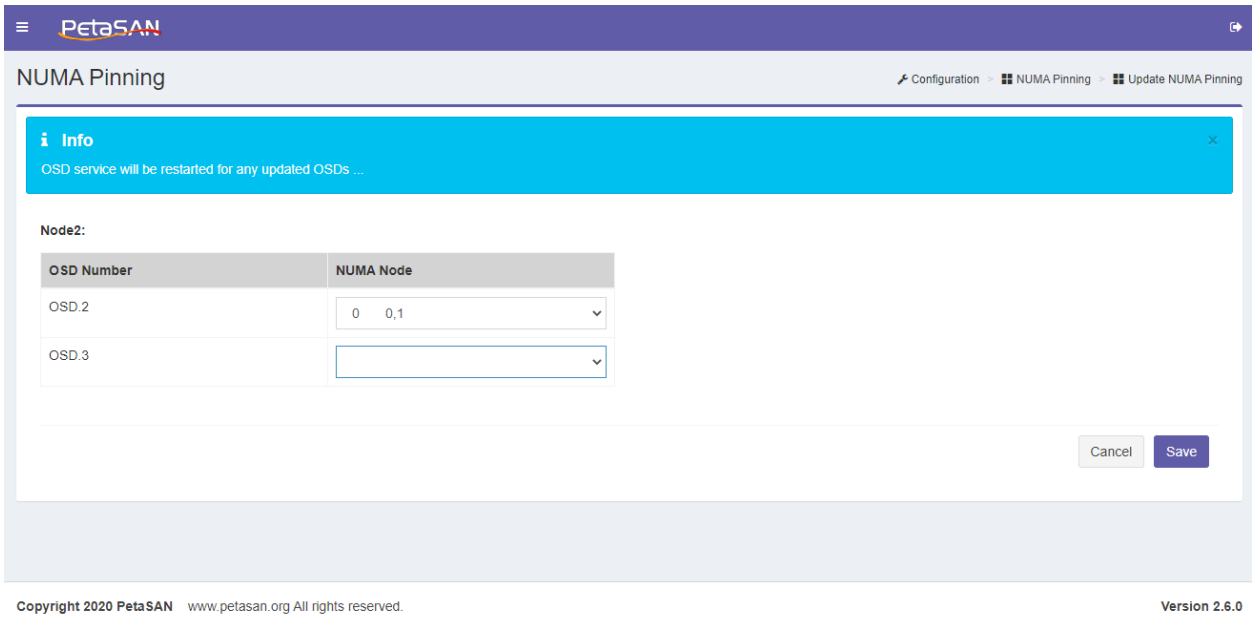

4.7. File Systems

➢ View File Systems

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

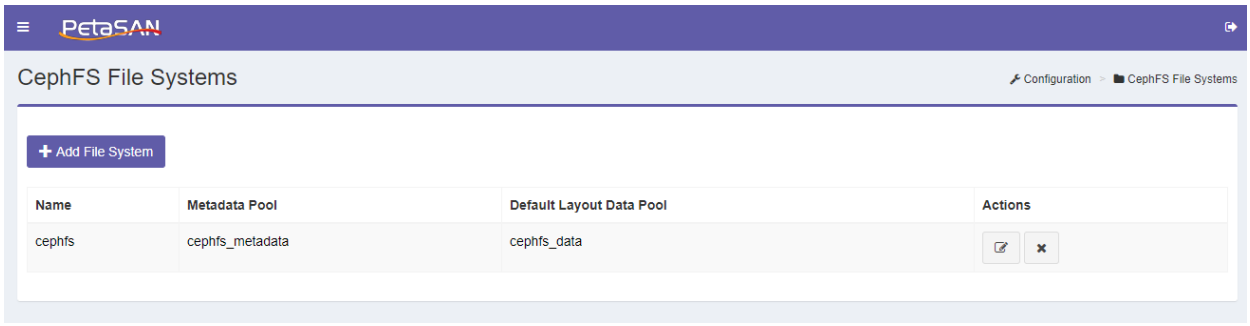

System views the following columns for the file systems

- **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 cehpfs

  - Save

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

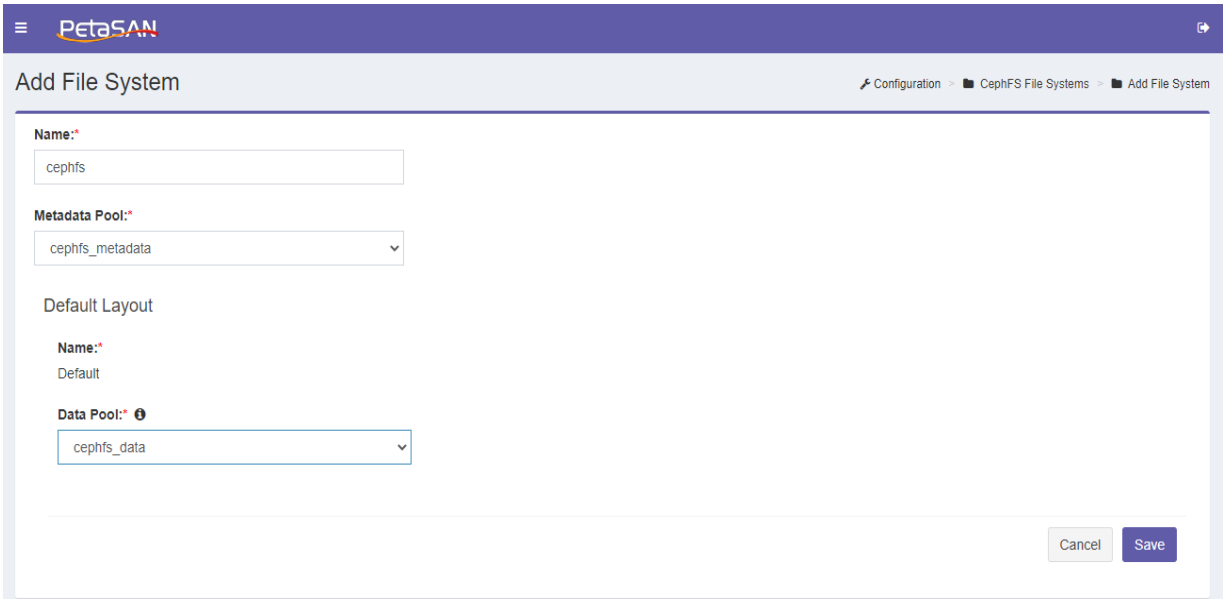

➢ Edit File System

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 .

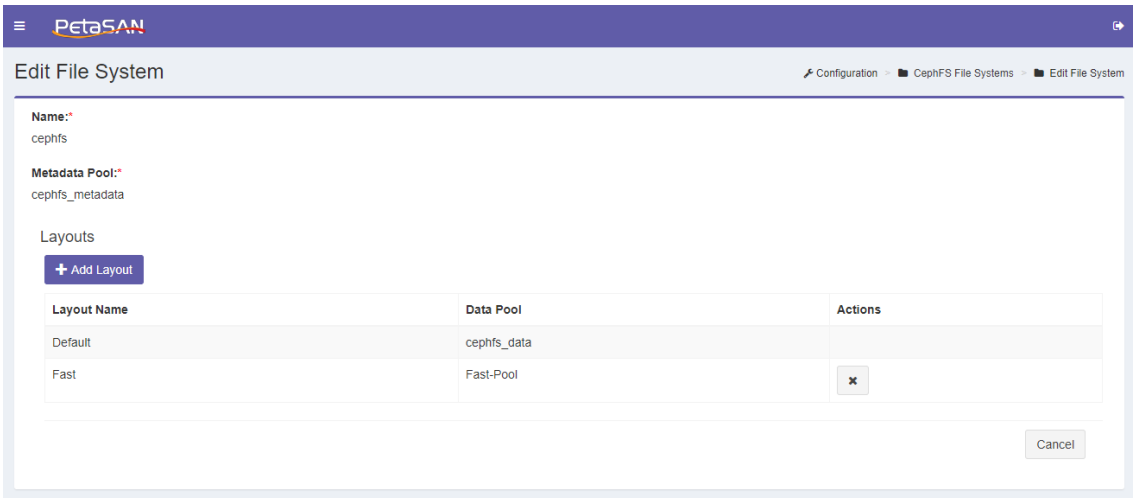

➢ Add layout

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.

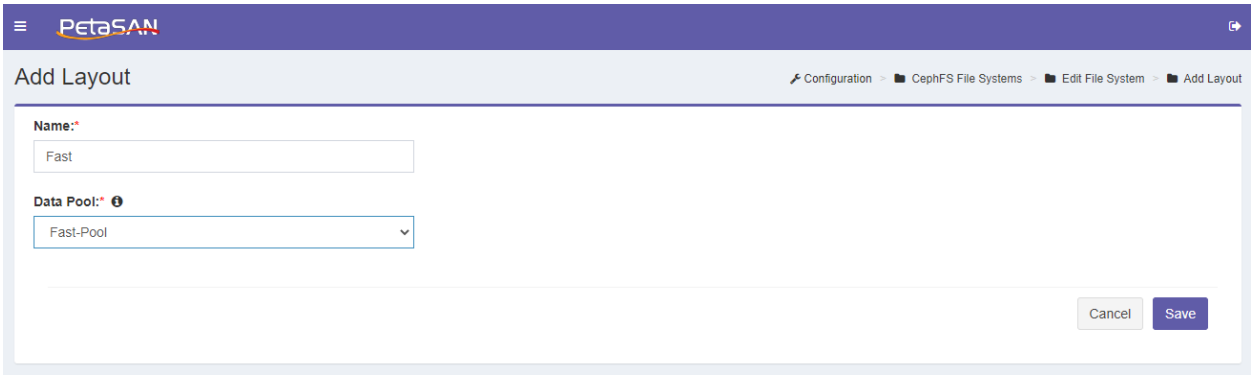

4.8. S3 Configuration

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

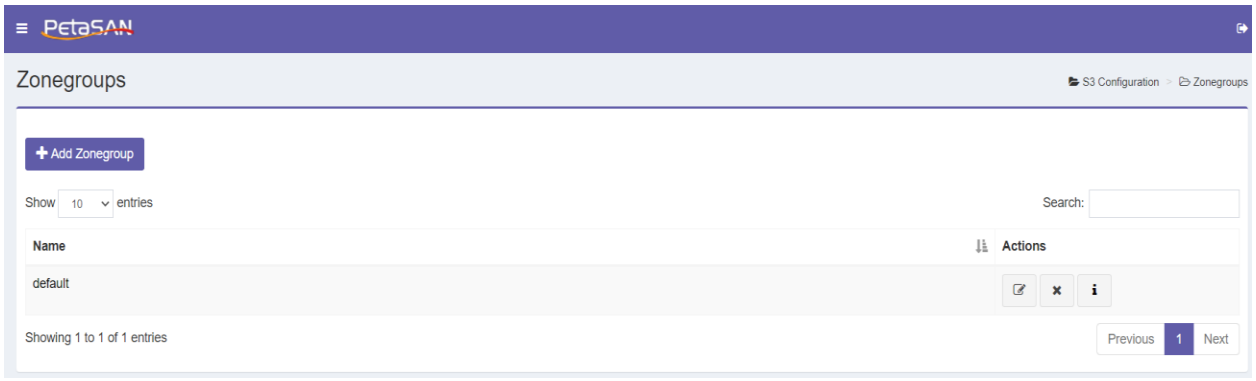

In case the site is not master site, system will view the Pull button as follows

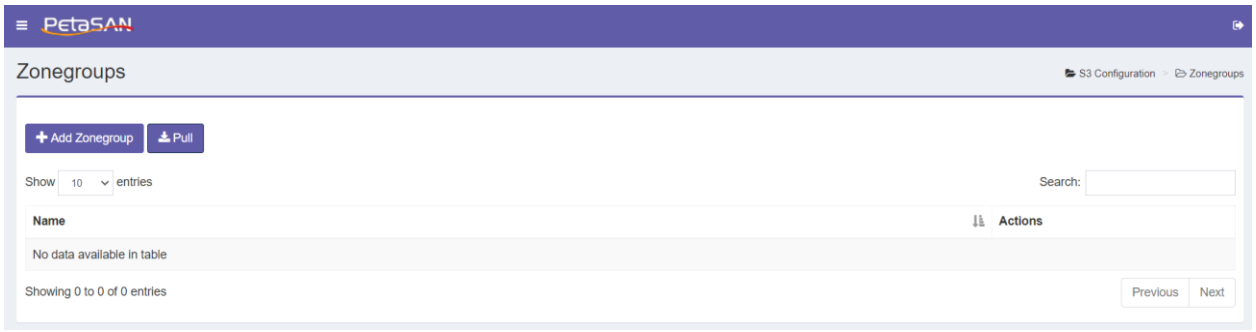

System views the following columns for the zonegroup.

- **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:

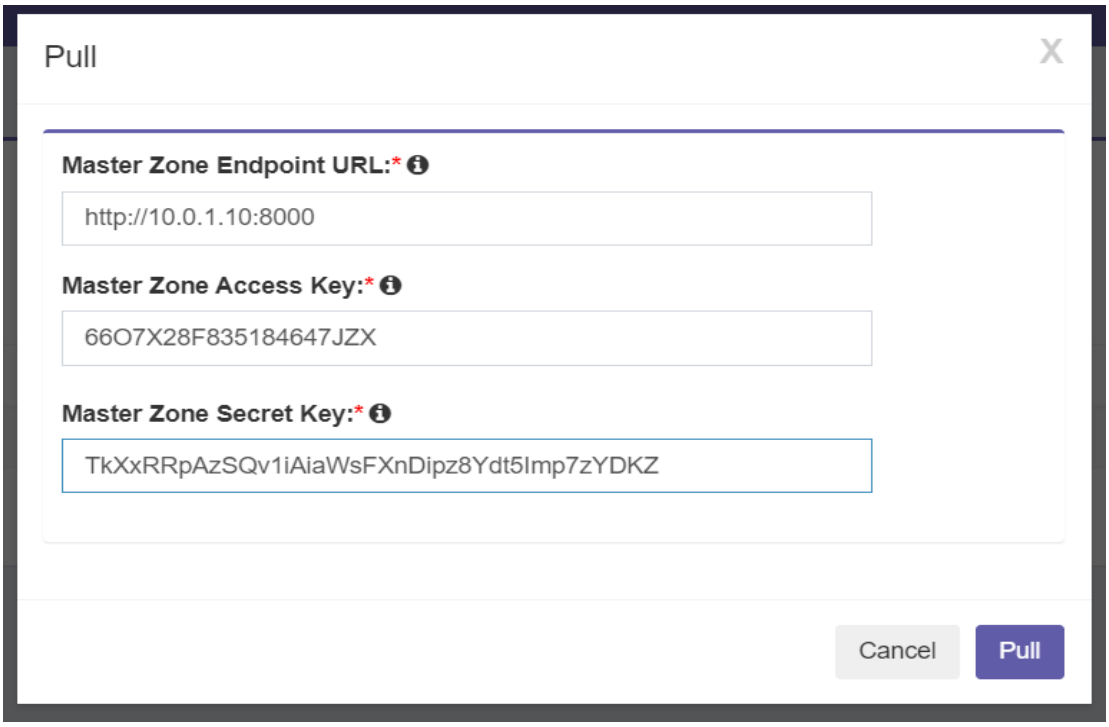

- **Master Zone URL**
  You need to specify the master zone's URL
- **Master Zone Access Key**
  You need to specify the system user key for the master zone.
- **Master Zone Secret Key**
  You need to specify the access key for the master zone.

- **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

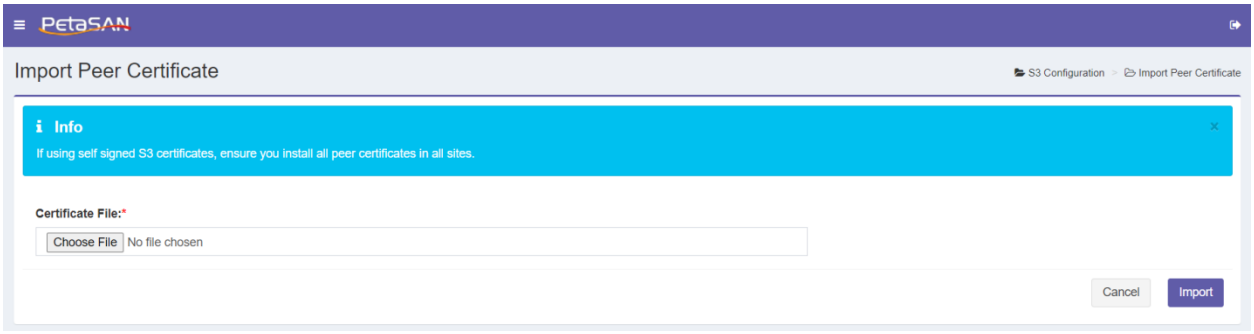

- **Actions**
  - Imports
    Imports the peer certificate.
  - Cancel
    Cancels the import.

➢ 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 STANDERD 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 that the default-placement
  - Add Storage Class

Adds new storage class under one of the existing placement targets

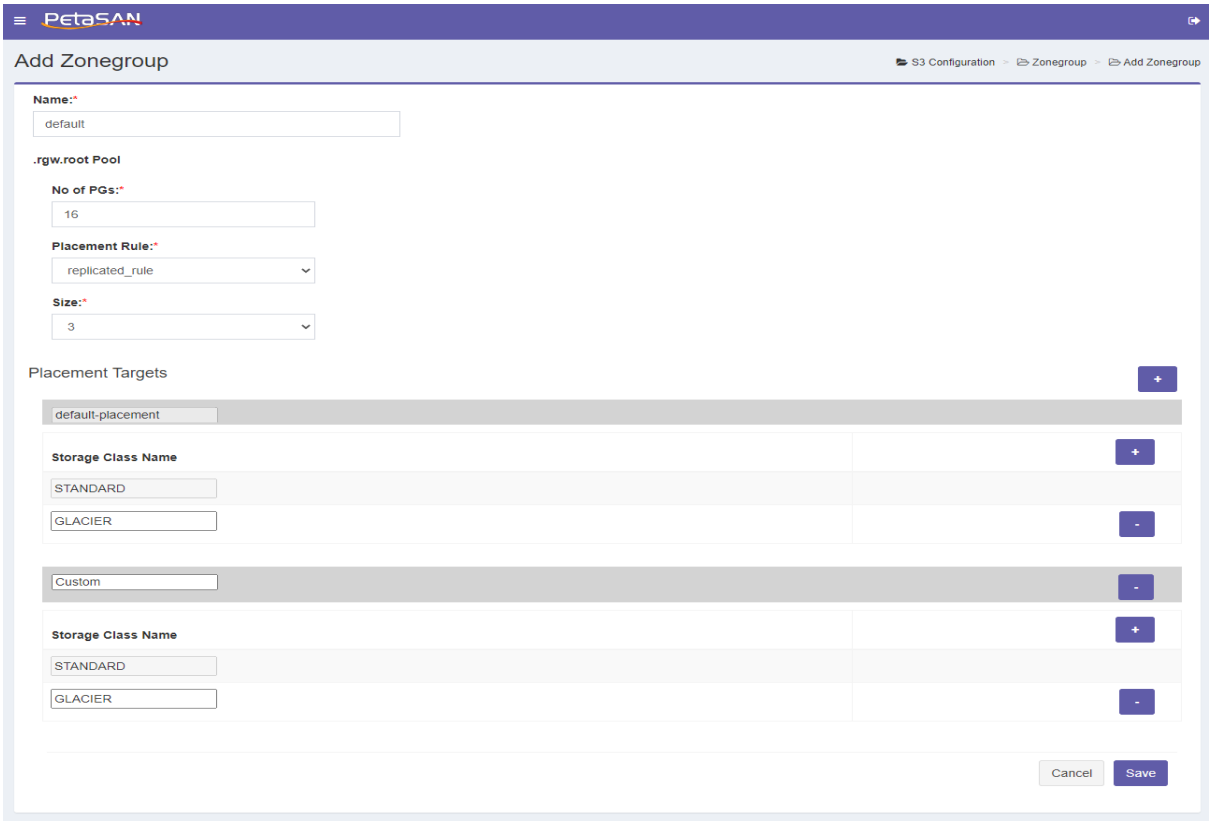

- Note:
  - You must enter the S3 Settings and make sure at least one node has the S3 role before adding new zonegroup.
  - You may find .rgw.root pool created before adding a zonegroup,this is because ceph creates it if any rgw command has been run, when you add a new zonegroup system will delete the 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

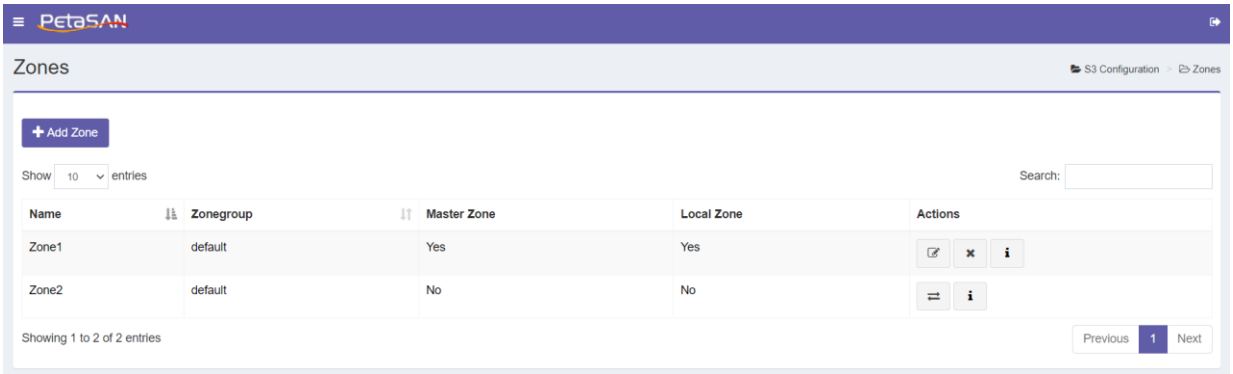

And in case of non master zone, system will view the zones list as follows:

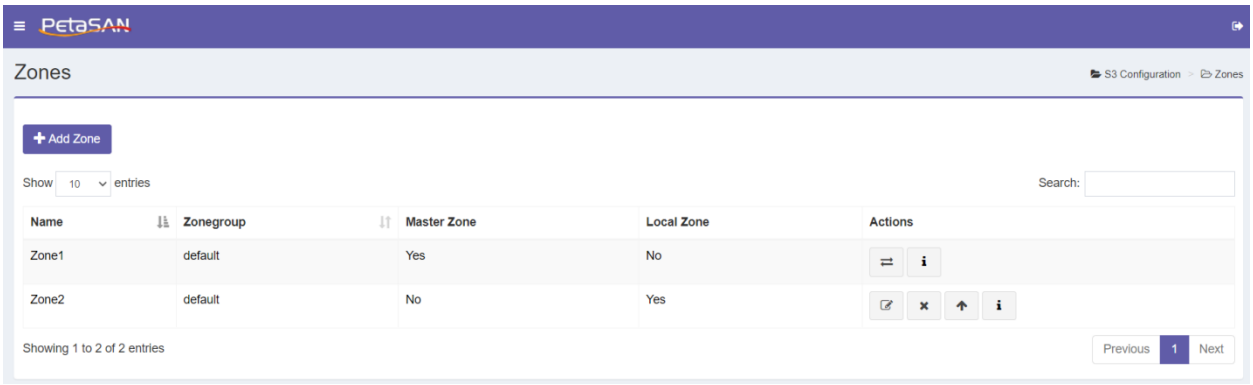

System views the following columns for the zones

- **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, my 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.

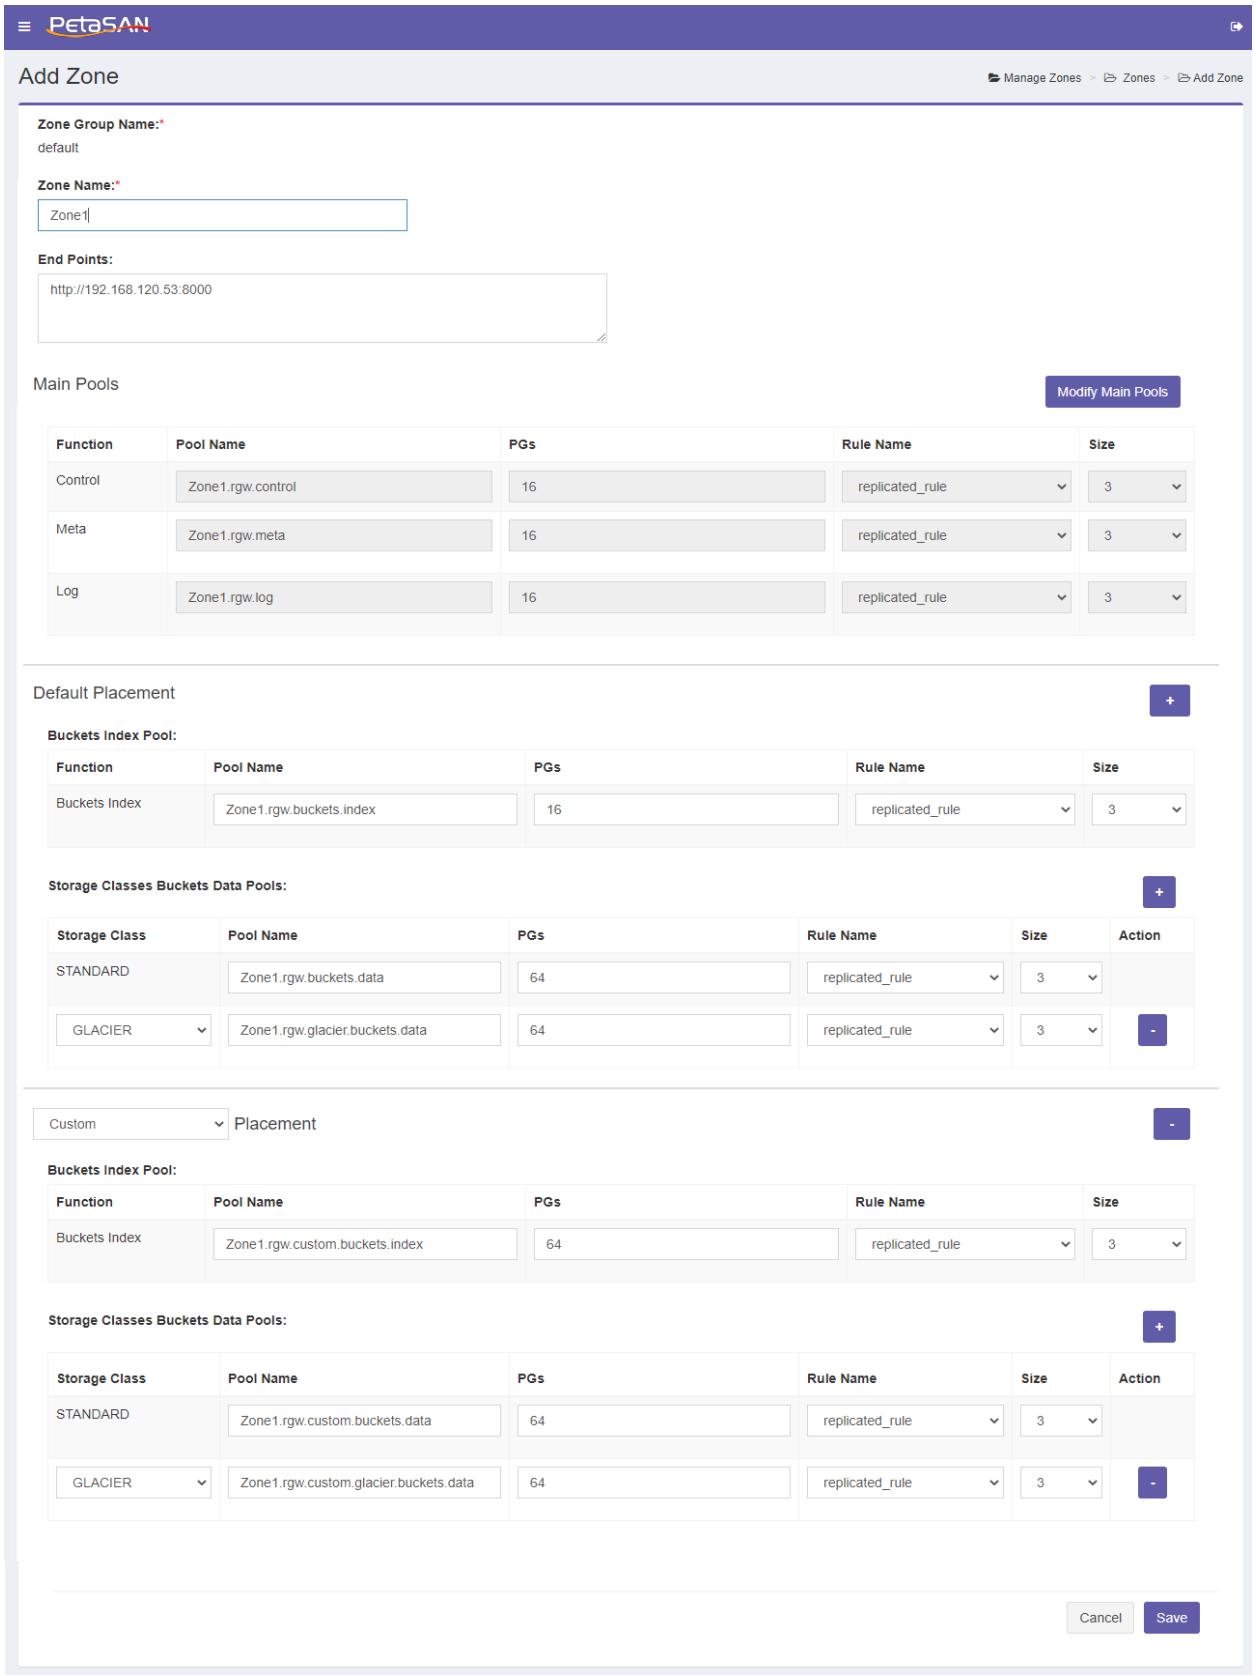
PetaSAN
Add Zone
Manage Zones > Zones > Add Zone
Zone Group Name:*
default
Zone Name:*
Zone1
End Points:
http://192.168.120.53:8000
Main Pools
Modify Main Pools
Function
Pool Name
PGs
Rule Name
Size
Control
Zone1.rgw.control
16
replicated_rule
3
Meta
Zone1.rgw.meta
16
replicated_rule
3
Log
Zone1.rgw.log
16
replicated_rule
3
Default Placement
Buckets Index Pool:
Function
Pool Name
PGs
Rule Name
Size
Buckets Index
Zone1.rgw.buckets.index
16
replicated_rule
3
Storage Classes Buckets Data Pools:
Storage Class
Pool Name
PGs
Rule Name
Size
Action
STANDARD
Zone1.rgw.buckets.data
64
replicated_rule
3
GLACIER
Zone1.rgw.glacier.buckets.data
64
replicated_rule
3
Custom
Placement
Buckets Index Pool:
Function
Pool Name
PGs
Rule Name
Size
Buckets Index
Zone1.rgw.custom.buckets.index
64
replicated_rule
3
Storage Classes Buckets Data Pools:
Storage Class
Pool Name
PGs
Rule Name
Size
Action
STANDARD
Zone1.rgw.custom.buckets.data
64
replicated_rule
3
GLACIER
Zone1.rgw.custom.glacier.buckets.data
64
replicated_rule
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

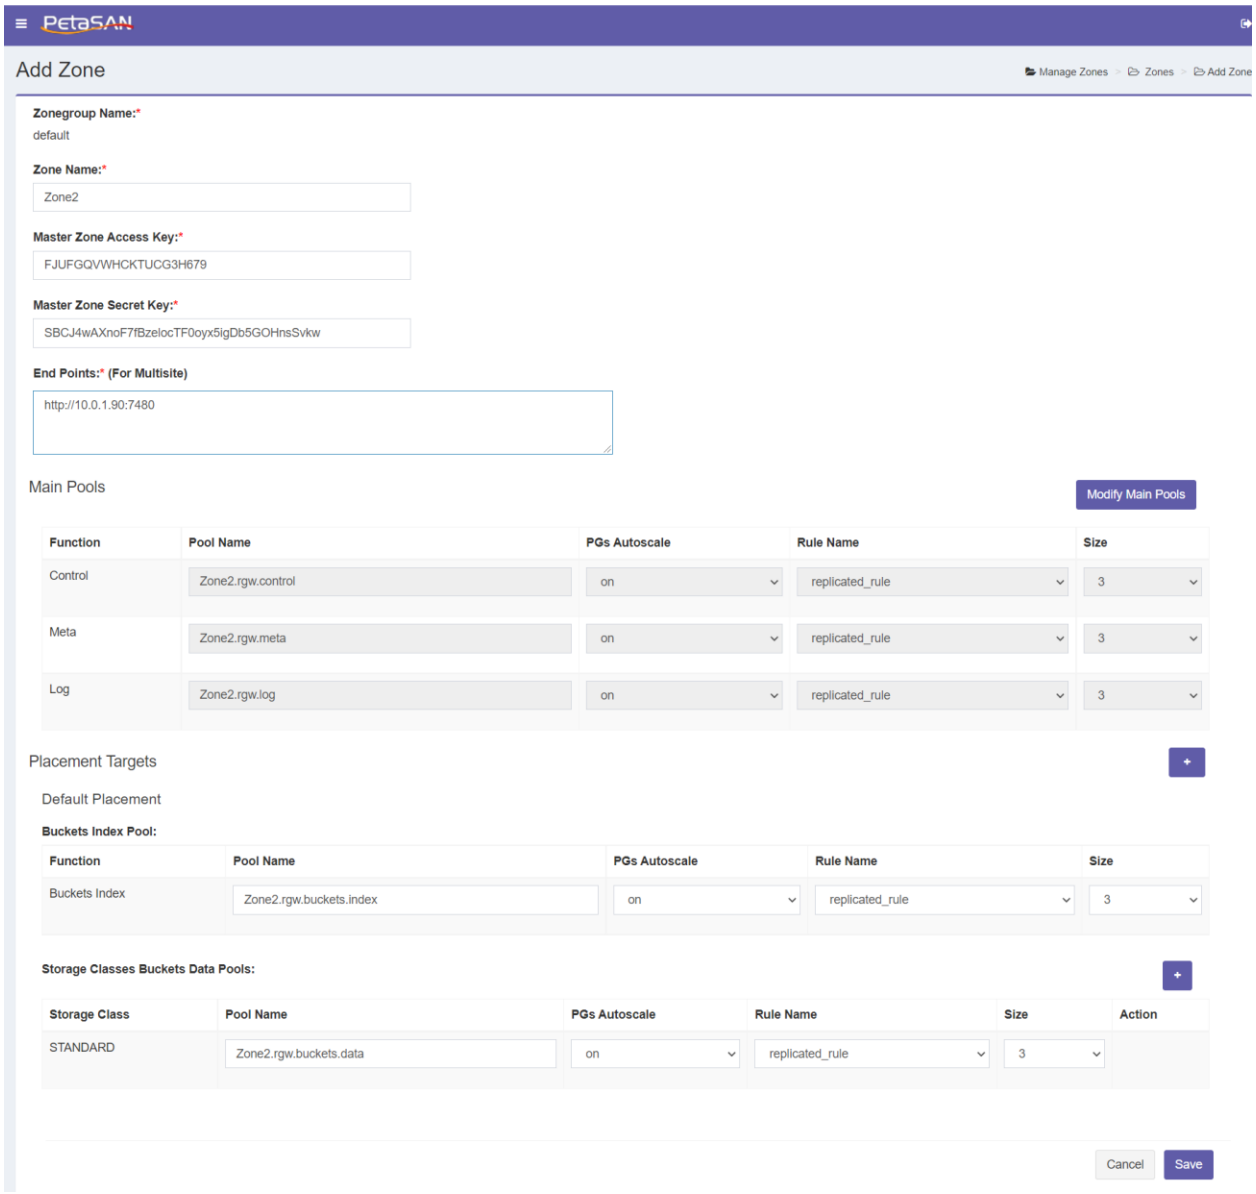

4.9. General Settings

➢ Update General Settings

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

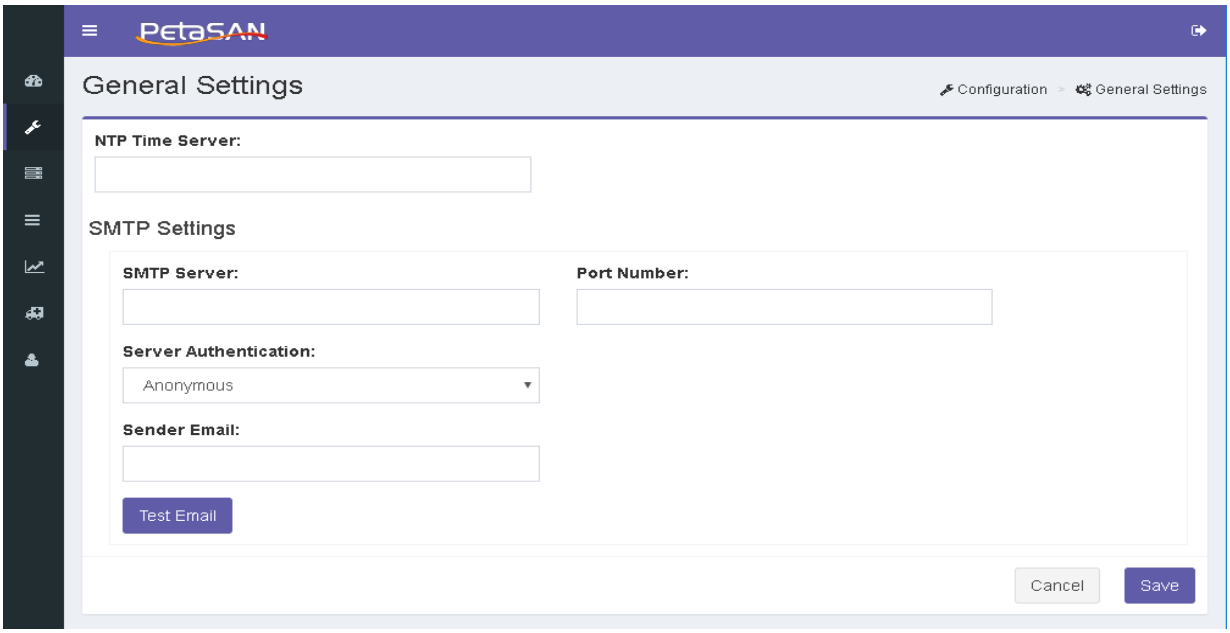

The form contains the following fields:

- **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.

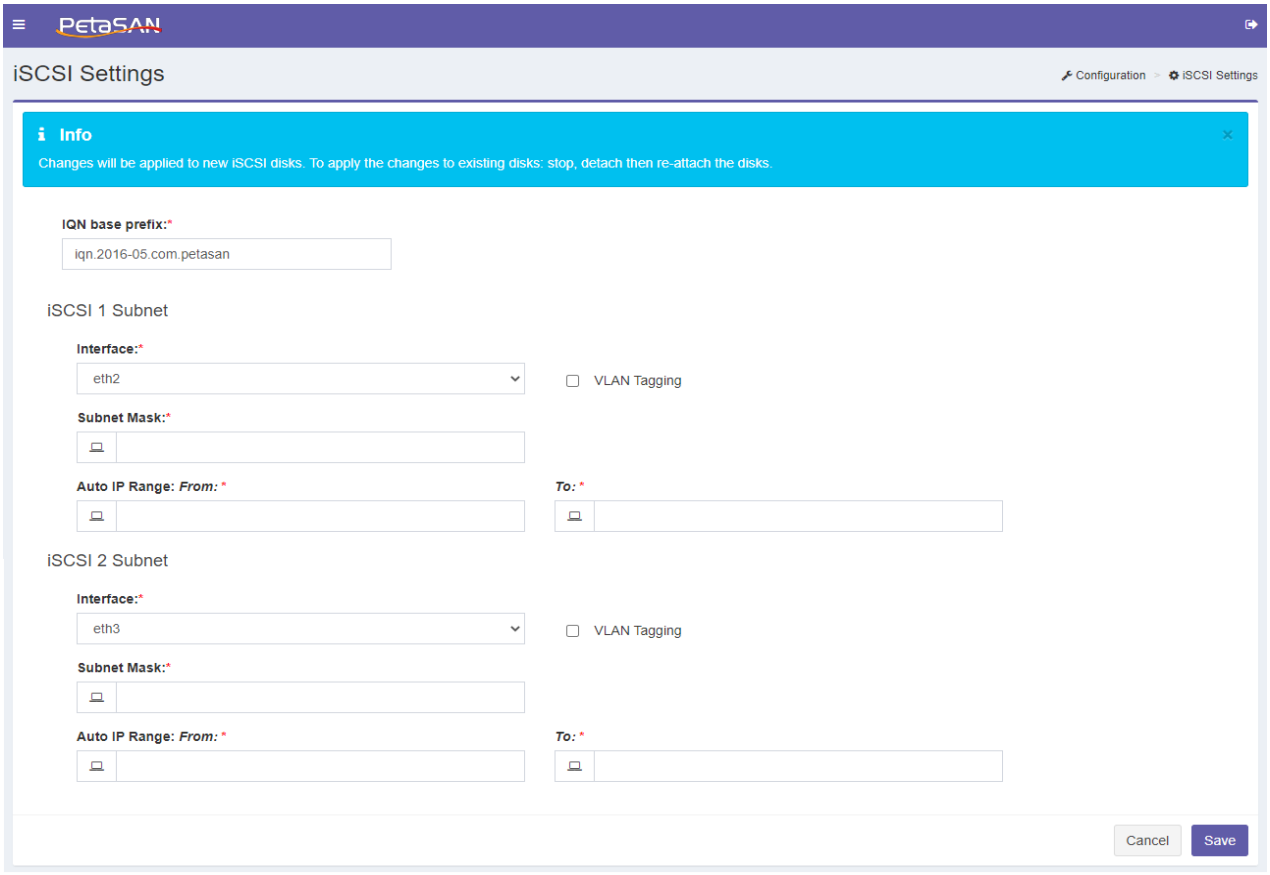

The form contains the following fields:

- **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.

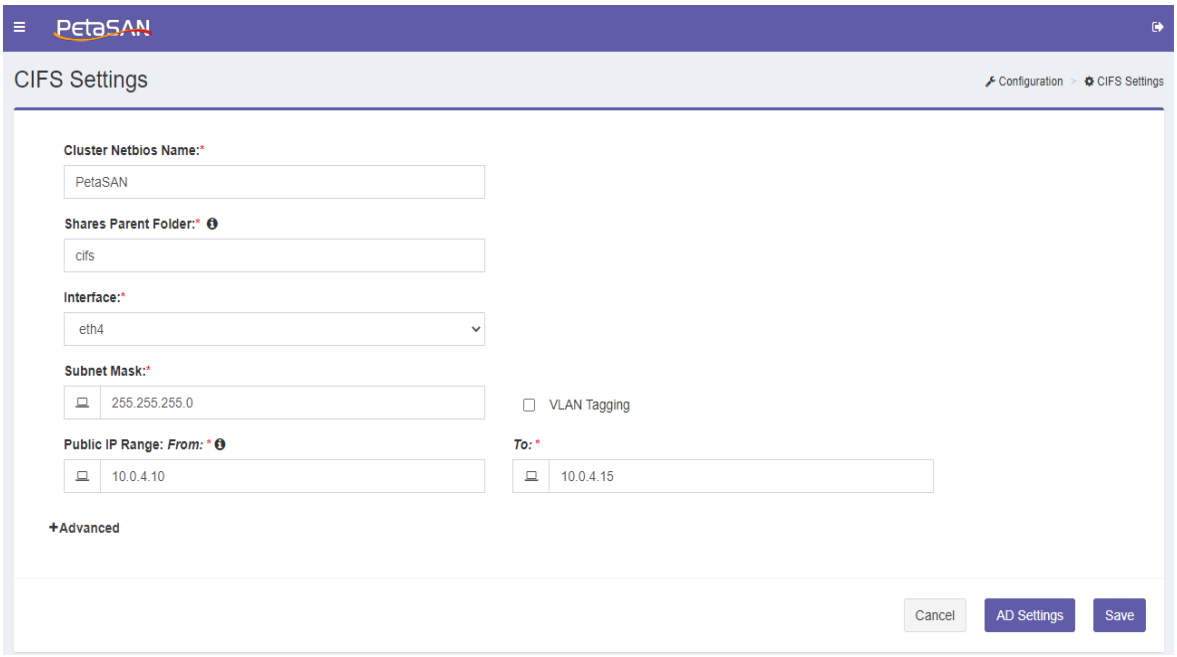

The CIFS settings form allows specifying the following fields

- **Cluster Netbios Name**
  The Nebios 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 usesthe 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

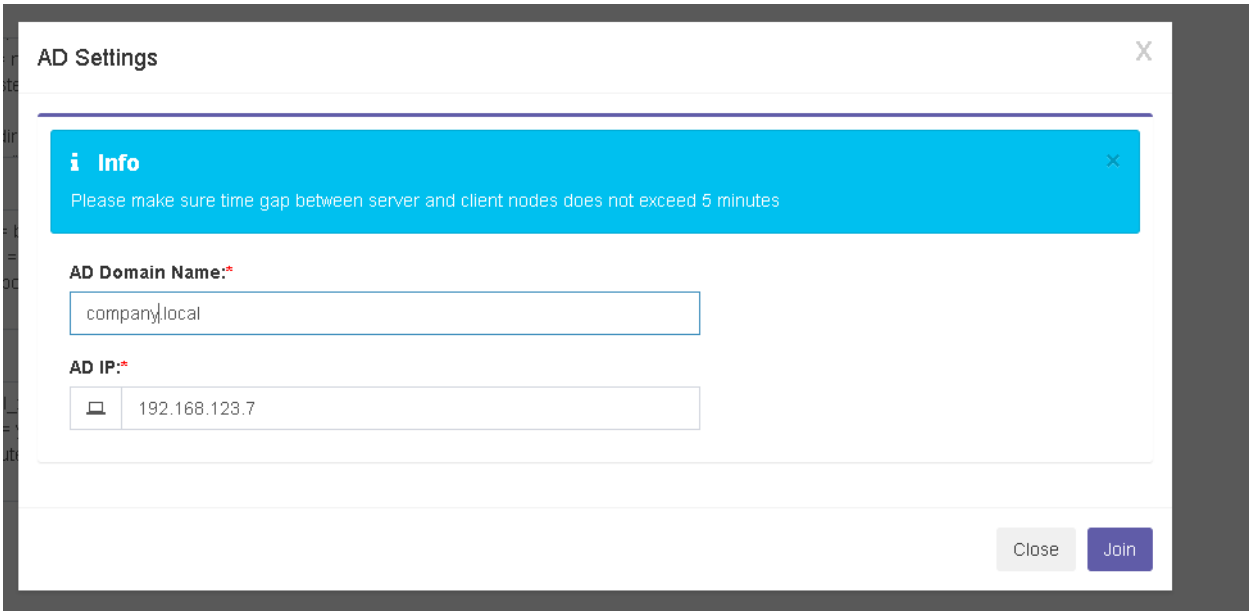

You will need to enter the Active Directory settings by entering the following fields:

- **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.

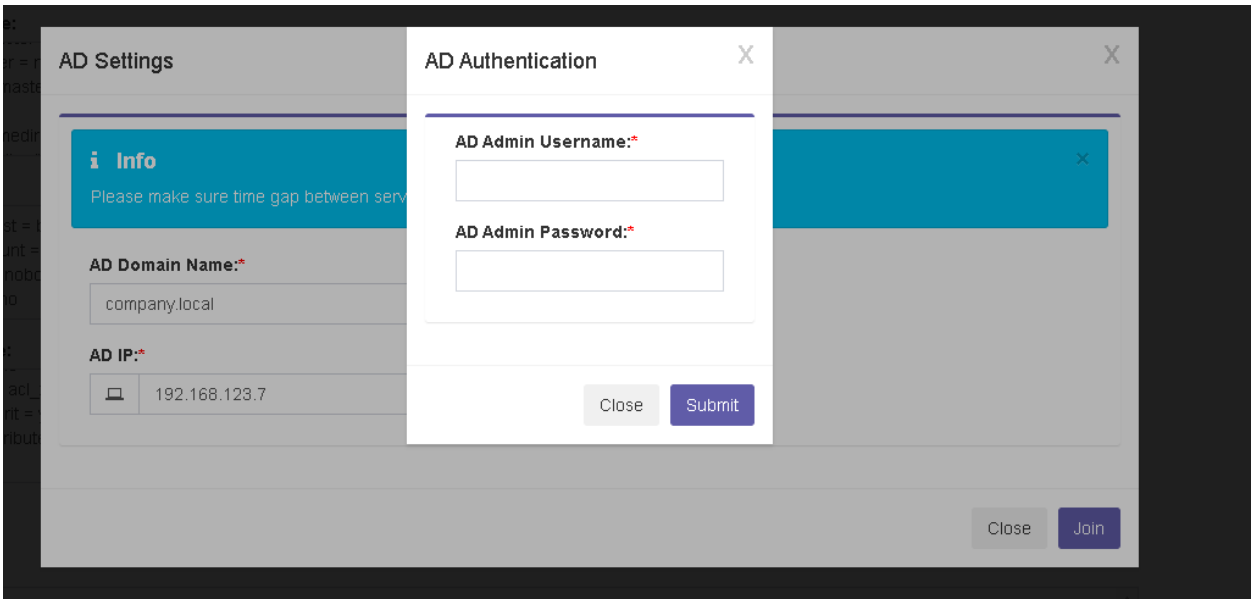
AD Settings
AD Authentication
Info
Please make sure time gap between serv
AD Admin Username:*
AD Admin Password:*
AD Domain Name:*
company.local
AD IP:*
192.168.123.7
Close
Submit
Close
Join

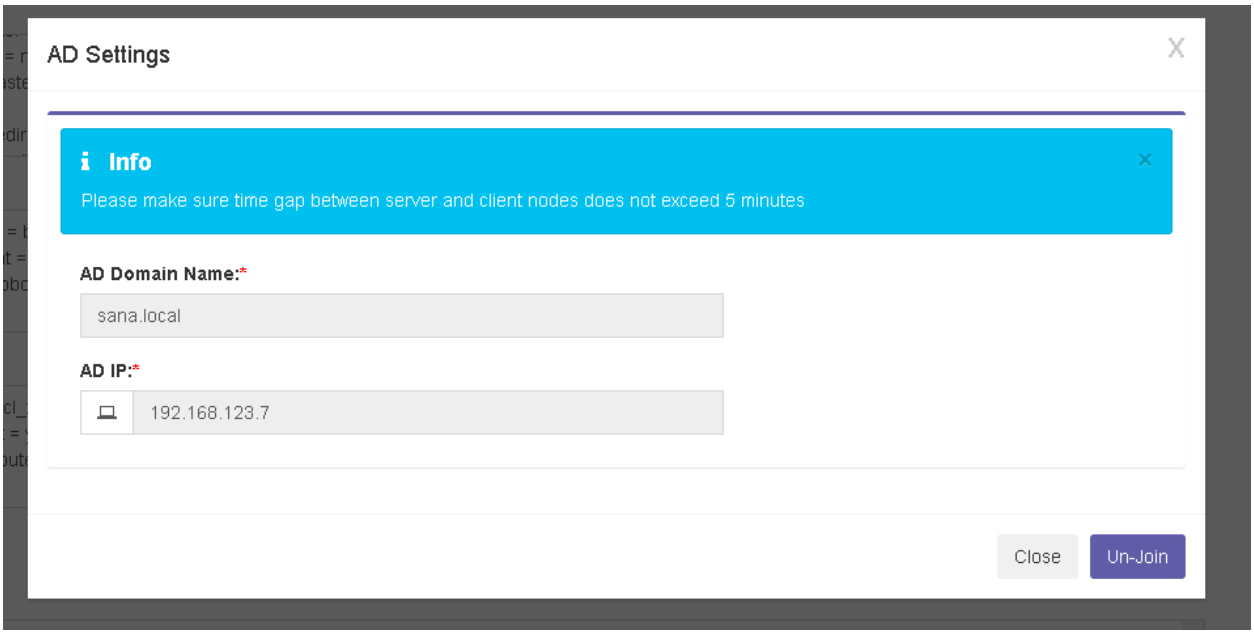
AD Settings
Info
Please make sure time gap between server and client nodes does not exceed 5 minutes
AD Domain Name:*
sana.local
AD IP:*
192.168.123.7
Close
Un-Join

4.12. NFS Settings

➢ Update NFS Settings

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

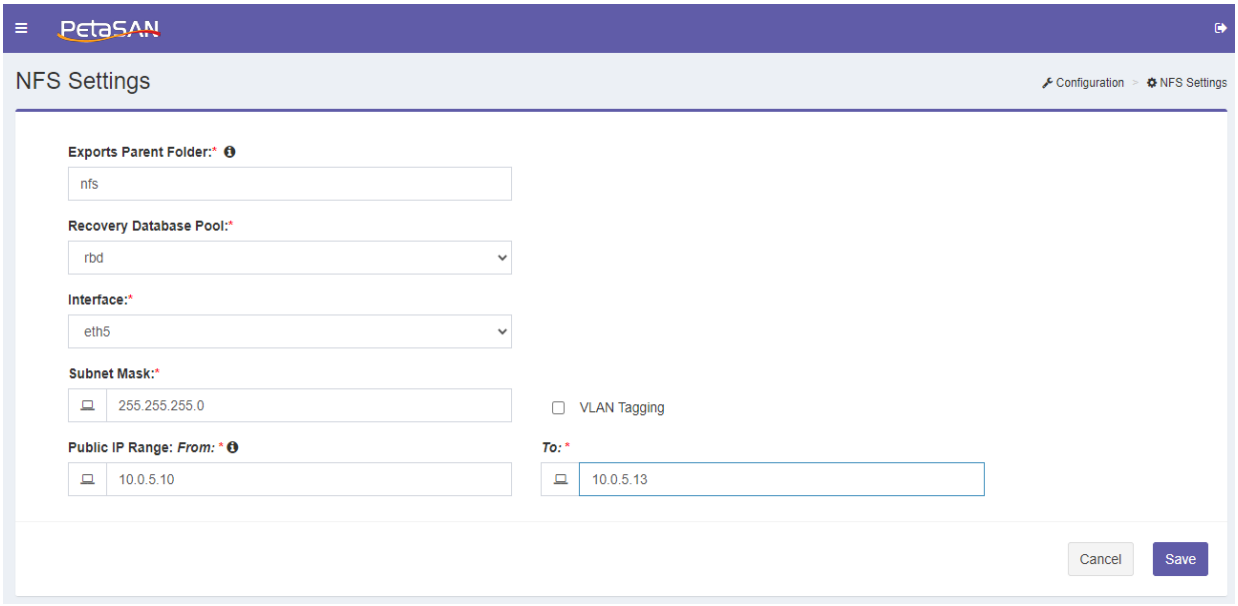

The NFS settings form allows specifying the following fields

- **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.

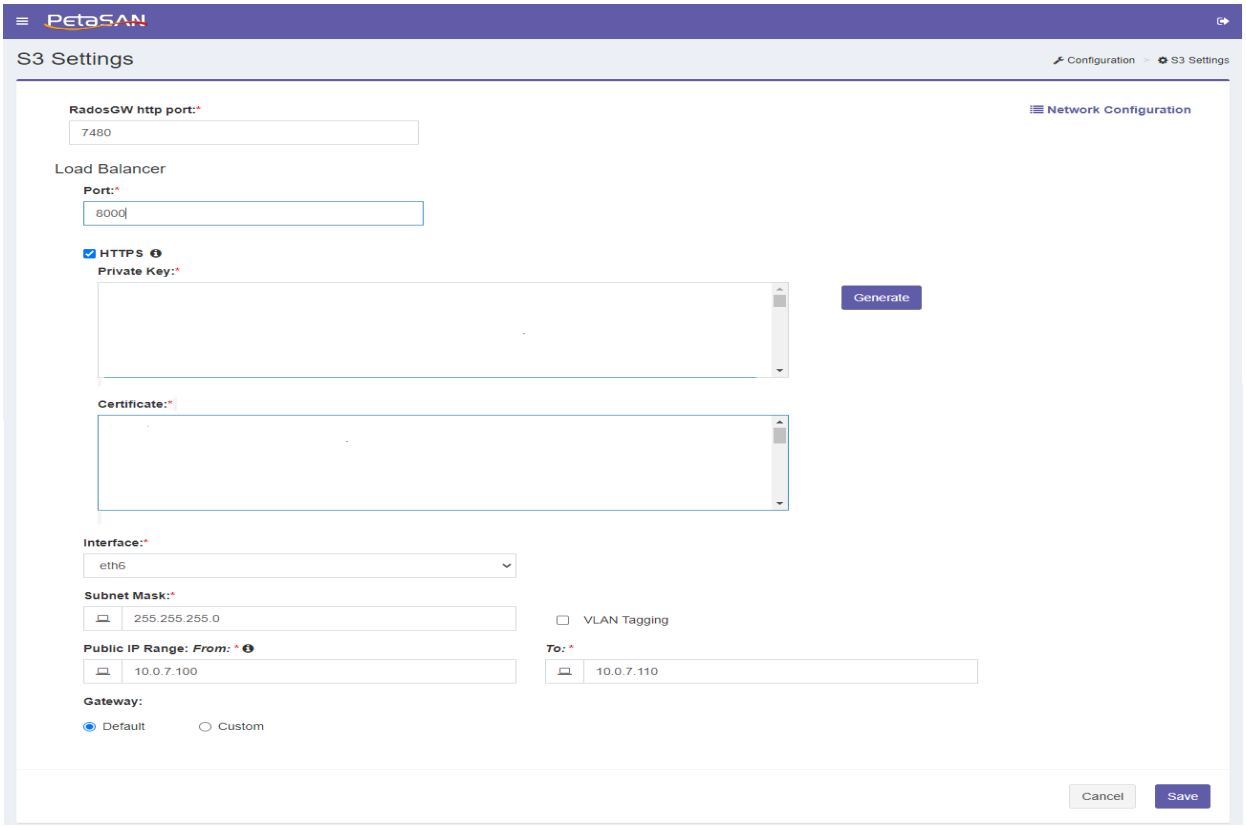

The S3 settings form allows specifying the following fields

- **RadosGW http port**
  The backend radosgw port, default is 7480.t.

- **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

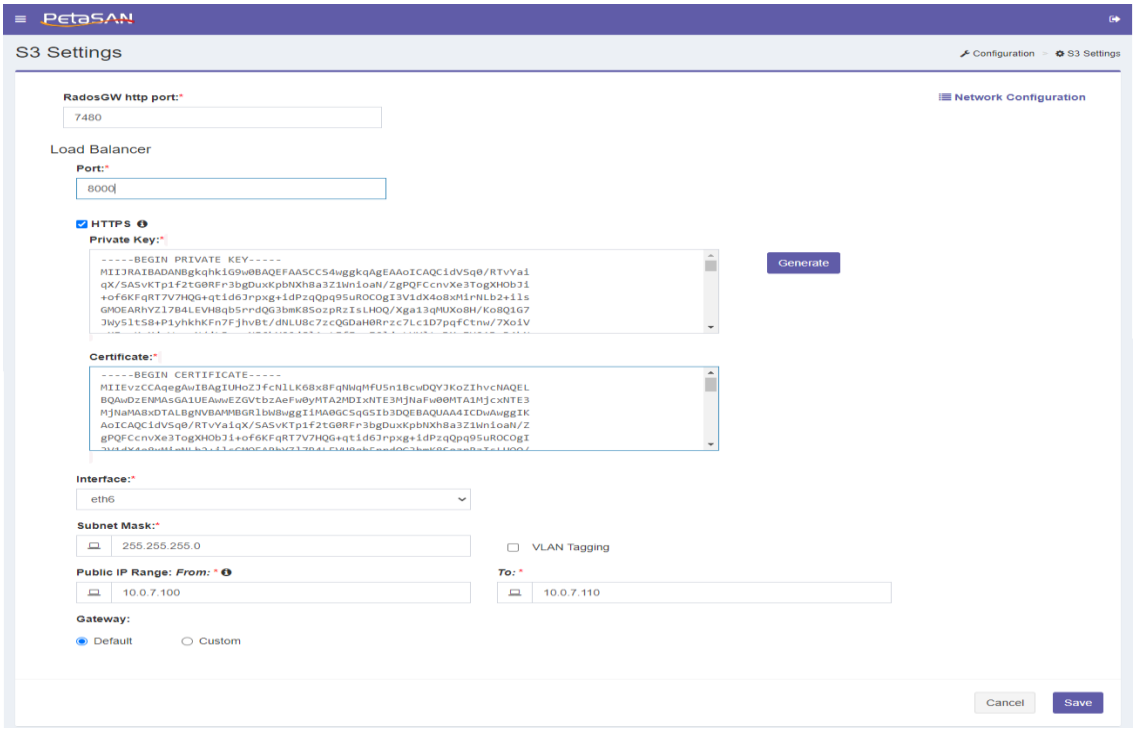

- **Interface**
  Interface for the S3 Load Balancer subnet ,it can be changed any time as long as all nodes hosting the S3 role have the selected Ethernet card.
- **Subnet Mask**
  The subnet mask of the S3 public subnet used by clients.
- **VLAN Tagging**
  Enable VLAN tagging for the S3 subnet, if 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**

- **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.

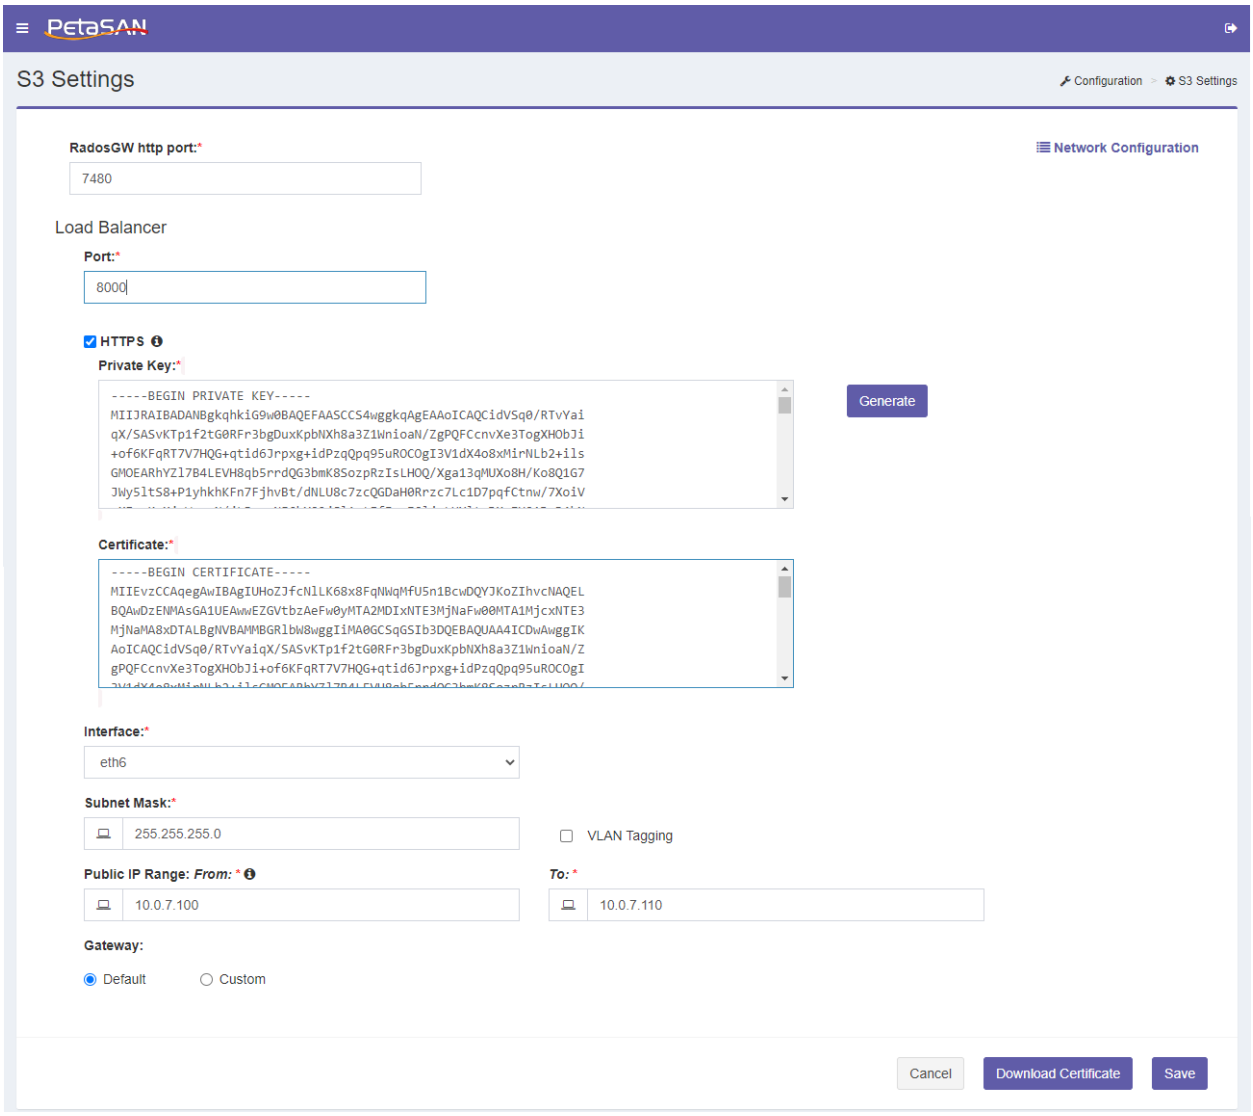

5. Manage Nodes

➢ Node List

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

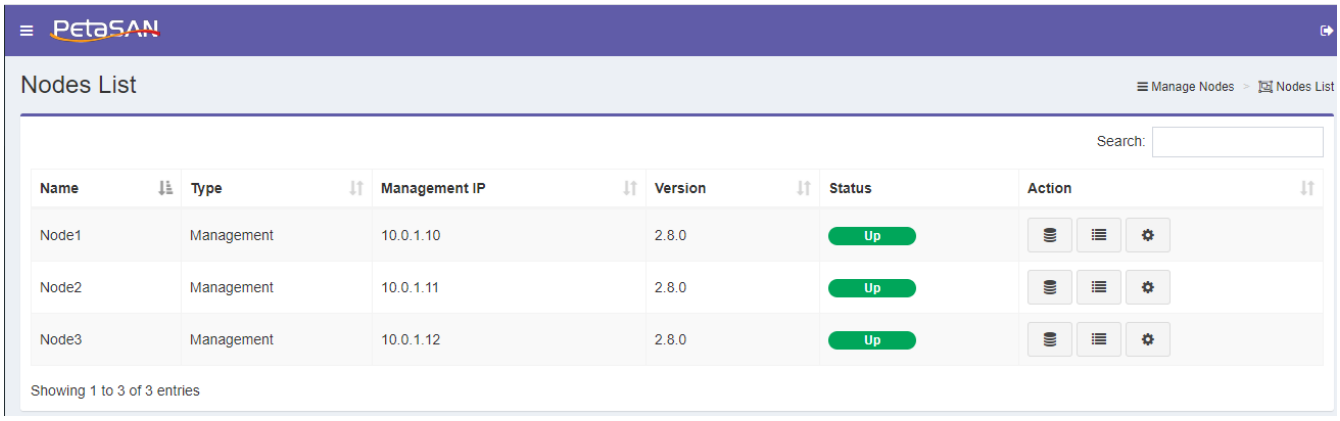

The node list views the following information about each node:

- **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.

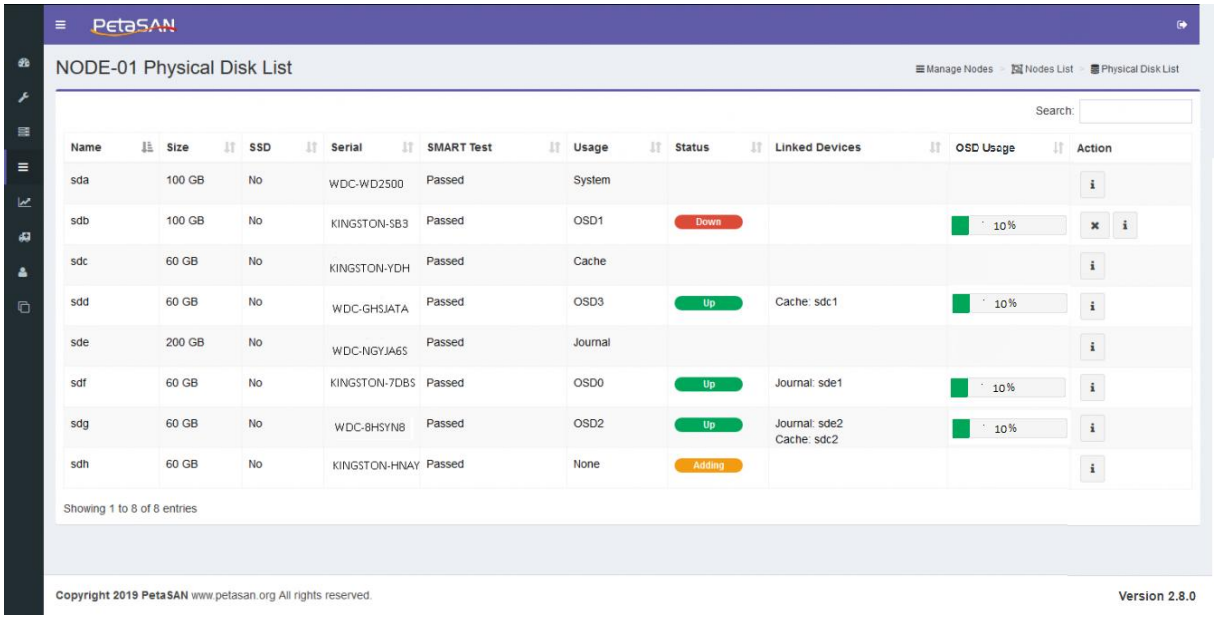

The node local disks list views the following information about each disk:

- **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 .

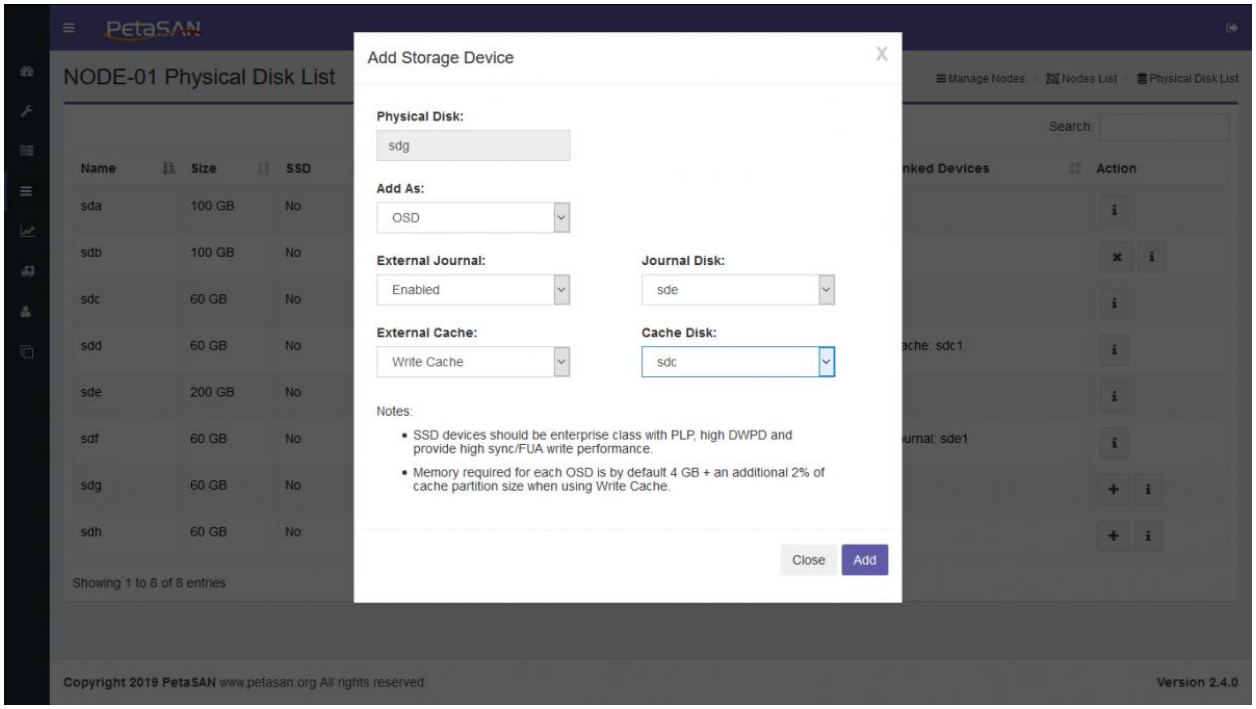

The system enables you to add the disk as

- 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.

➢ View Log

The system views the system log of the node.

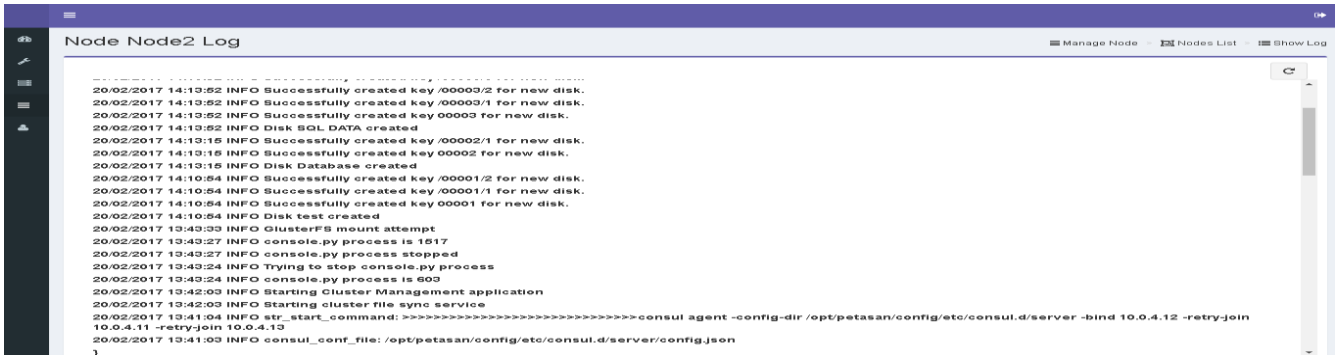

You can use the refresh button to view the latest log entries.

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

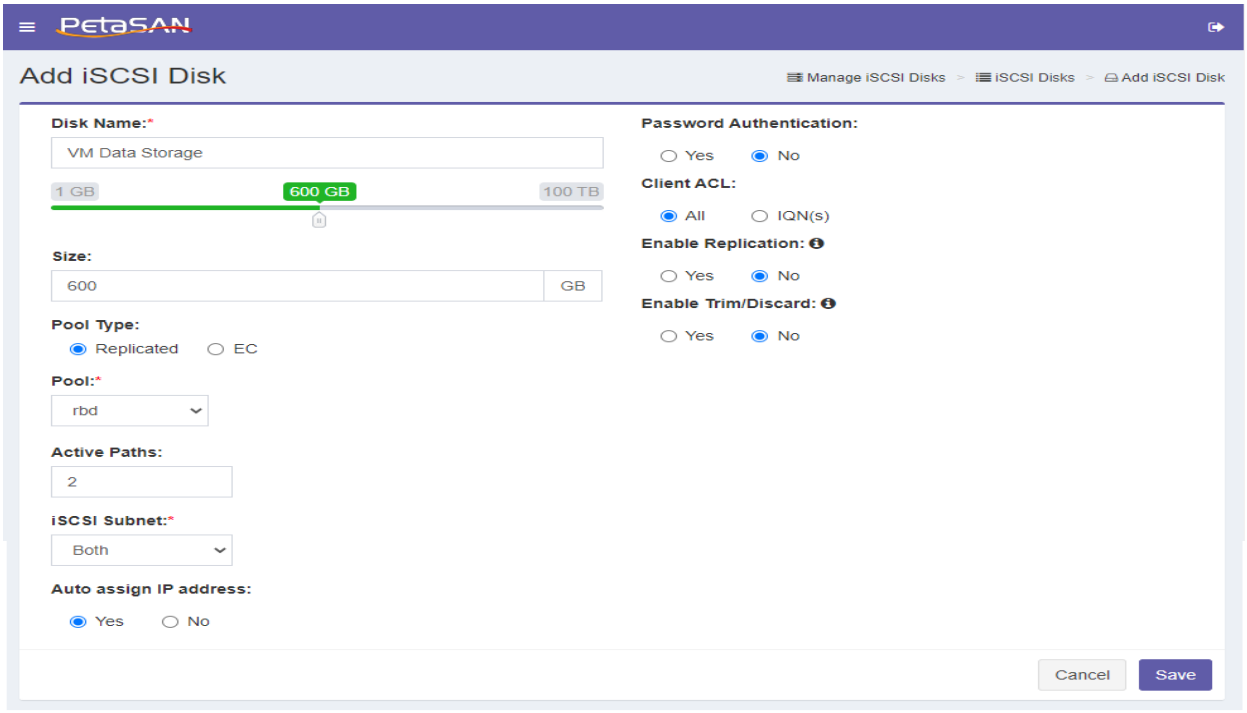

The Add iSCSI disk from contains the following fields:

- **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

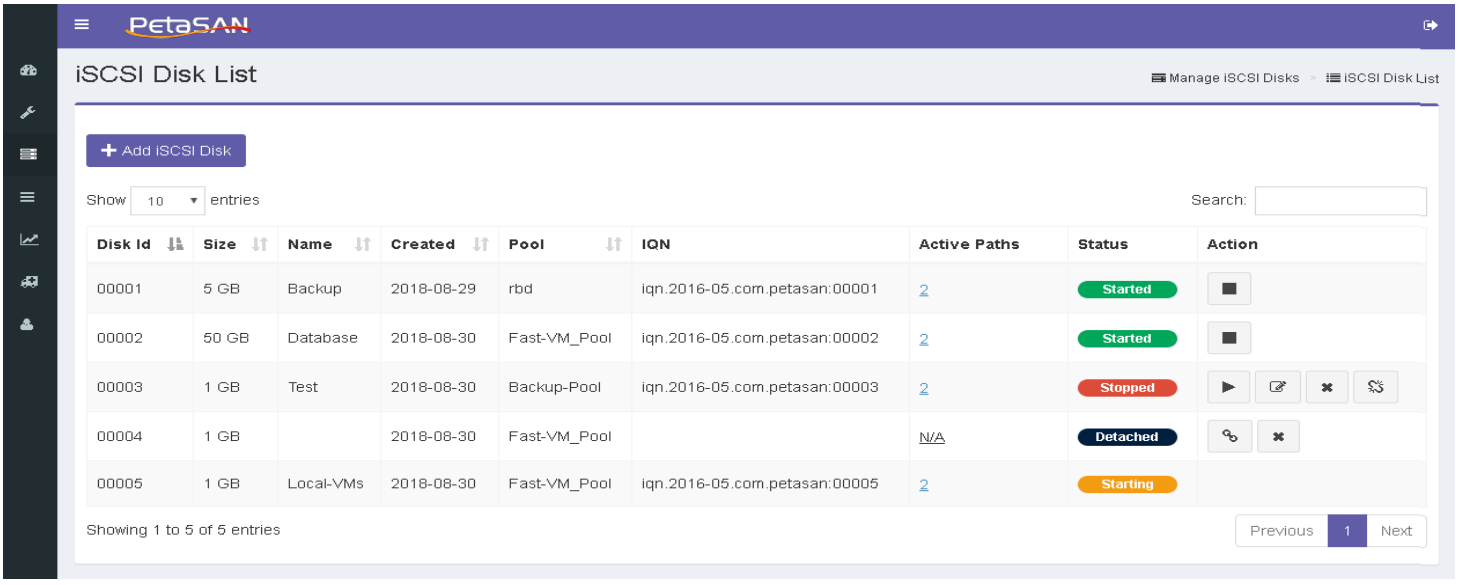

The iSCSI Disk List displays the following information for each disk:

- **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.

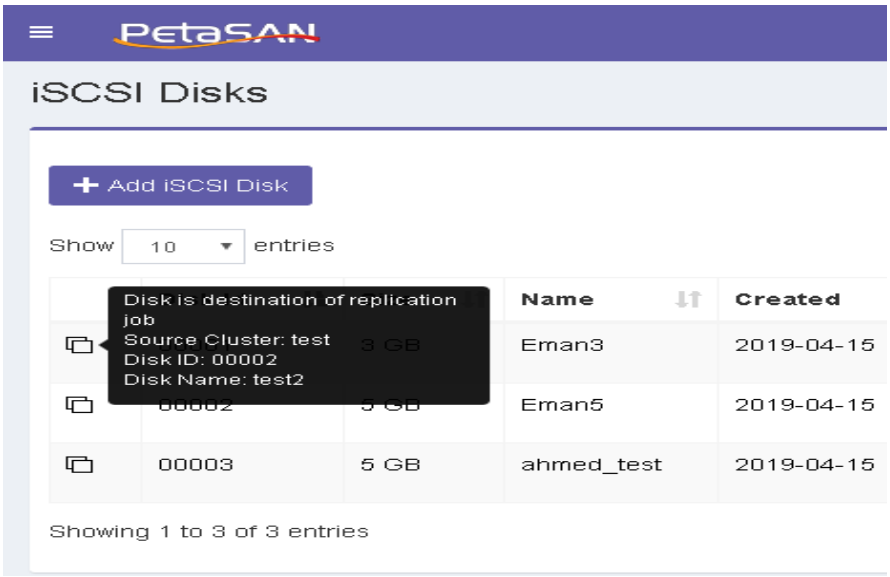

- **Disk Id**
  The unique disk id assigned to the disk by the system
- **Size**
  The disk size
- **Name**
  The disk name.
- **Created**
  The date when the disk was created
- **Pool**
  The storage pool storing the disk image
- **IQN**
  The IQN assigned to the disk
- **Active Paths**
  The number 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

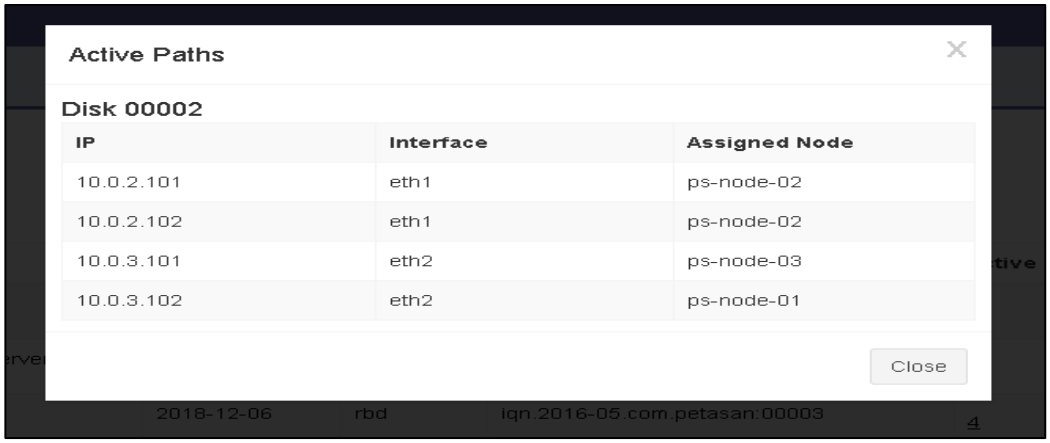

- **Status**
  Display the current status of the disk which can be one of the following:
  - Started
    At least one of the paths has been assigned and clients can connect to the disk.
  - Stopped
    All disk paths have been stopped and clients cannot connect to the disk.
  - Detached
    Disk image 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.

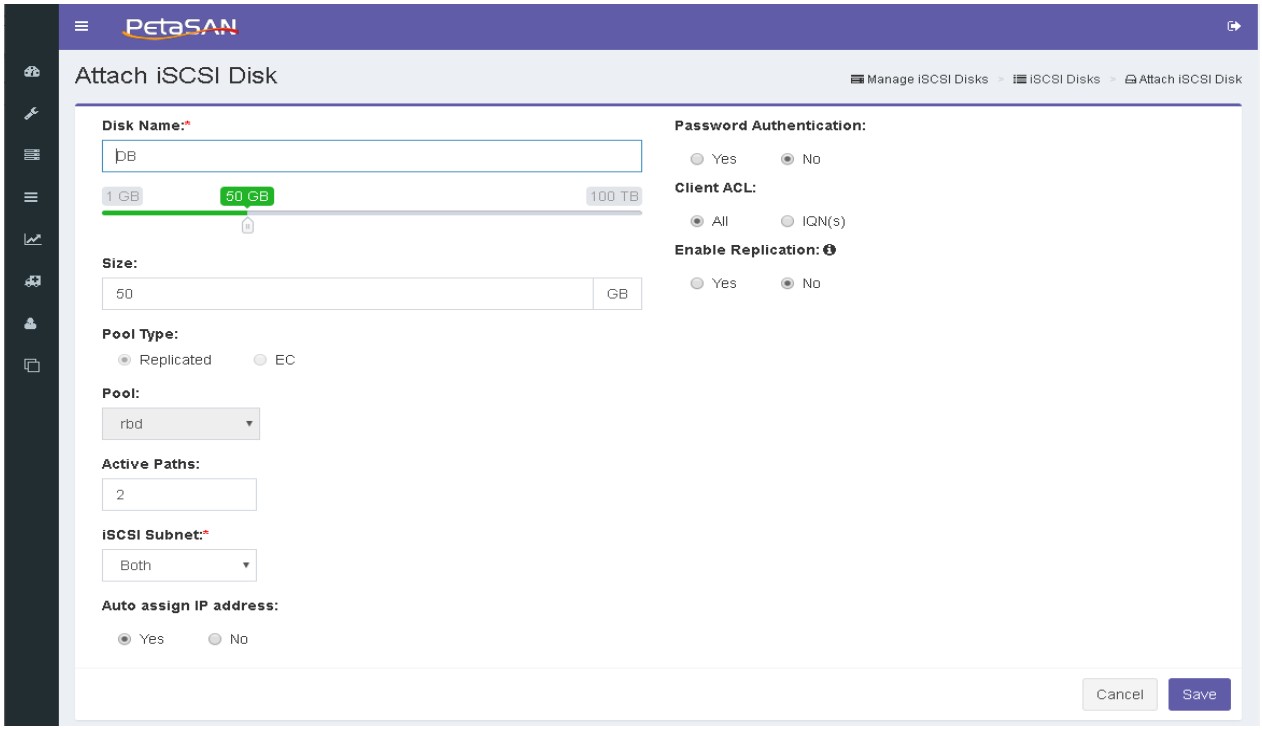

- Edit

  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.*

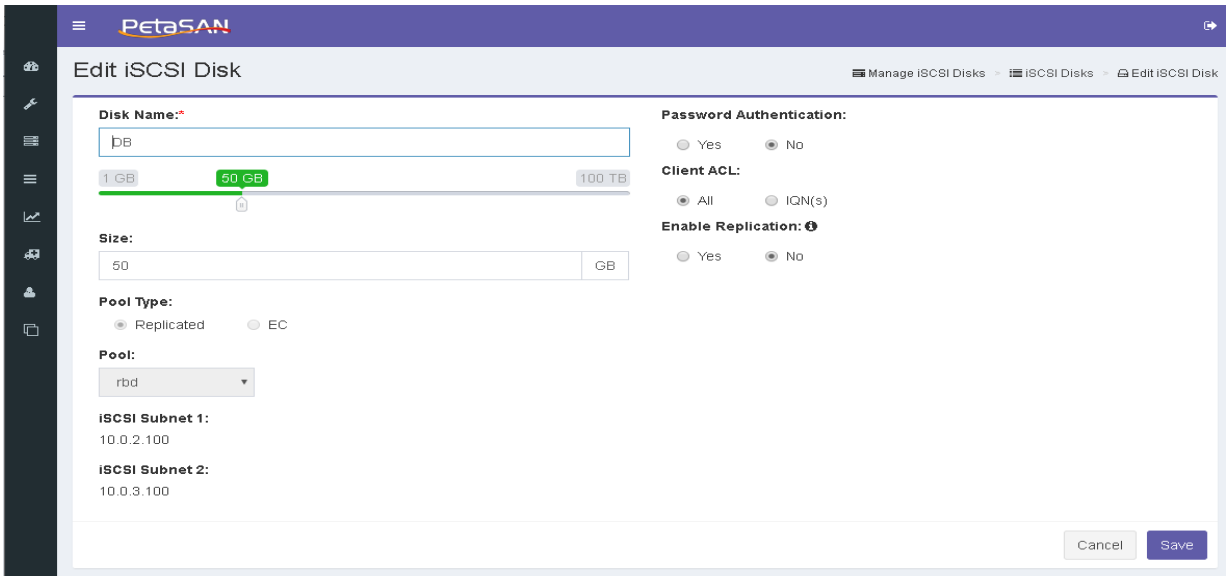

- Delete
  Deletes the disk and all its data from the system and it is available only when the disk is stopped or detached.

➢ *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.

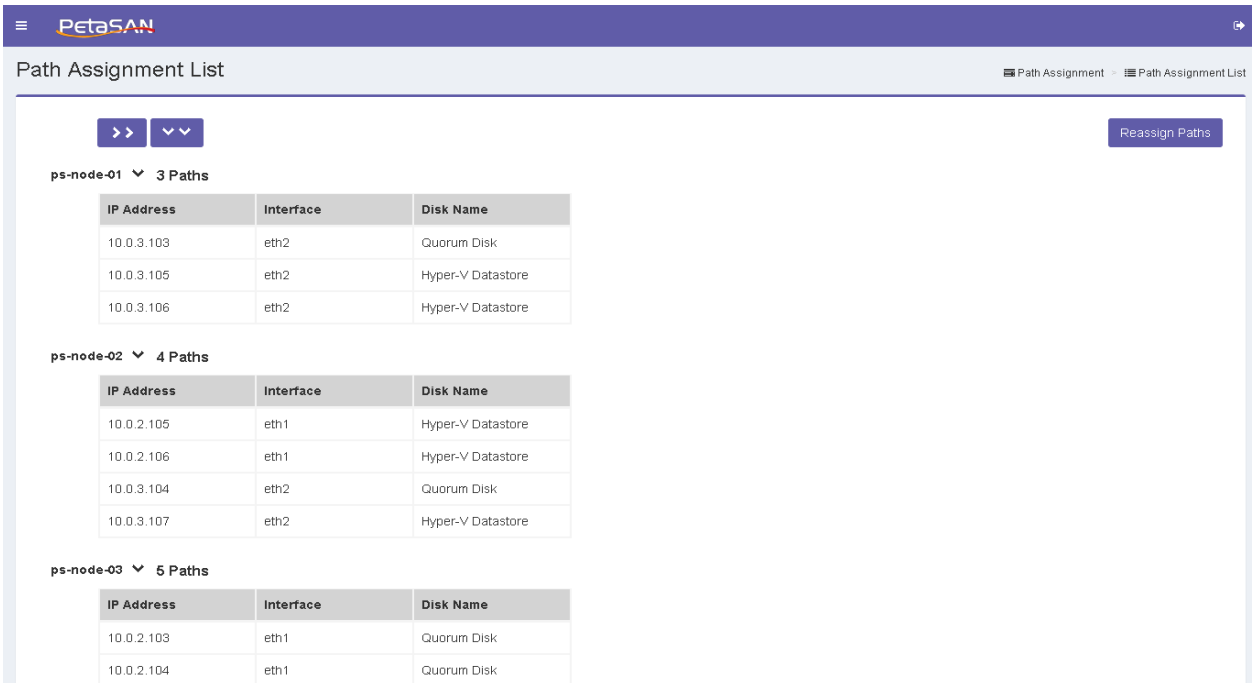

You can expand, collapse nodes or use the expand all or collapse all buttons

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

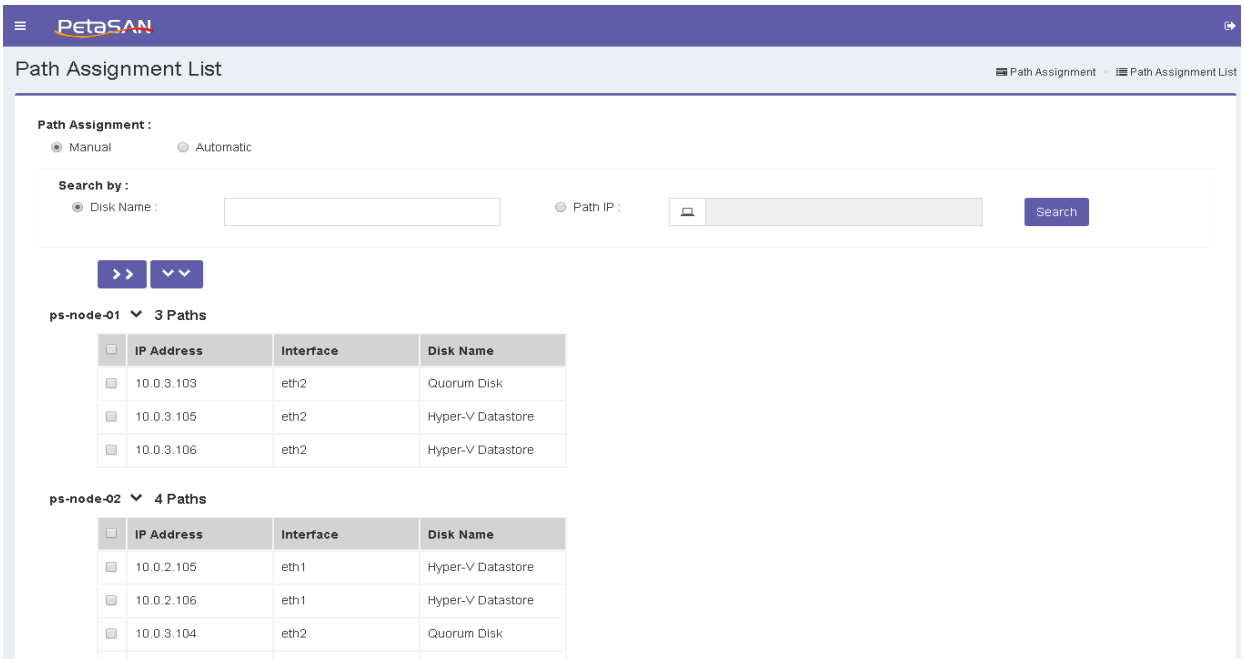

You can search for specific paths using part of the disk name or using specific path IP

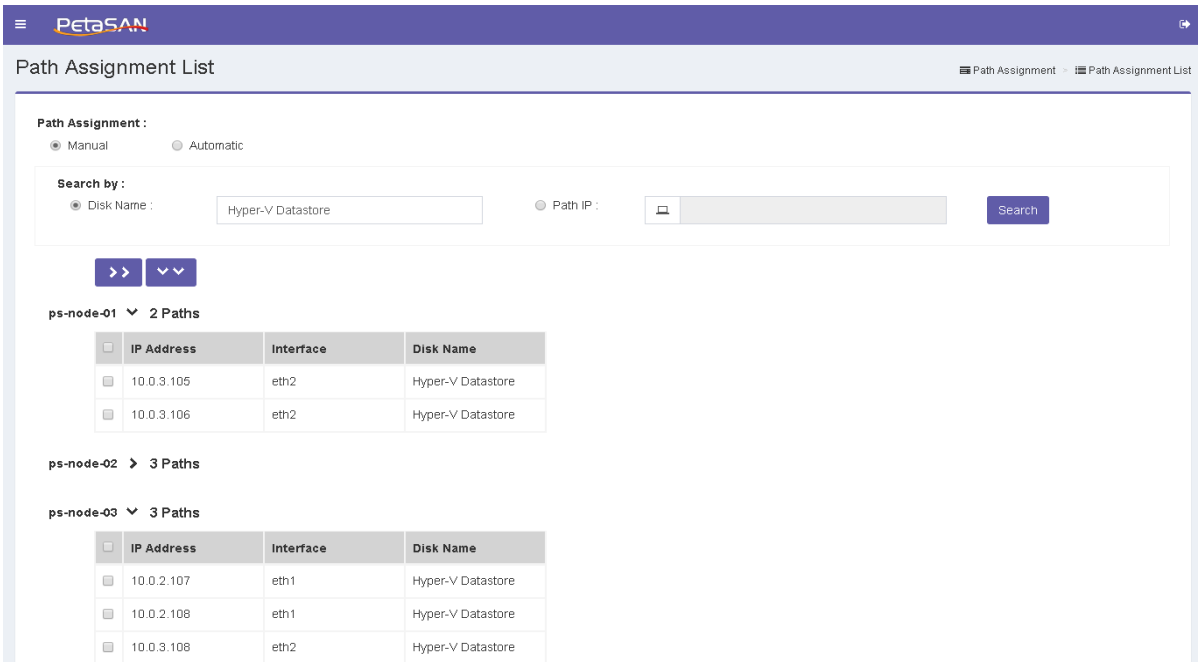

You can select one or more paths and select the destination node. You can select Auto which will assign the selected paths to the node(s) containing the least number of paths.

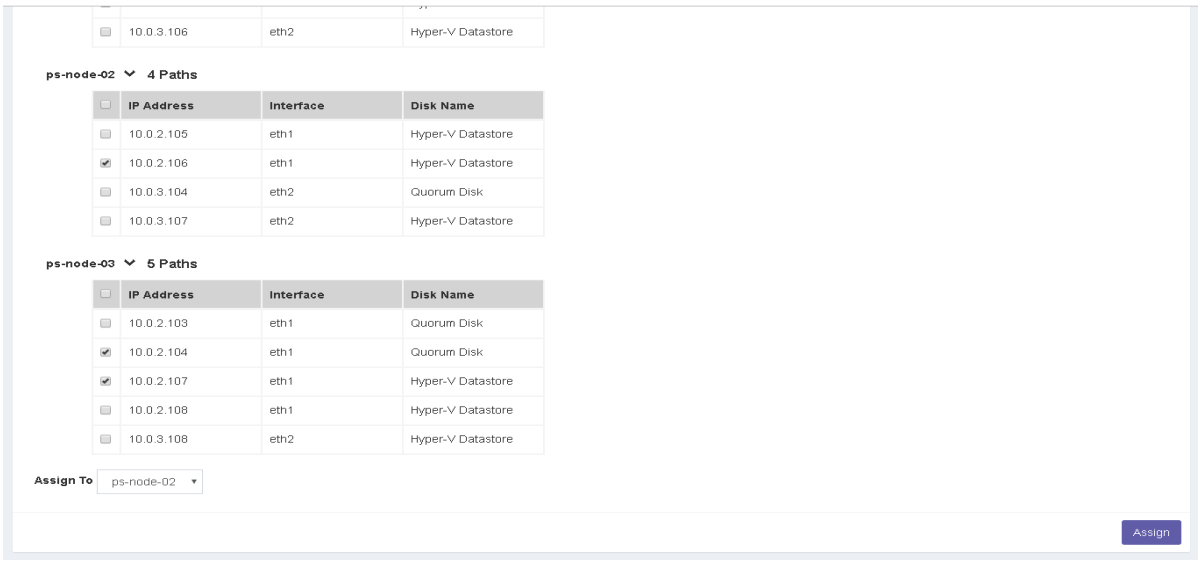

You can select one or more paths then select the destination node

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,

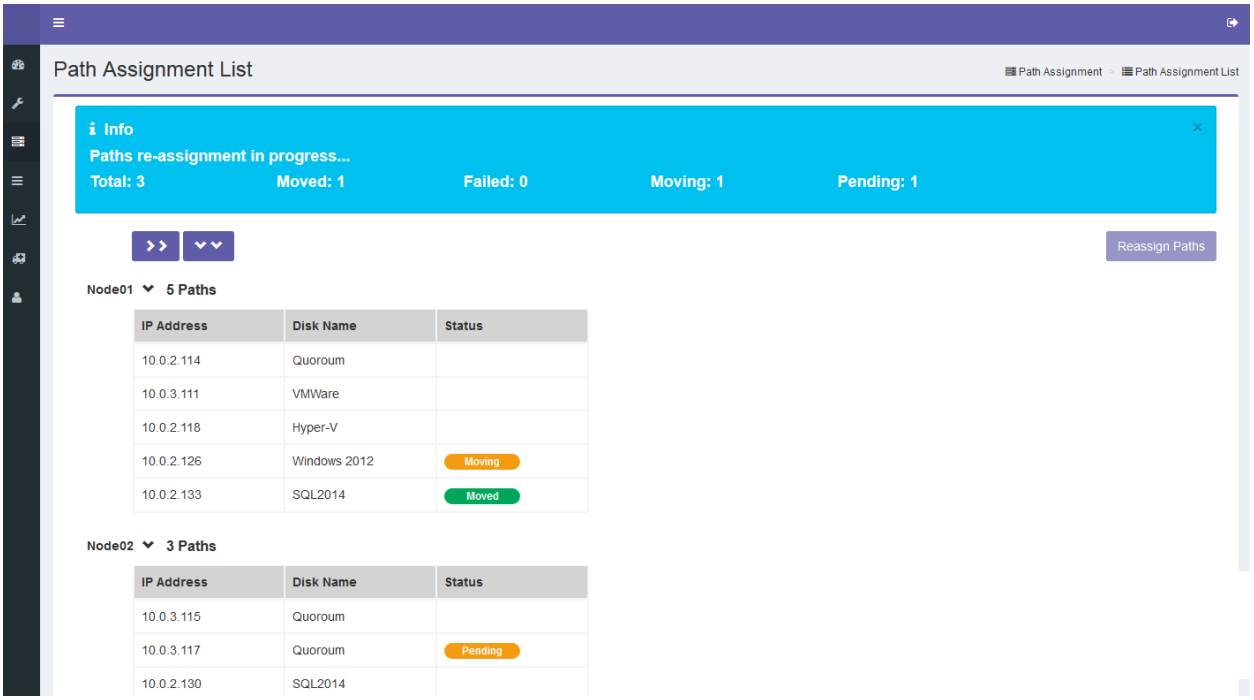

The status can be one of the following:

- 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.

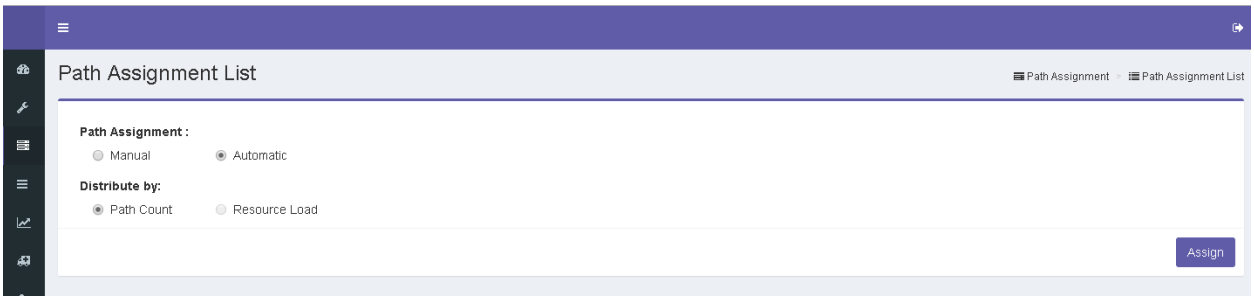

Note: In case you are using VLANs for the iSCSi subnets, system will view the VLAN id next to the interface name ex: eth1.2.

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

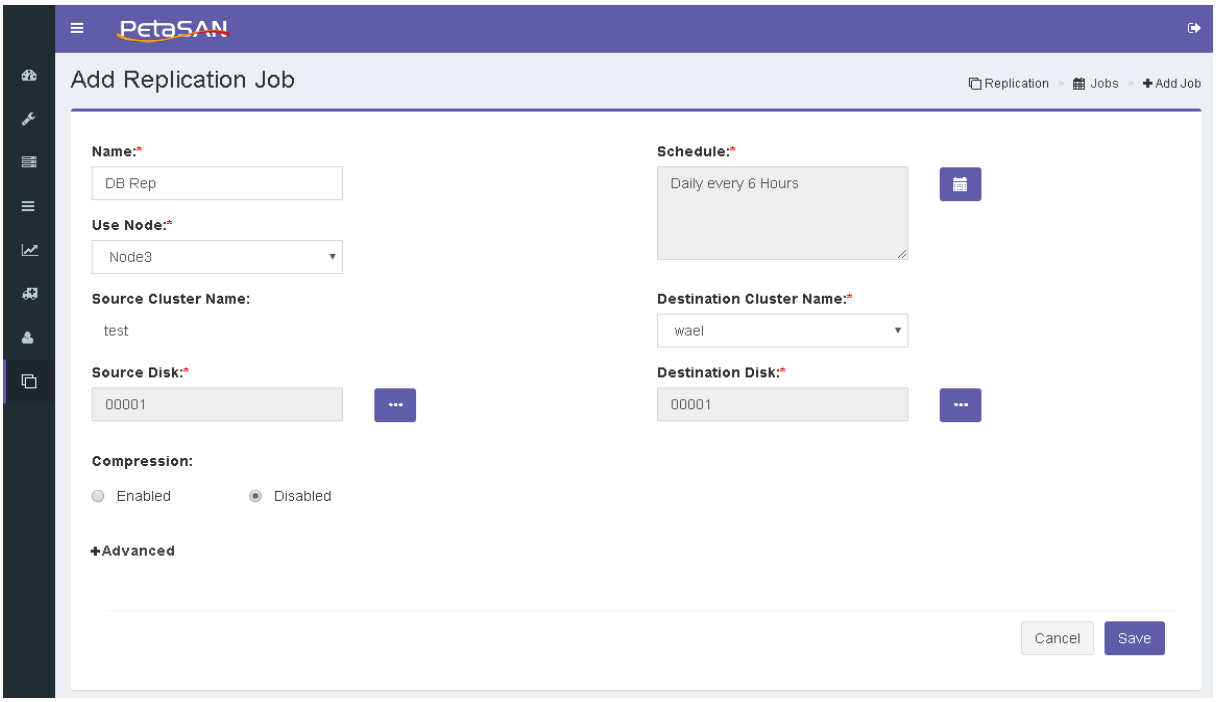

The form requires the following information:

- **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.

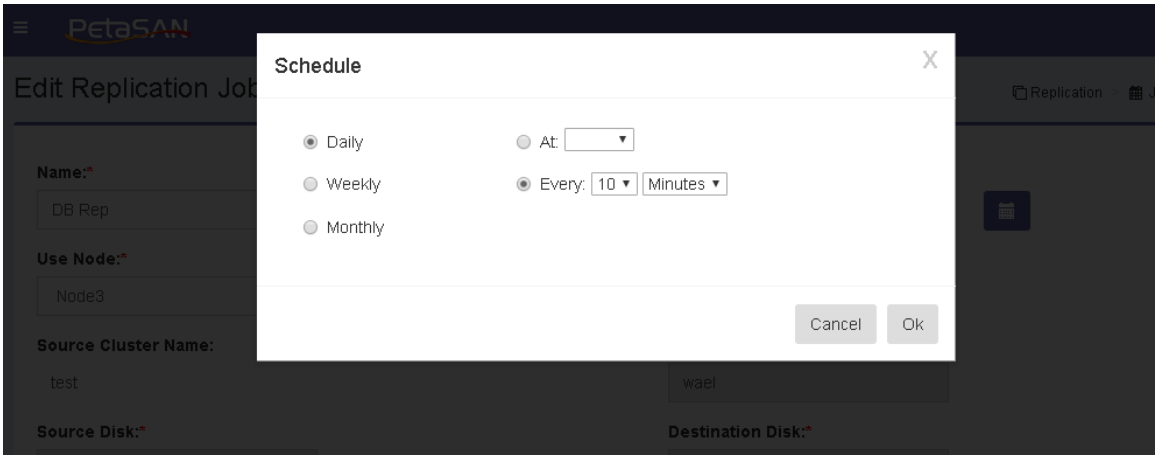

First you should select the job frequency from one of the following options:

- **Daily**

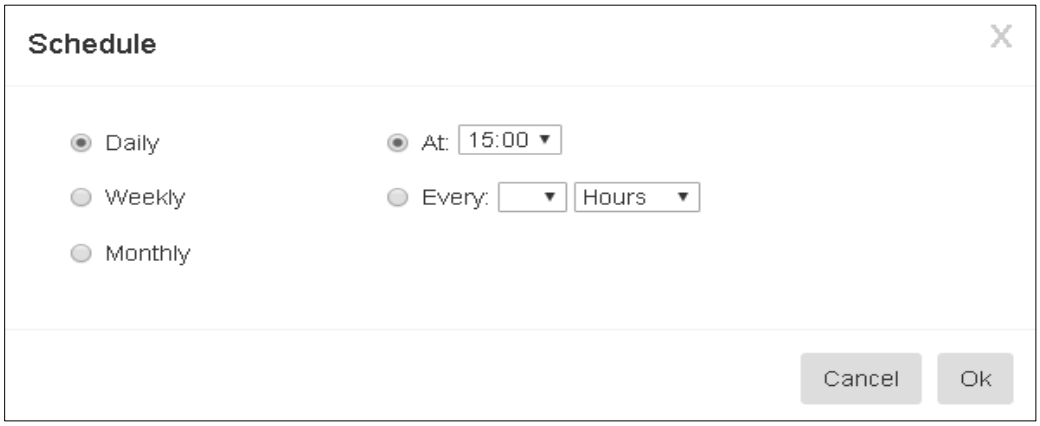

  - **At**
    Job will run daily at a specific hour, you select the hour in 24 hours format
  - **Every**
    Job will run every x minutes or x hours.
- **Weekly**

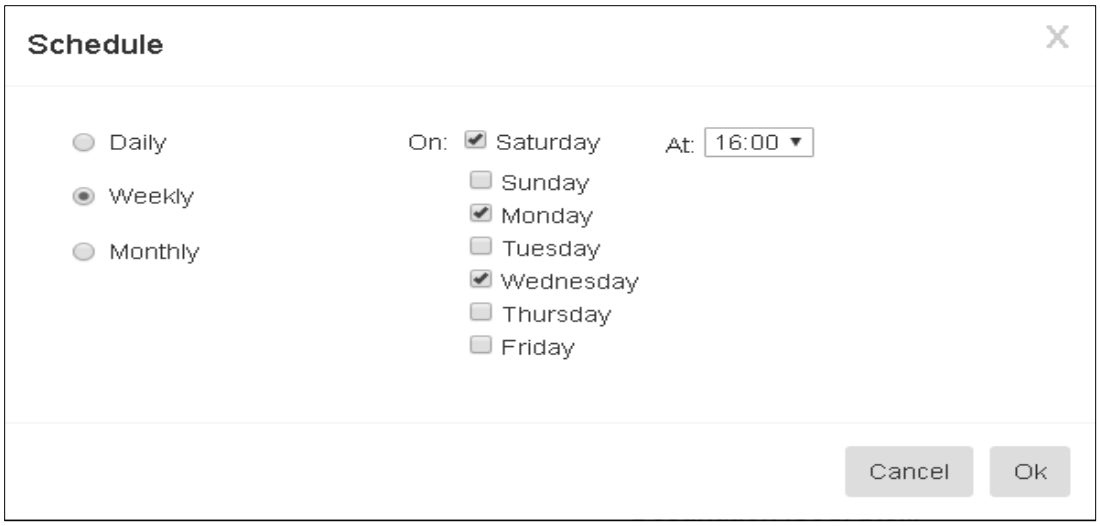

- **Week Days**
  You can select the weeks days the job will run on.
- **At**
  Job will run daily at a specific hour, you will select the hour in 24 hours format.

- **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.

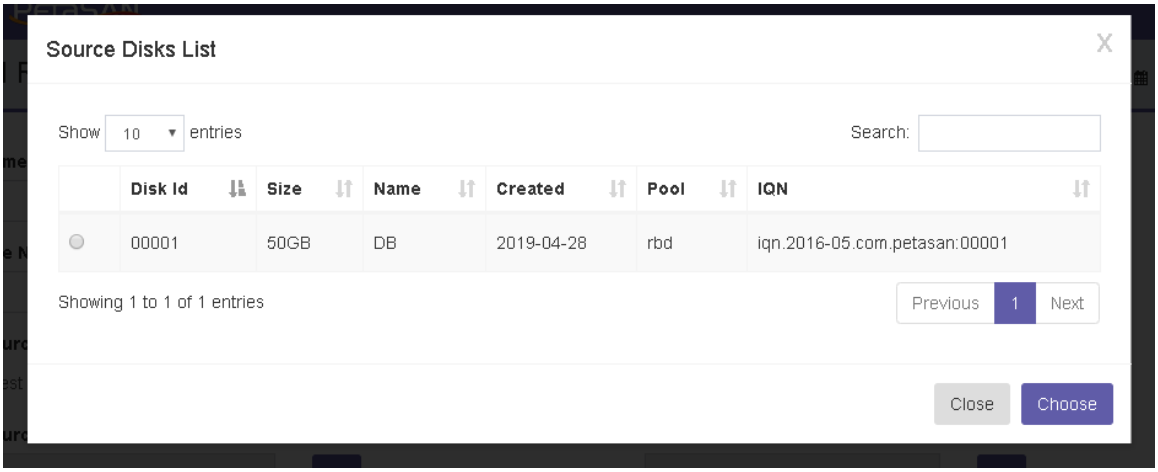

- **Destination Cluster Name**
  The destination cluster name which you can select from the destination clusters you defined; refer to Manage Destination Clusters for more information
- **Destination Disk**
  The system enables you to select an iSCSI disk from the list of replication target disks on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

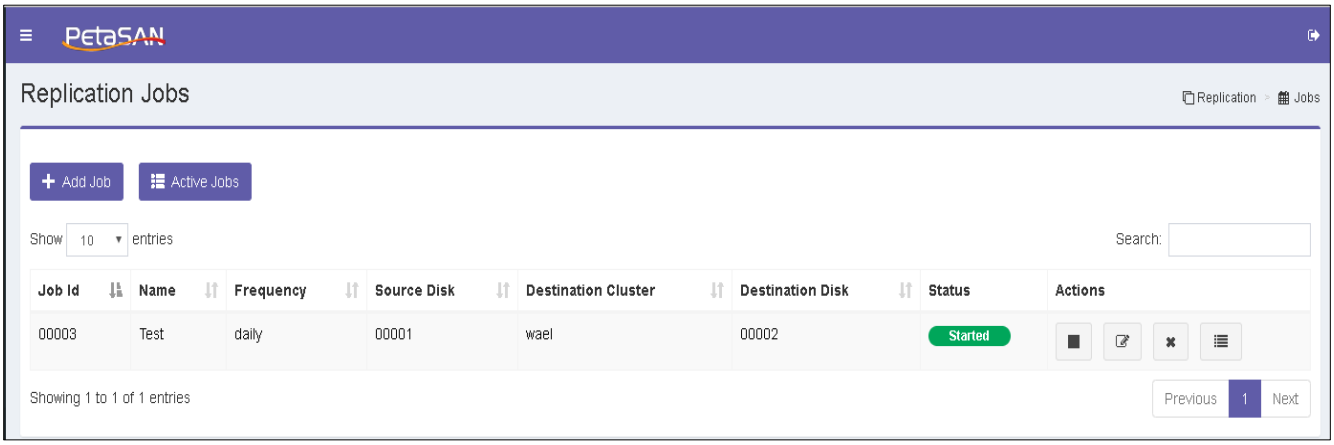

The Jobs List contains the following information for each Job:

- **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.

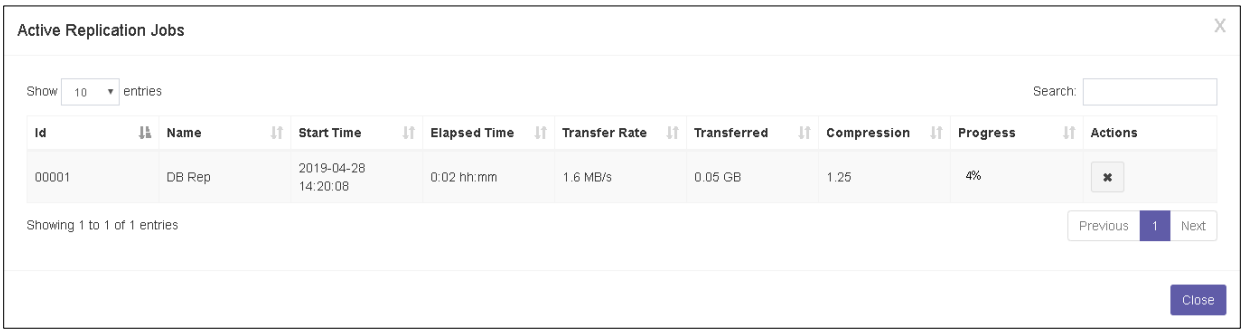

The Active Jobs List displays the following information for each Job:

- **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.

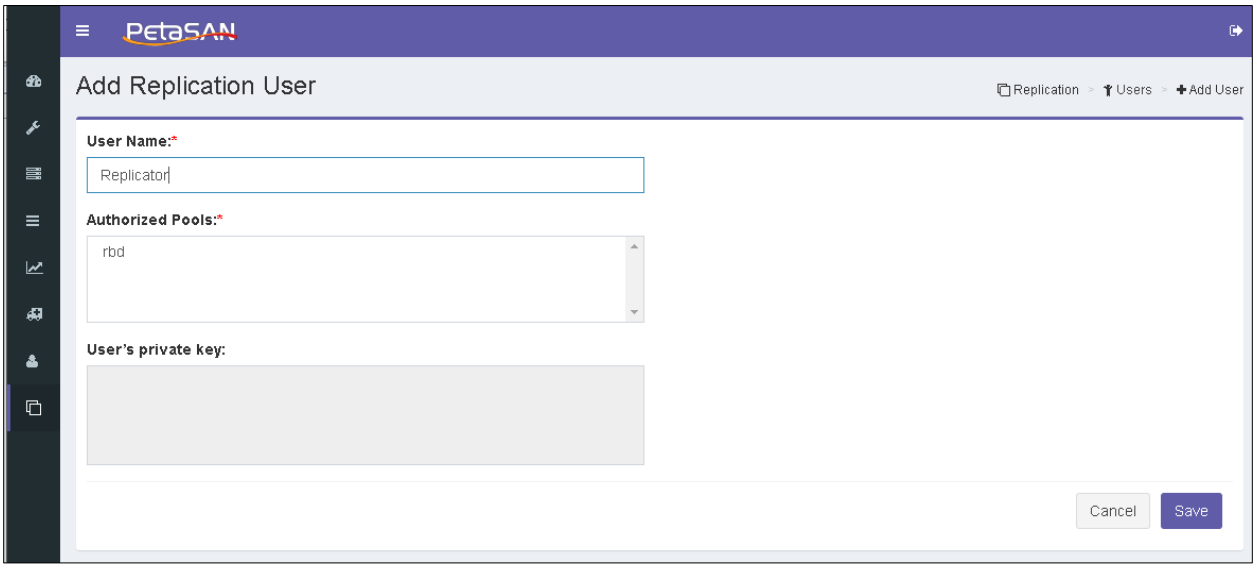

The form requires the following information:

- **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.

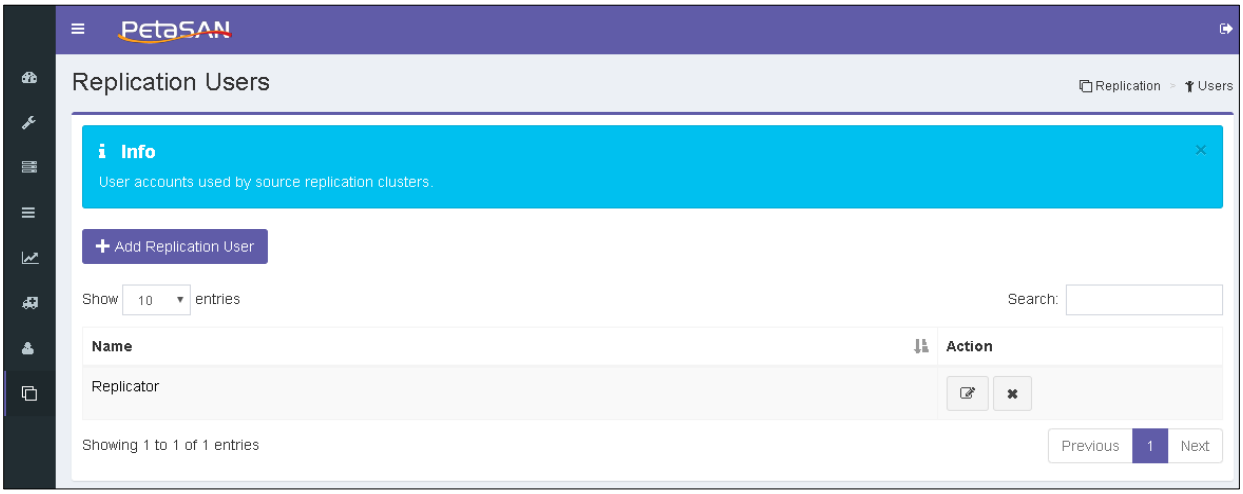

The Replication Users list shows the following information for each user:

- **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.

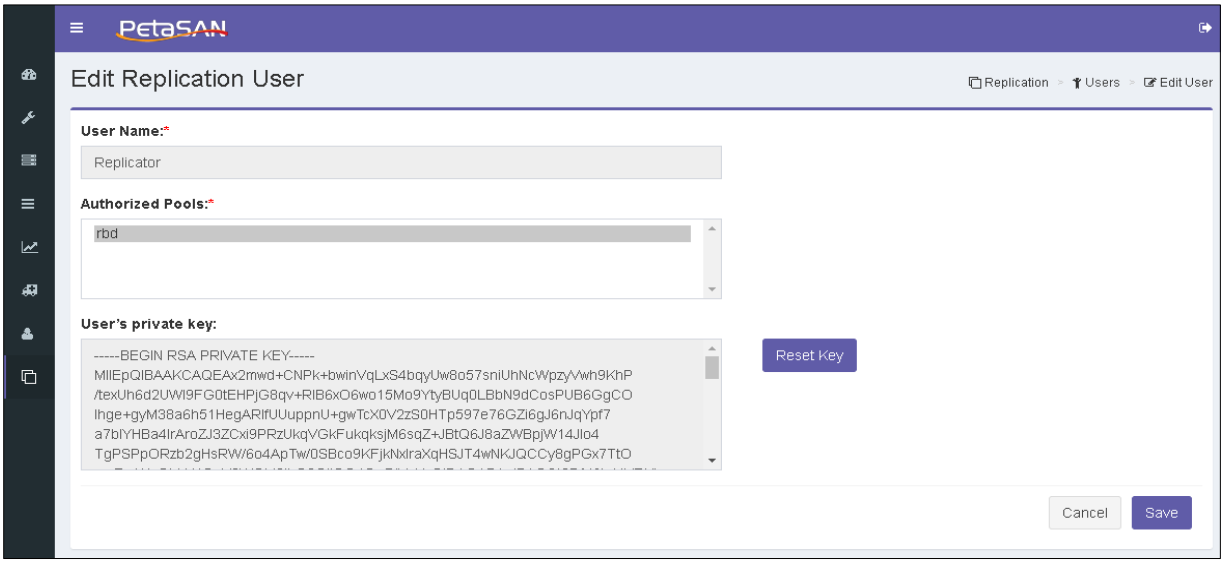

  - Delete
    The system deletes the selected user from the system OS and from Ceph users.

➢ *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

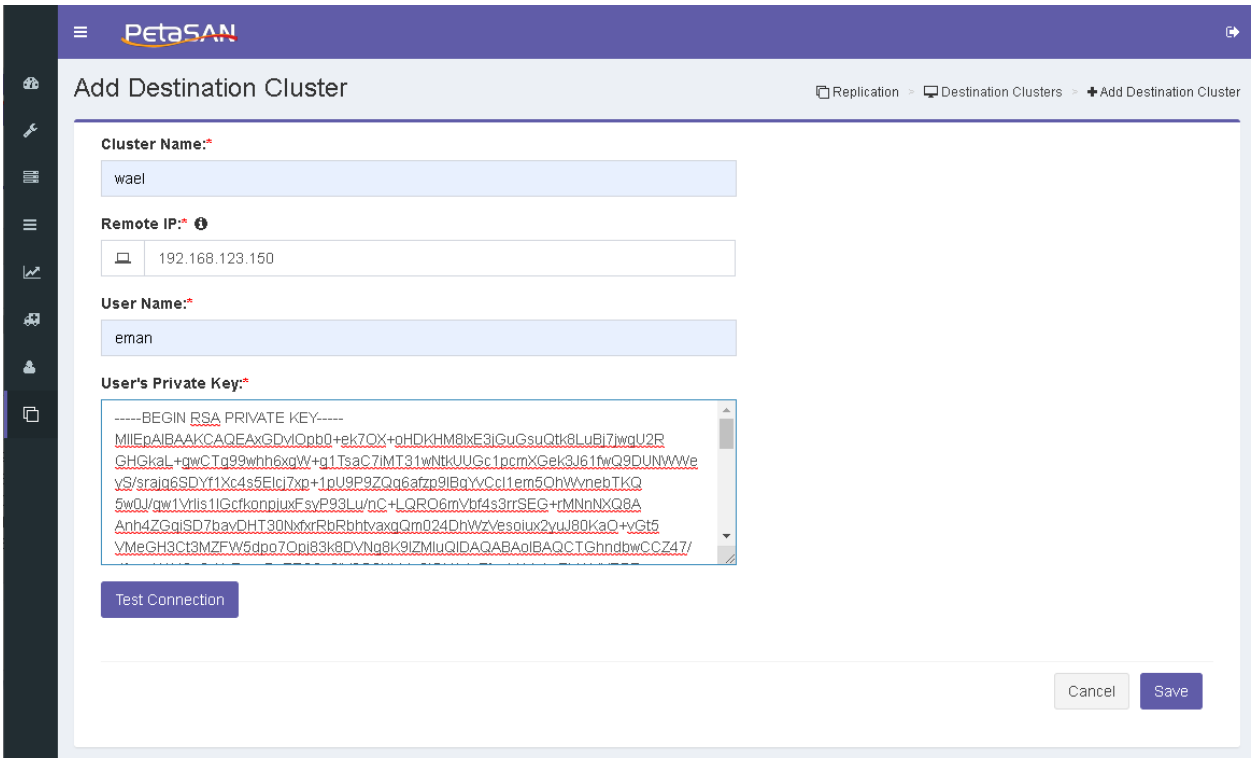

The Destination Cluster from contains the following fields:

- **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 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

➢ View CIFS Shares

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.

- **Browsable**
  Indicates if the share is browsable by clients.
- **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.
- **Actions**
  - Add
    Opens the new share form
  - Delete
    Delete the current share including all stored data.

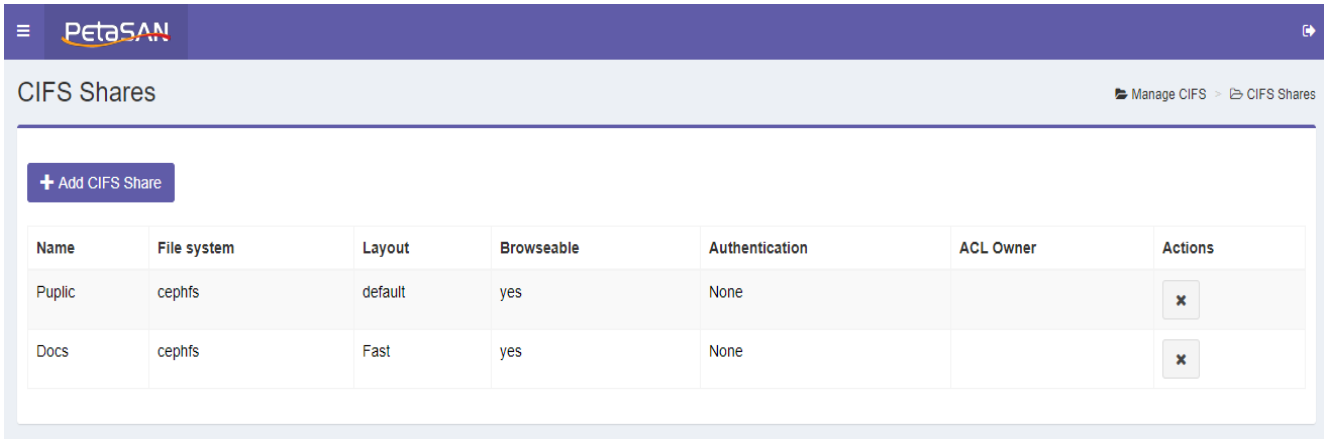

➢ Add CIFS Share

The system opens the CIFS Share form

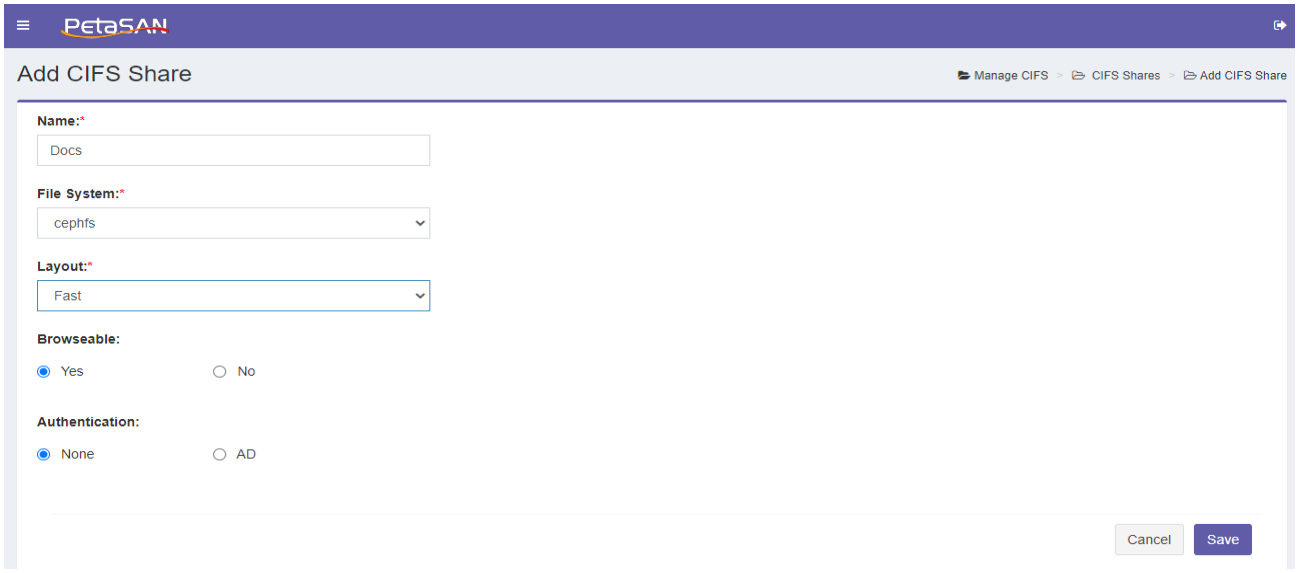

The CIFS Share form contains the following fields:

- **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.
- **Browsable**
  Indicates if the share is browsable by clients.
- **Authentication**
  You can select no authentication or AD; in case of you select AD you will be required to enter the ACL owner.
- **Actions**
  - Save
    Saves the added share

7.2. CIFS Status

➢ View CIFS Status

You can view the current CIFS status

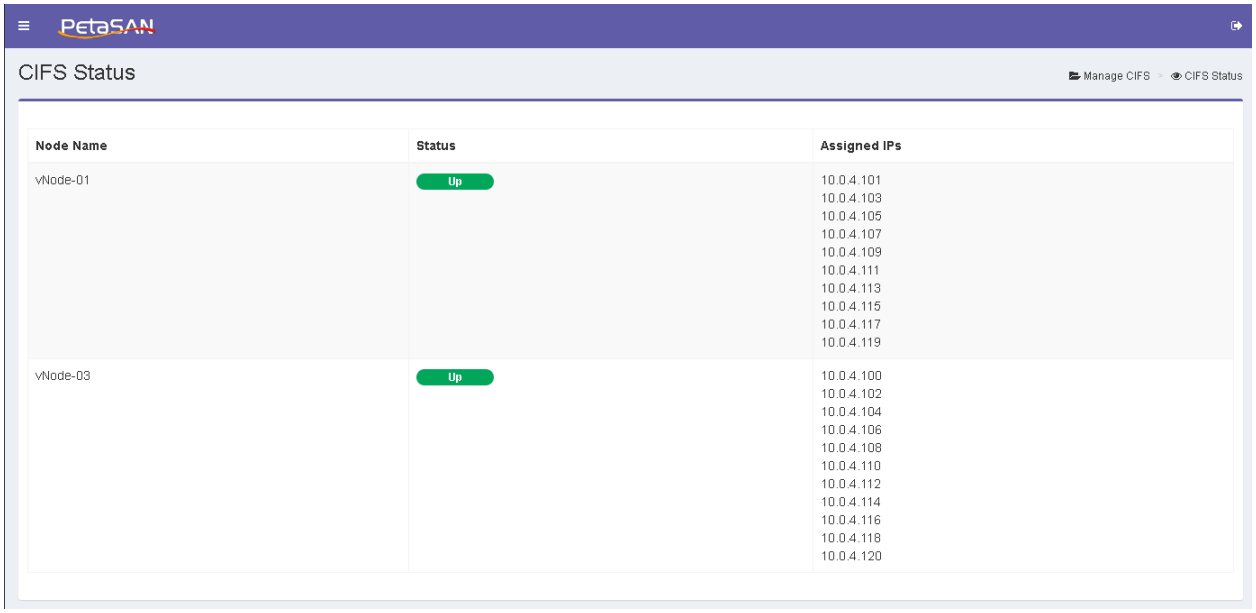

The system will view the following columns:

- **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

8. Manage NFS

8.1. NFS Exports

➢ View NFS Exports

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

- **Name**
  Export name
- **File System**
  The file system in which the export is created.
- **Layout**
  The layout in which the export is created.

- **Actions**
  - Add
    Opens the new export form
  - Delete
    Deletes the current export including all stored data.

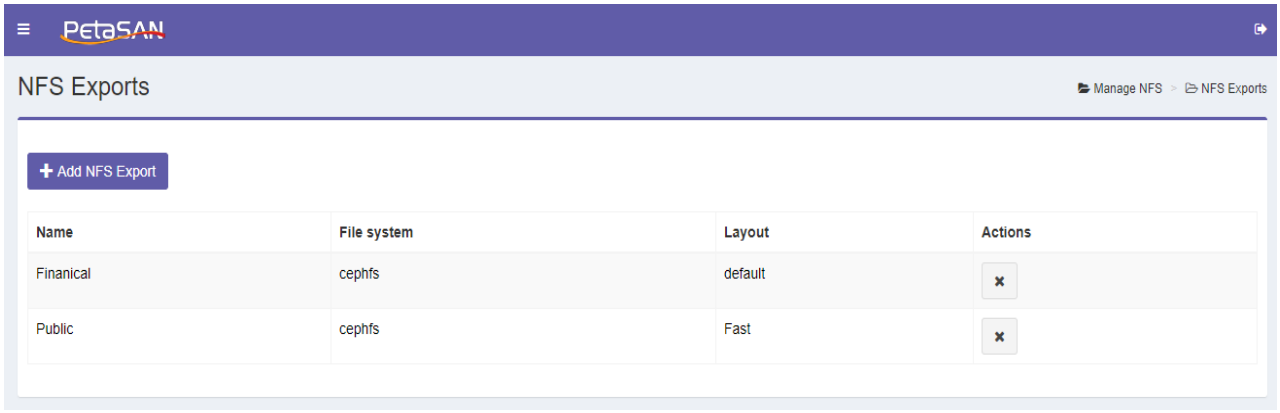

➢ **Add NFS Export**

The system opens the NFS export form

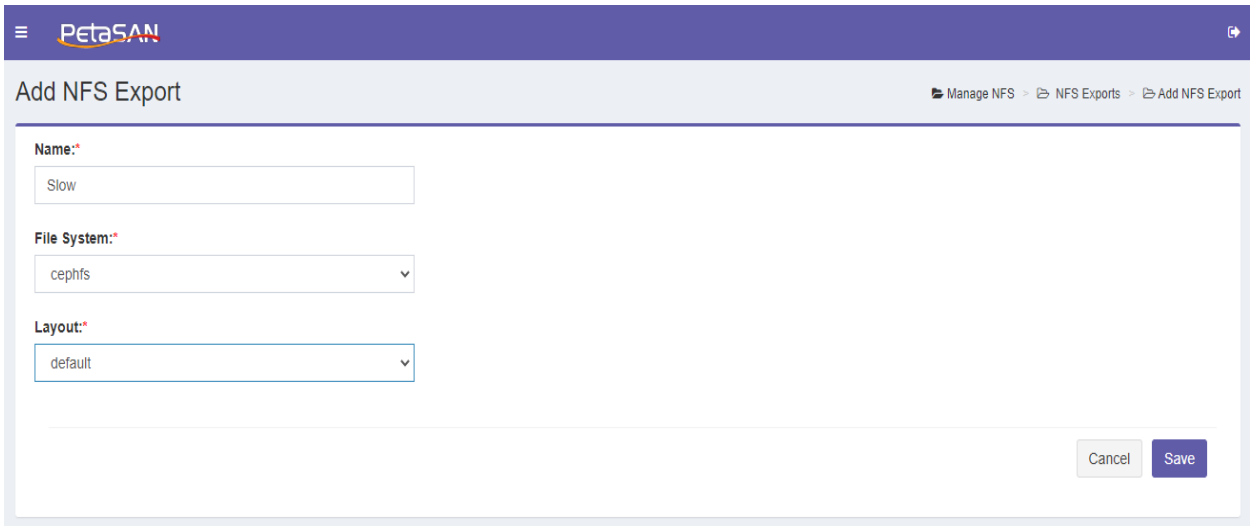

The NFS export form contains the following fields:

- **Name**

  The export name
- **File System**

  The file system on which the export will be created.

- **Actions**
  - Save

    Saves the added export

8.2. NFS Status

➢ **View NFS Status**

You can view the current NFS status

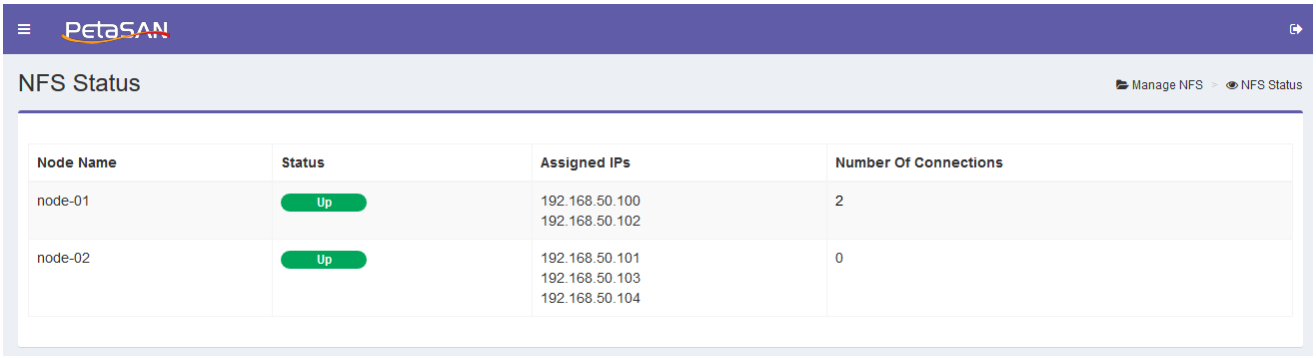

The system will view the following columns:

- **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 connecions
- **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

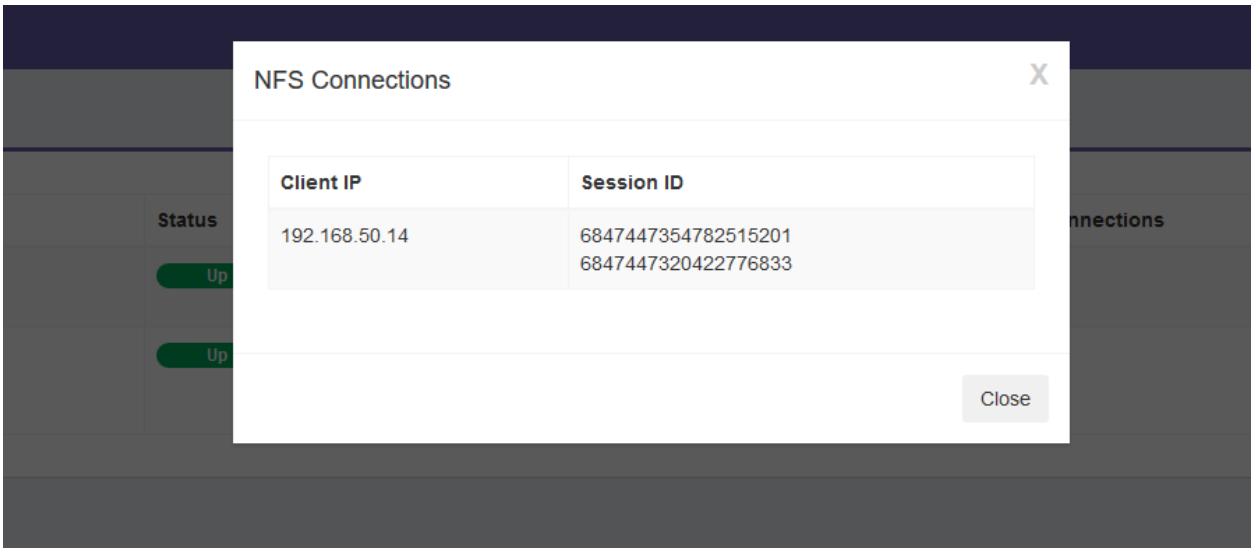

The system will view the following columns for each active connection:

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

9. Manage S3

9.1. S3 Users

➢ View S3 Users List

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

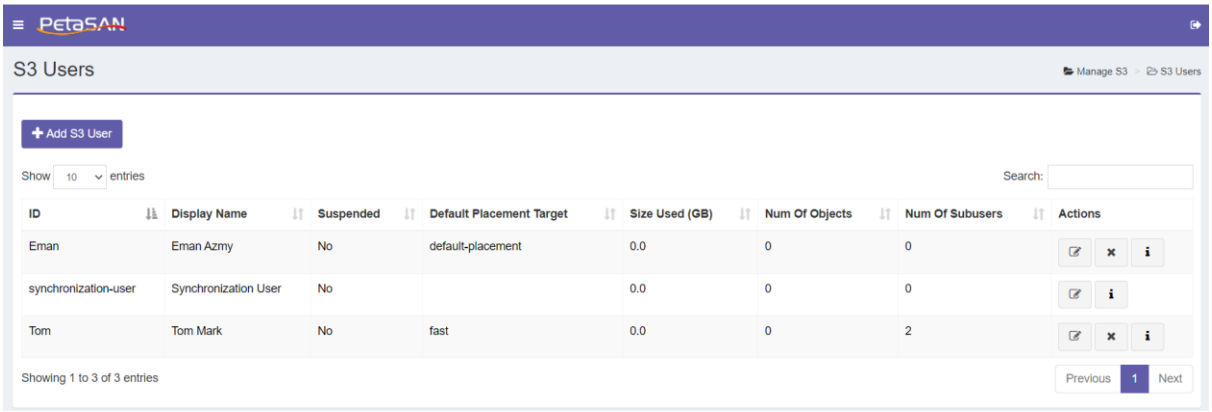

- **ID**
  The user ID
- **Display Name**
  The user name.
- **Suspended**
  The status of the user suspended or not.
- **Default Placement Target**
  The user's default placement target.
- **Size Used (GB)**
  The number of Gigabytes stored by the user.
- **Num of objects**
  The number of objects stored 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

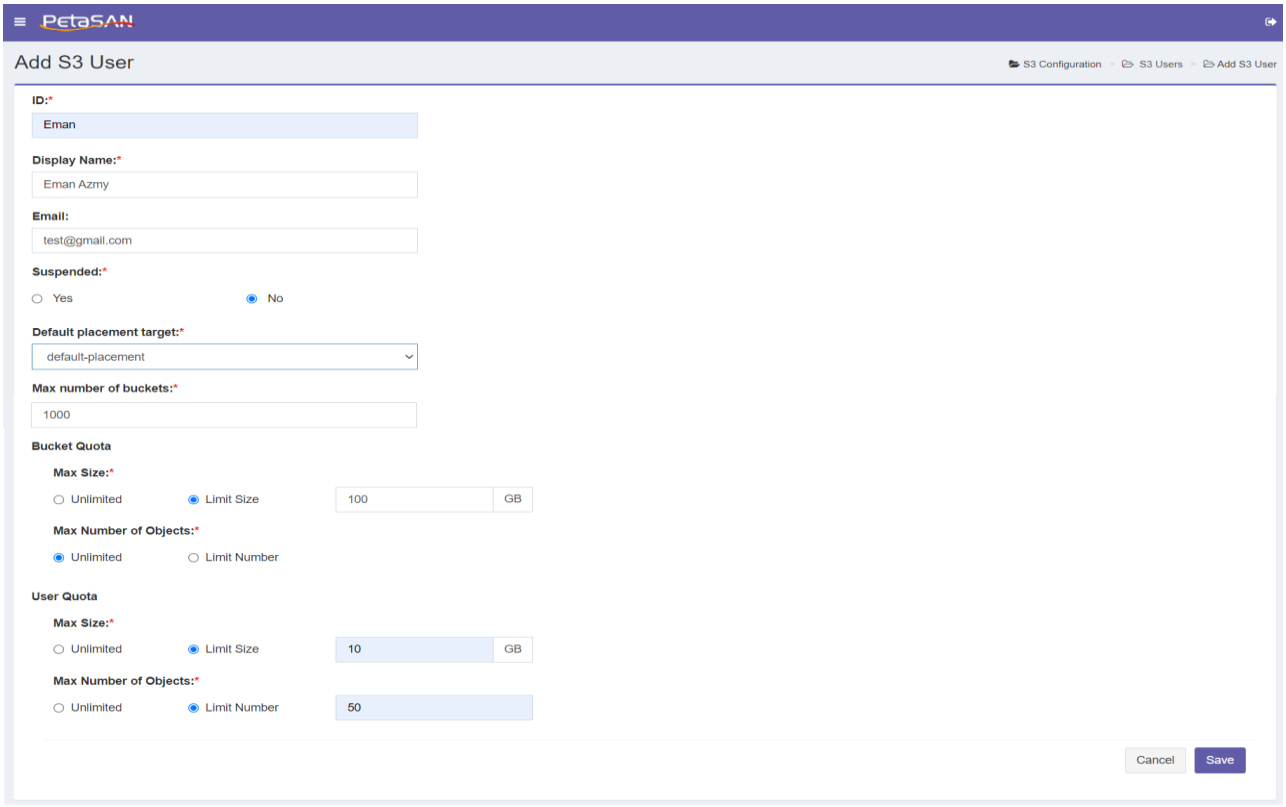

The S3 User form contains the following fields:

- **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 f 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 f 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.

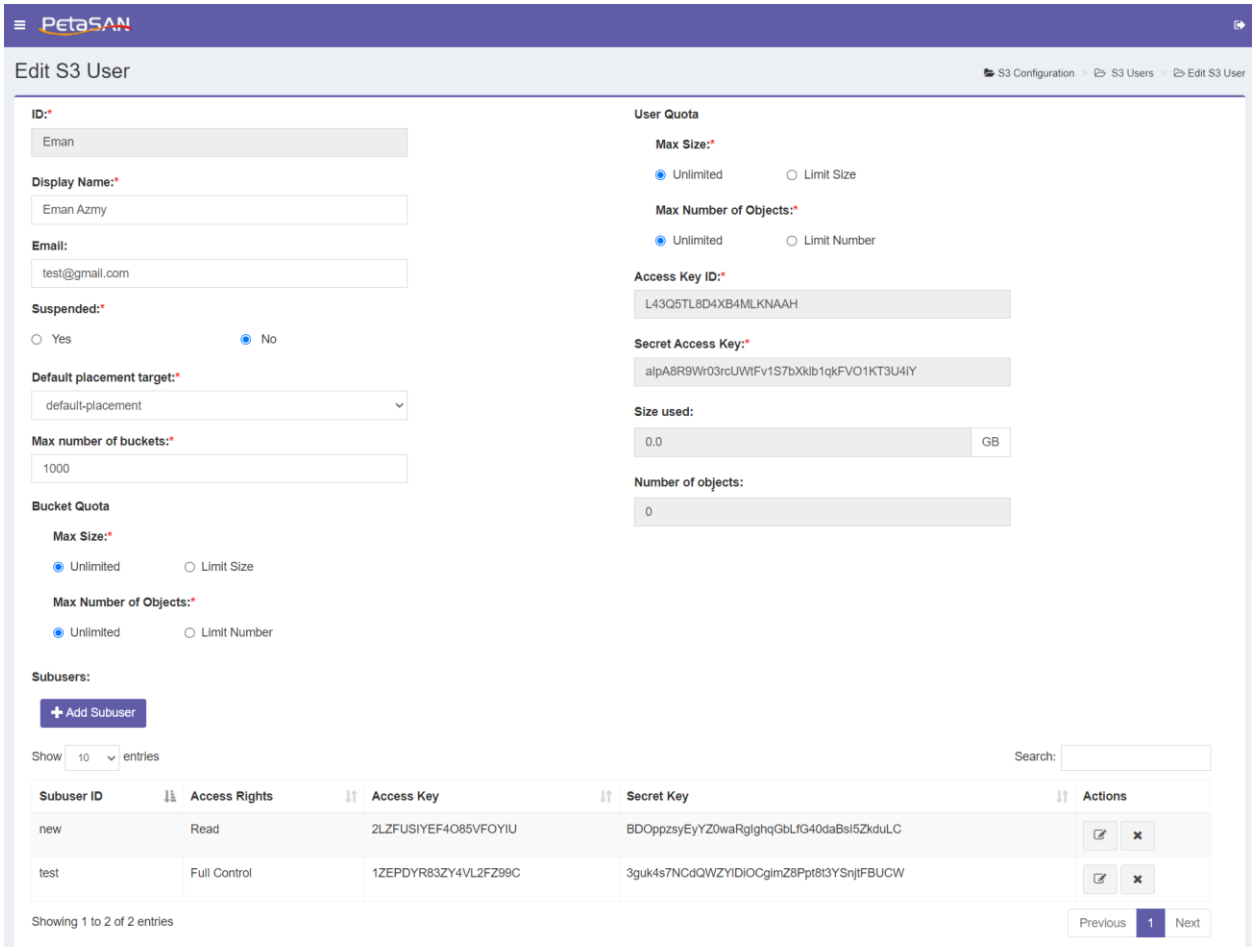

User can add subusers under the user and when added system will view the following columns for each subuser:

- **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.

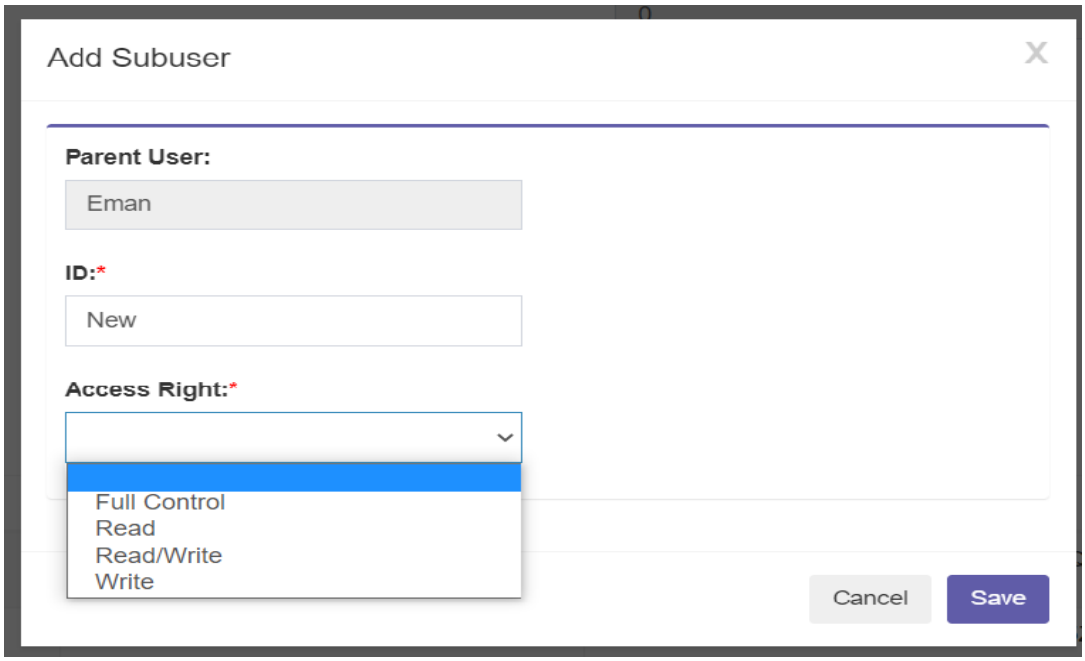

After saving the subuser system will redirects the user to the Edit subuser form showing the subuser's keys

➢ 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.

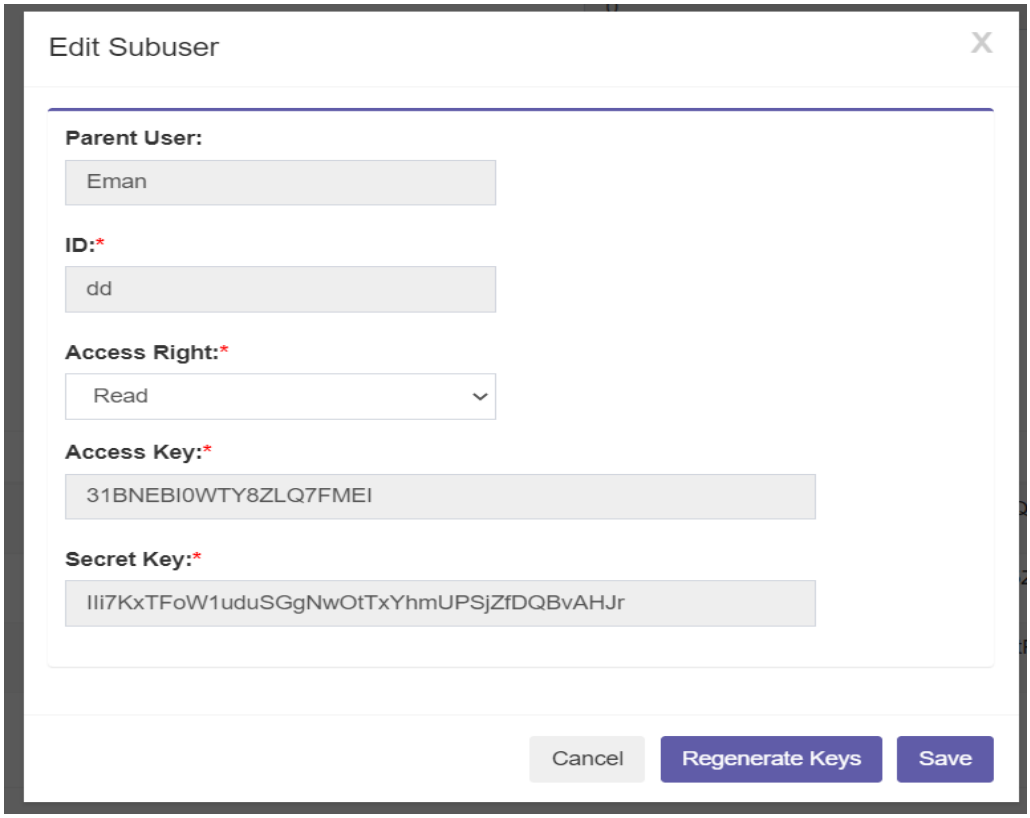

Note:

You can regenerate the subuser's access and secret keys

9.2. S3 Status

➢ View S3 Status

You can view the current S3 status

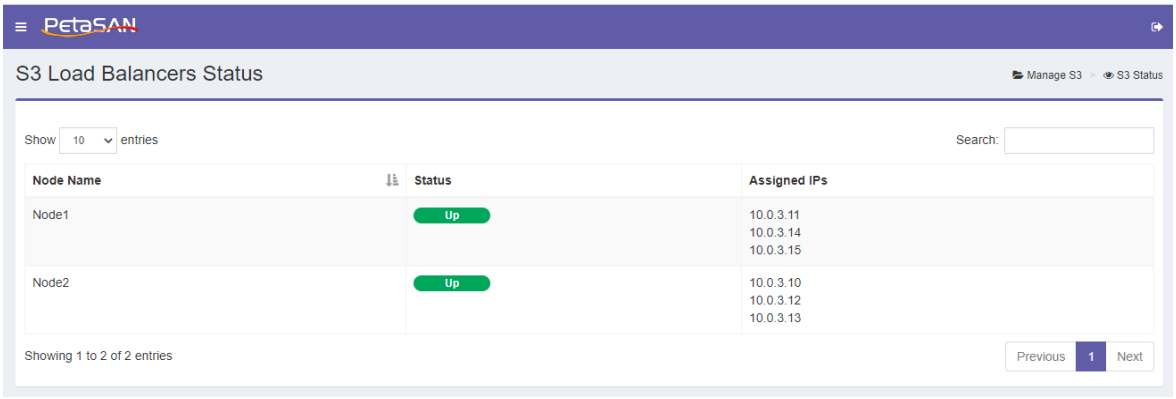

The system will view the following columns:

- **Node Name**
  Name of theS3 Server Node
- **Status**

Node status, up, down

- **Assigned IPs**
  IPs assigned to each node

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

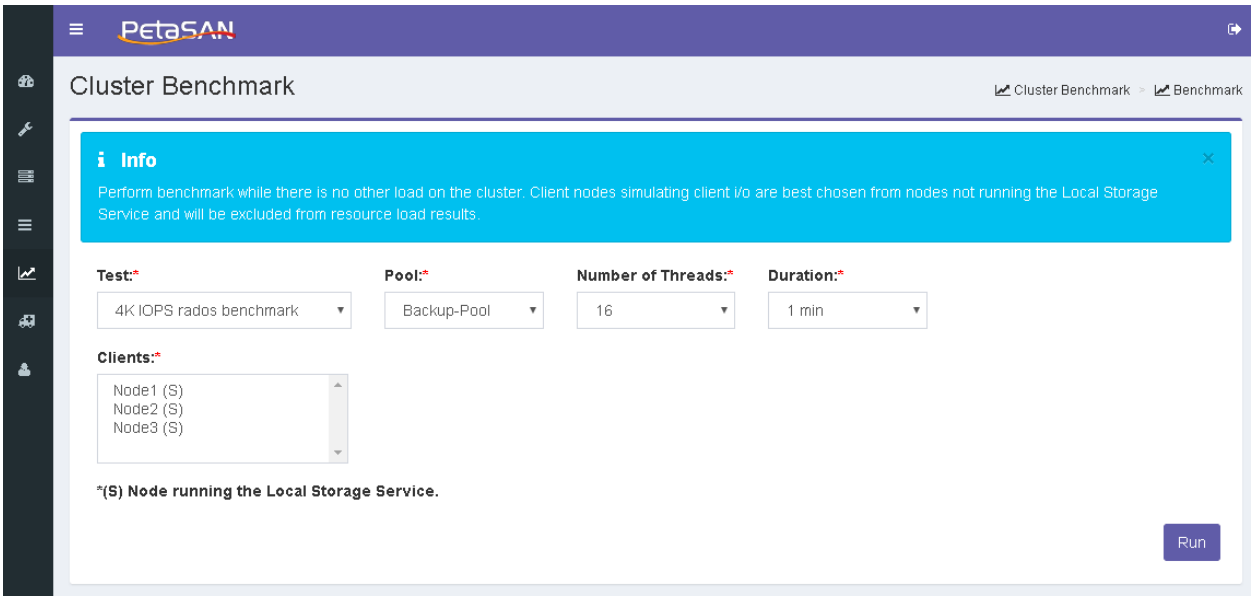

You can customize your test using the following fields:

- **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

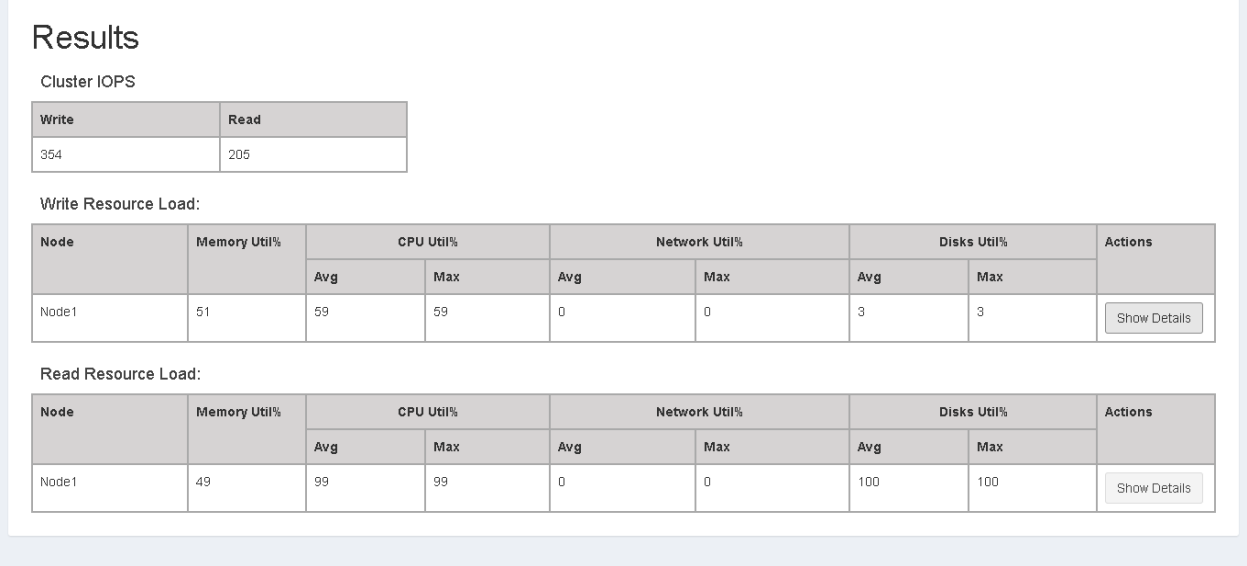

- **Cluster IOPS/Throughput**

  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	rxmcst/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

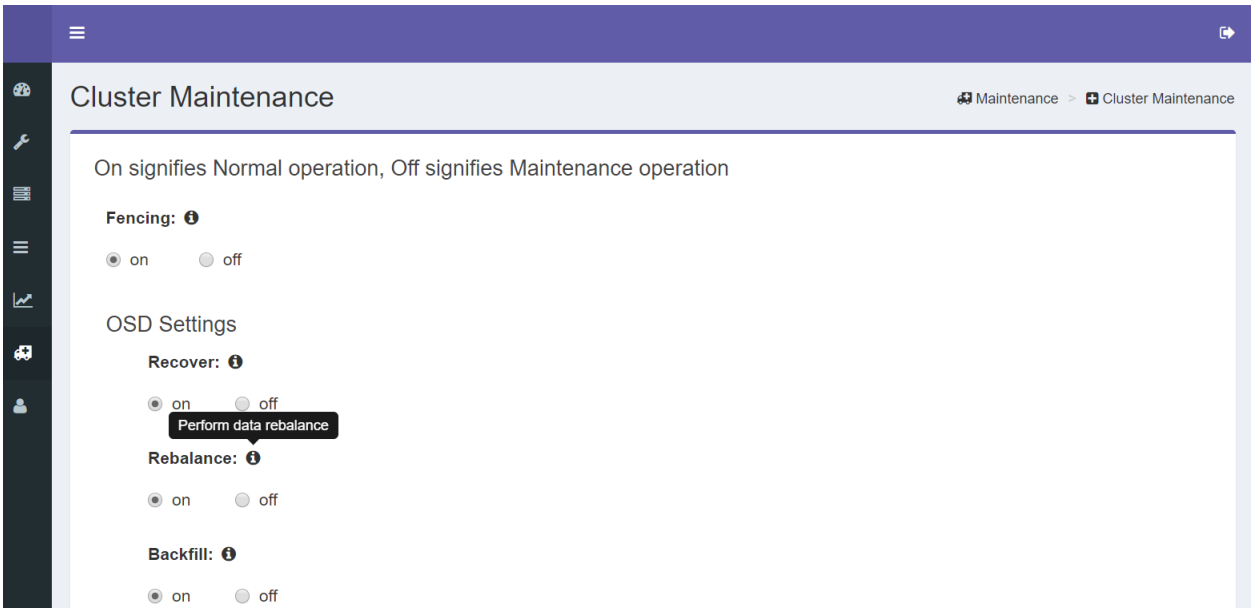

➢ *Notes:*

- *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

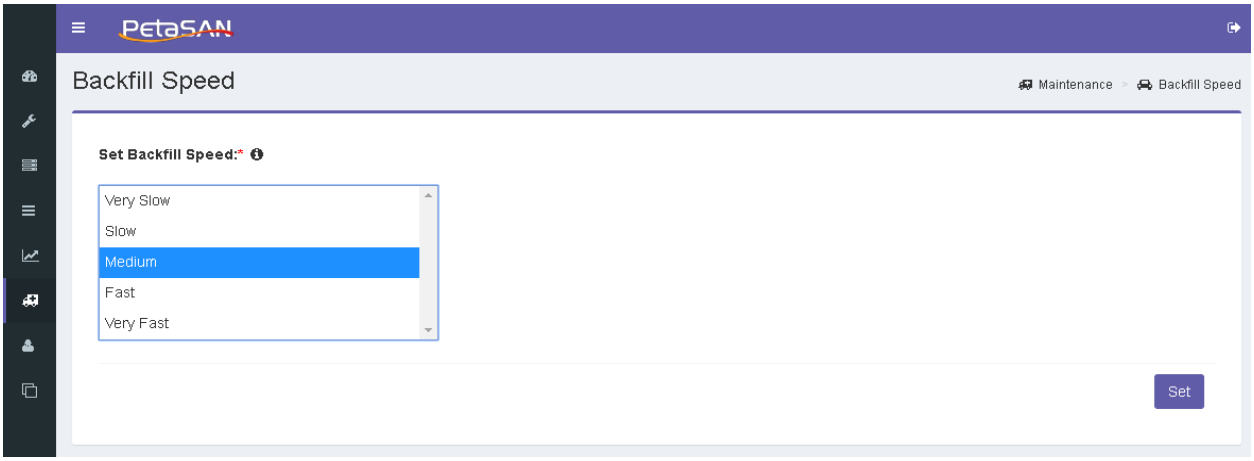

11.3. Scrub Speed

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.

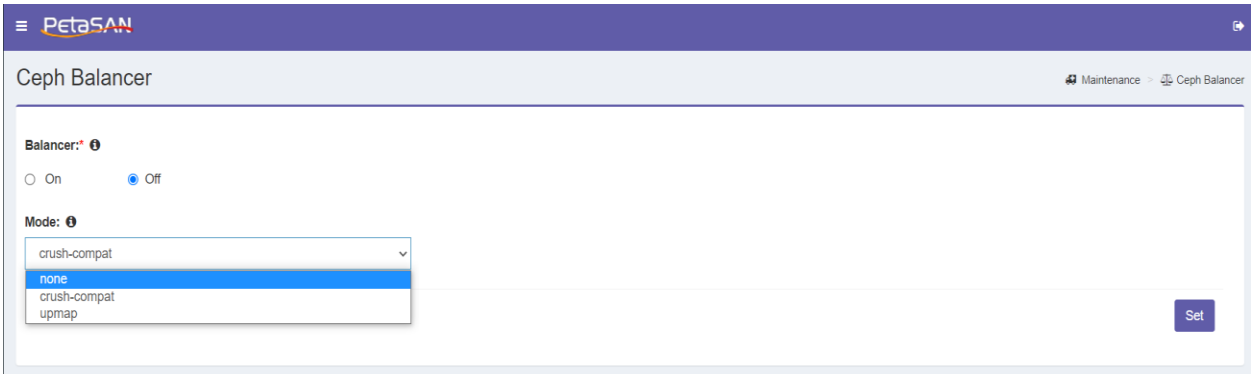

11.5. OSD CRUSH Weight

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.

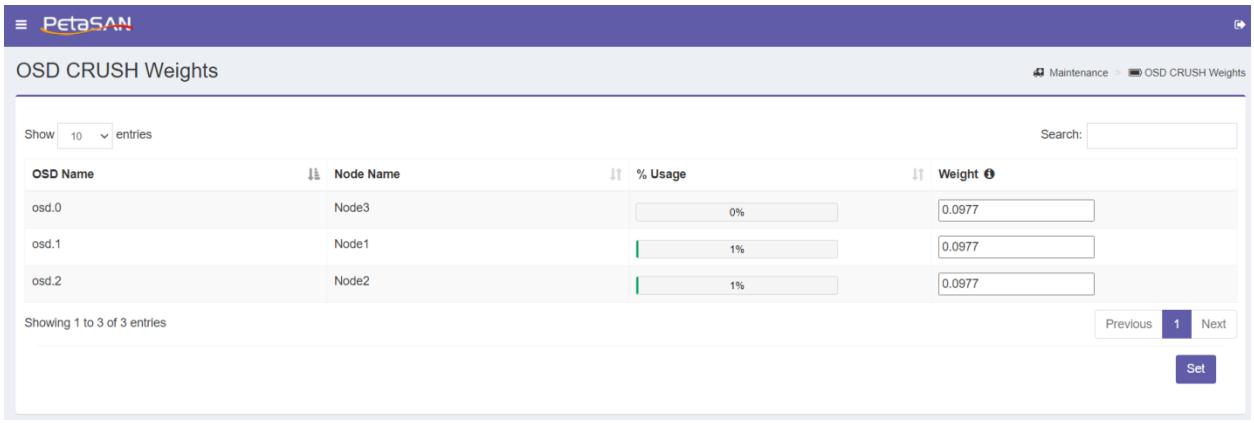

12. Manage Users

➢ Add User

You can add new user by using the Add User Form

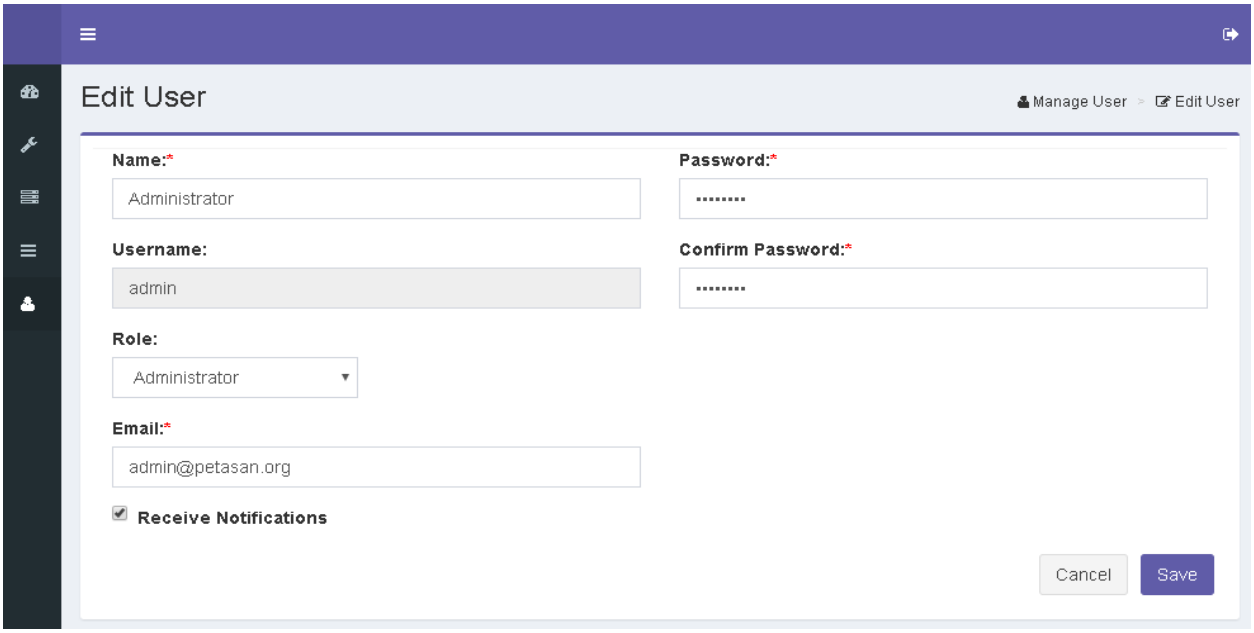

The form requires the following information:

- **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

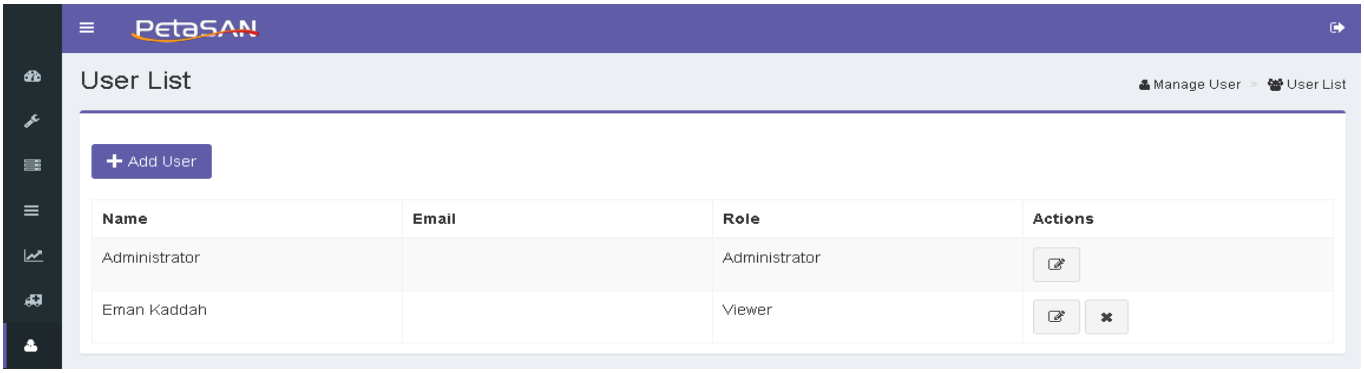

- **Actions**

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

  - Edit

    Enables updating the user information except the username.

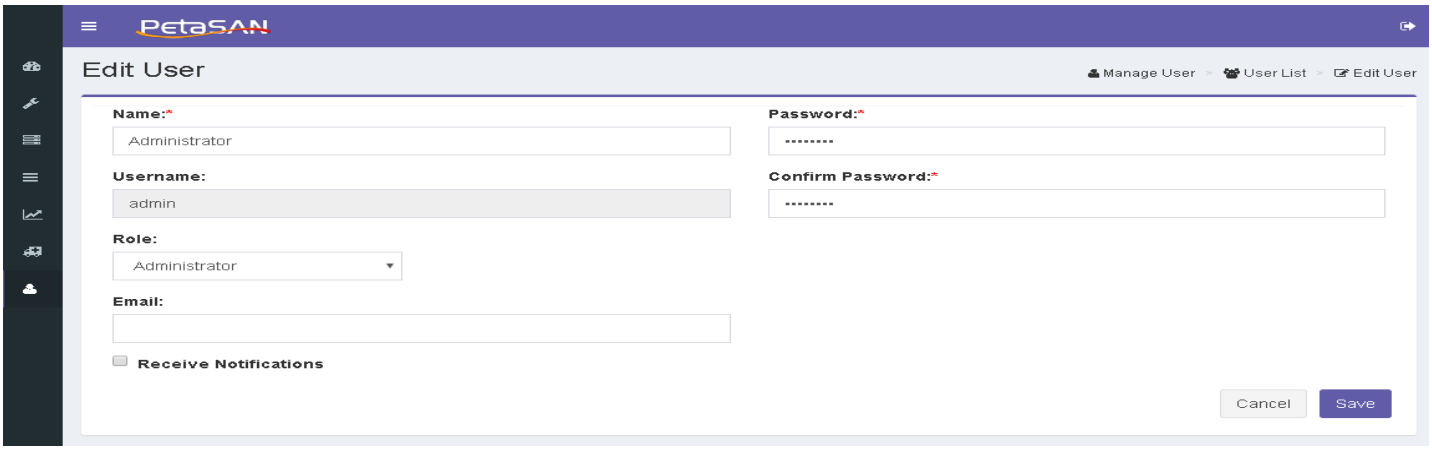

  - Delete

    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