



Quick Start Guide

Version 1.8



Revision History

Date	Version	Description
10-10-2016	1.0	Initial version
19-10-2016	1.1	Update the initial Cluster Management password
16-11-2016	1.2	Update screen shots related to replace management node feature
01-06-2017	1.3	Add Time Setup and Cluster NIC Settings screen shots
29-10-2017	1.4	Add the cluster tuning page and update deployment wizard and dashboard screenshots
08-01-2018	1.5	Update node services and dashboard screen shots
16-12-2018	1.6	Update screen shots based on Release 2.2
10-03-2020	1.7	Updates based on Release 2.5
27-07-2020	1.8	Format change, support Release 2.6
27-06-2021	1.9	Updates based on Release 2.8 (S3)



Contents

1. Purpose	4
2. Planning the cluster network	4
3. Node Installation	5
4. Node Deployment	8
Node 1	8
Node 2	12
Node 3	15
Nodes 4+	17
5. iSCSI Setup	18
6. CIFS/SMB Setup	22
7. NFS Setup	25
8. S3 Setup	28



1. Purpose

The purpose of this guide is to quickly get you up and running using PetaSAN. It is recommended to be used as the first introductory guide. It will go through the main stages: cluster planning, node installation, node deployment and cluster building, setup for iSCSI, CIFS, NFS and S3 services.

2. Planning the cluster network

PetaSAN requires several subnets:

Management: This is used for management traffic.

Backend: This is a core backend subnet, internally it is used by several internal components such as the Ceph storage engine, Consul service mesh platform, GlusterFS shared configuration and stats data.

iSCSI 1,2: These are 2 subnets for iSCSI client access using MPIO.

CIFS/SMB: Subnet used for CIFS/SMB client access.

NFS: Subnet used for NFS client access.

S3: Subnet used for S3 client access.

These subnets need to be separate and do not overlap. They could still share network interfaces, so it is not required to have a separate interface for each. It is important to plan these networks prior to PetaSAN deployment. Note that much of the network settings can be changed even after cluster deployment, with the exception of the Backend network which is more complex to change.

In this guide, we will setup 3 nodes each with 4 interfaces, using the following configuration:

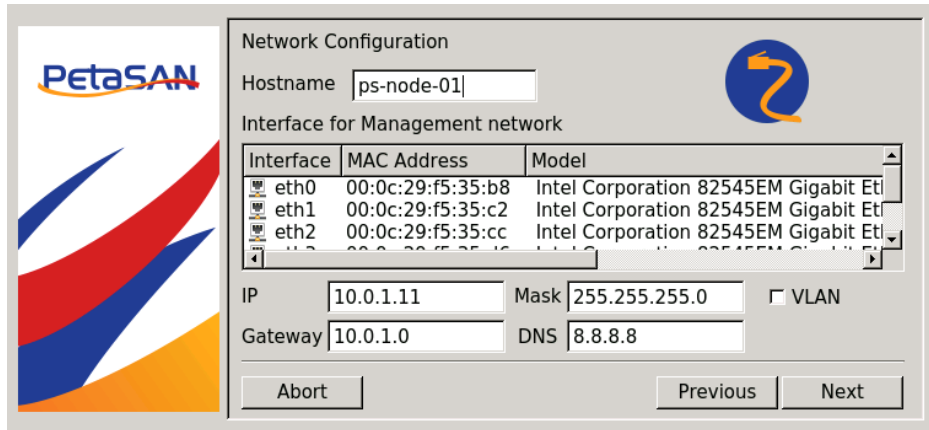
	First node	Second node	Third node
Hostname	ps-node-01	ps-node-02	ps-node-03
Management Interface	eth0		
Management IP	10.0.1.11	10.0.1.12	10.0.1.13
Backend Interface	eth1		
Backend IP	10.0.2.11	10.0.2.12	10.0.2.13
iSCSI 1 Interface	eth2		
iSCSI 1 IPs	Shared virtual IPs: 10.0.3.100 to 10.0.3.110		
iSCSI 2 Interface	eth3		
iSCSI 2 IPs	Shared virtual IPs: 10.0.4.100 to 10.0.4.110		
CIFS/SMB Interface	eth2		
CIFS/SMB IPs	Shared virtual IPs: 10.0.5.100 to 10.0.5.110		
NFS Interface	Eth2		
NFS IPs	Shared virtual IPs: 10.0.6.100 to 10.0.6.110		
S3 Interface	Eth2		
S3 IPs	Shared virtual IPs: 10.0.7.100 to 10.0.7.110		

3. Node Installation

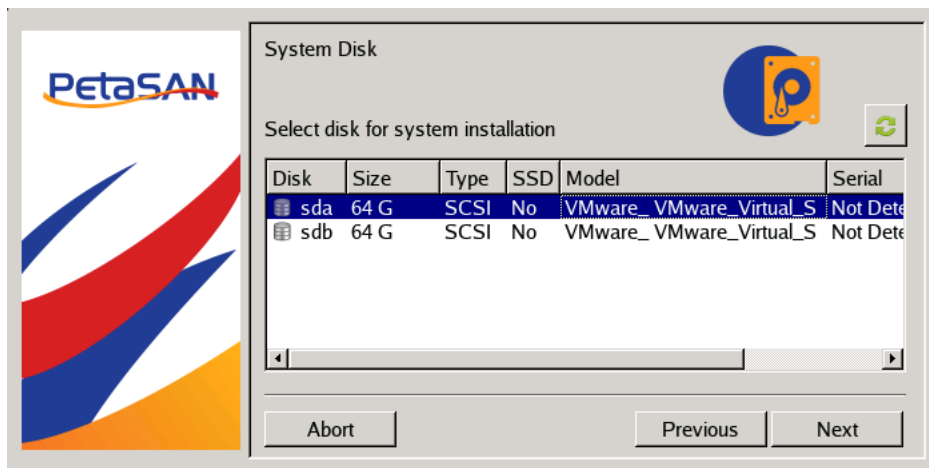
The installation iso can be burned to USB using widely available USB tools such as rufus www.rufus.org.

There are 3 main settings to perform during the installer:

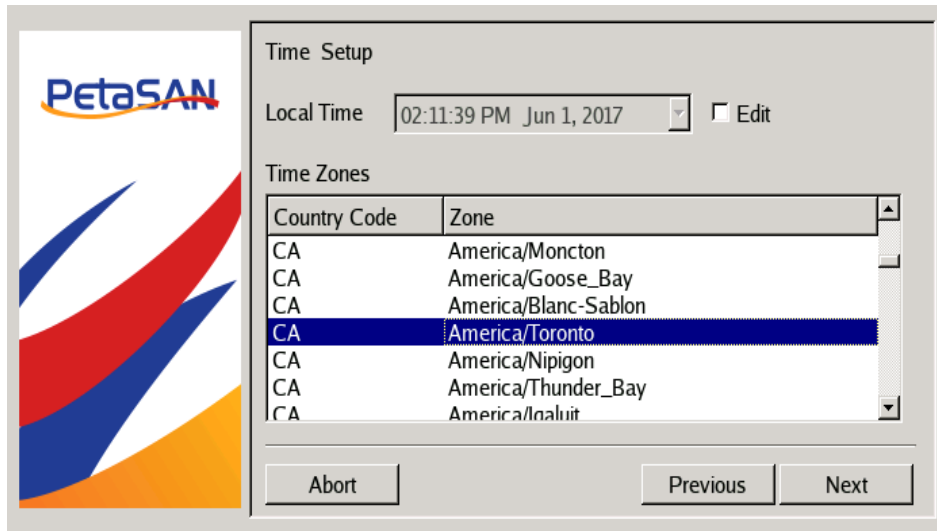
Network configuration, here we define the hostname of the node and setup its management interface.



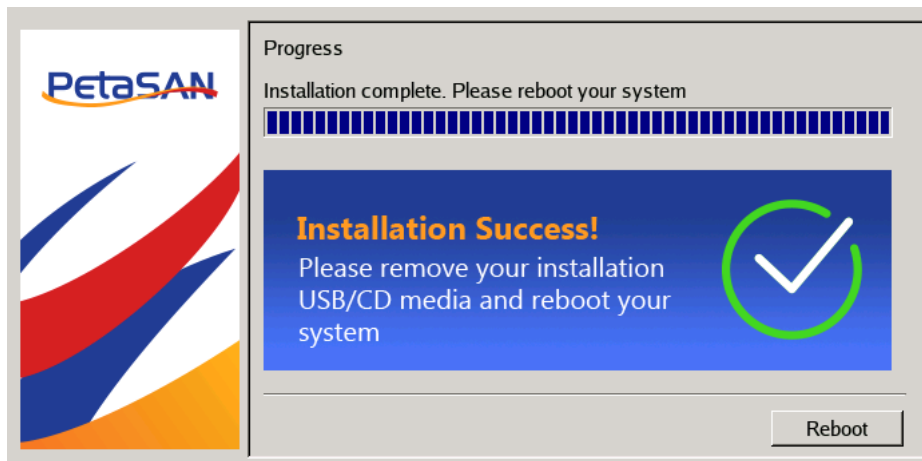
System Disk selection, here we select the disk to install the PetaSAN system.



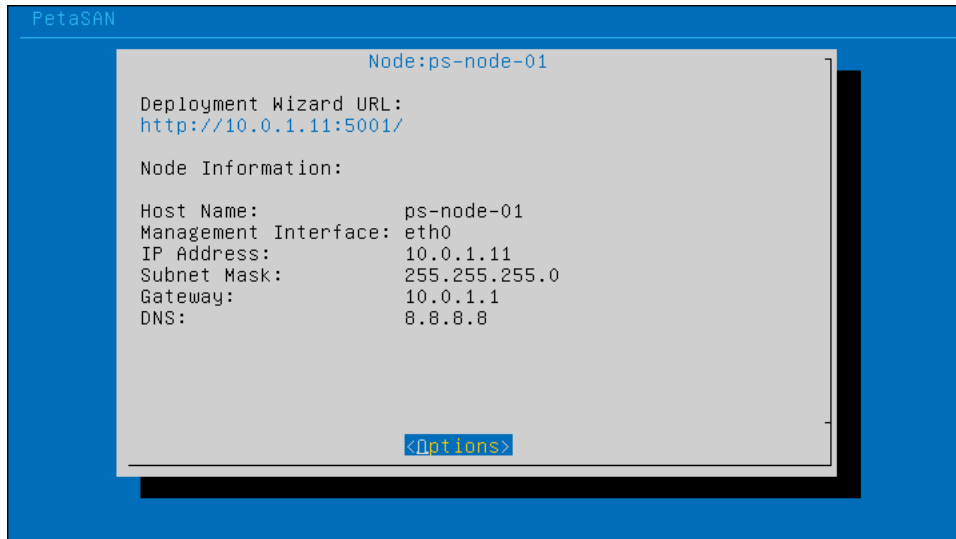
Time Setup, adjusting the machine time and time zone.



Upon successful completion, we need to remove the install media before rebooting



On reboot a console will be display node data as well list the url for the Deployment Wizard



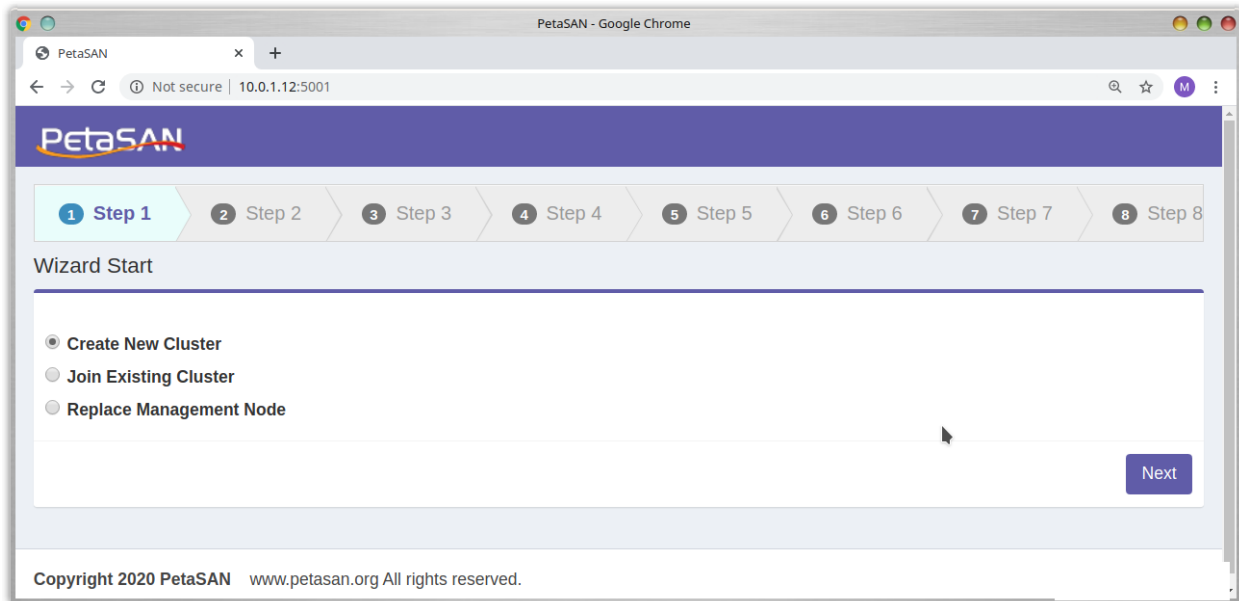
We need to repeat the above steps for installing the remaining nodes:

	First node	Second node	Third node
Hostname	ps-node-01	ps-node-02	ps-node-03
Management Interface	eth0		
Management IP	10.0.1.11	10.0.1.12	10.0.1.13

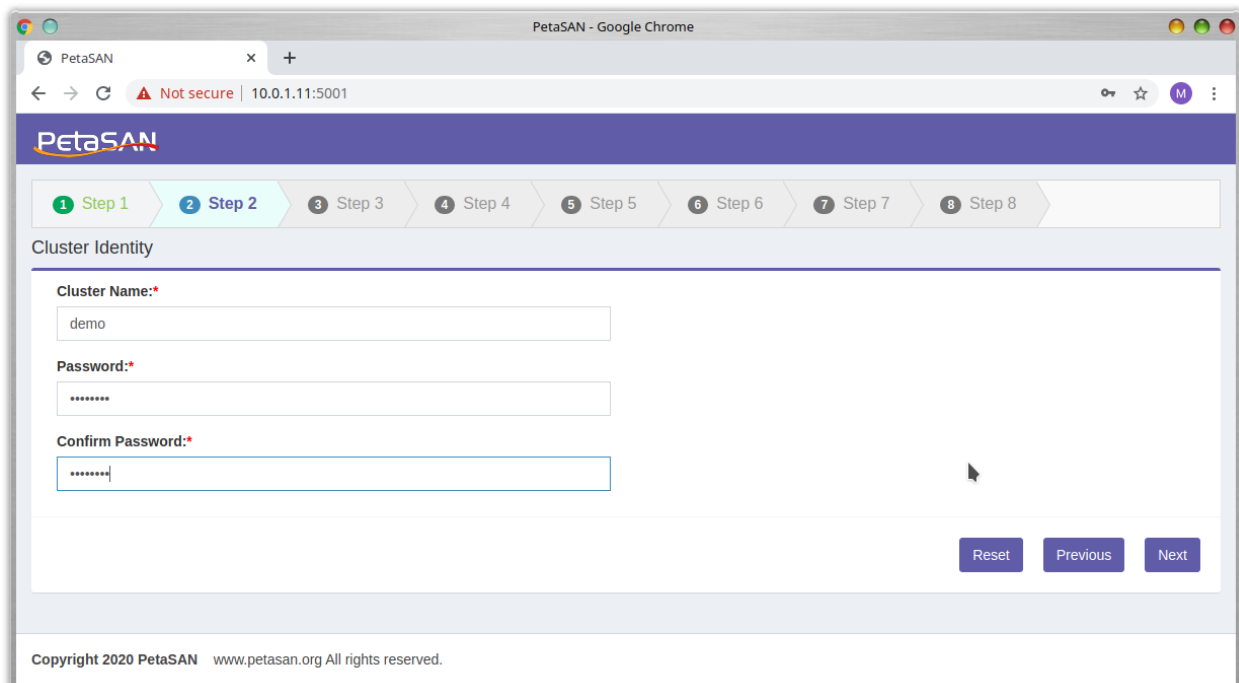
4. Node Deployment

Node 1

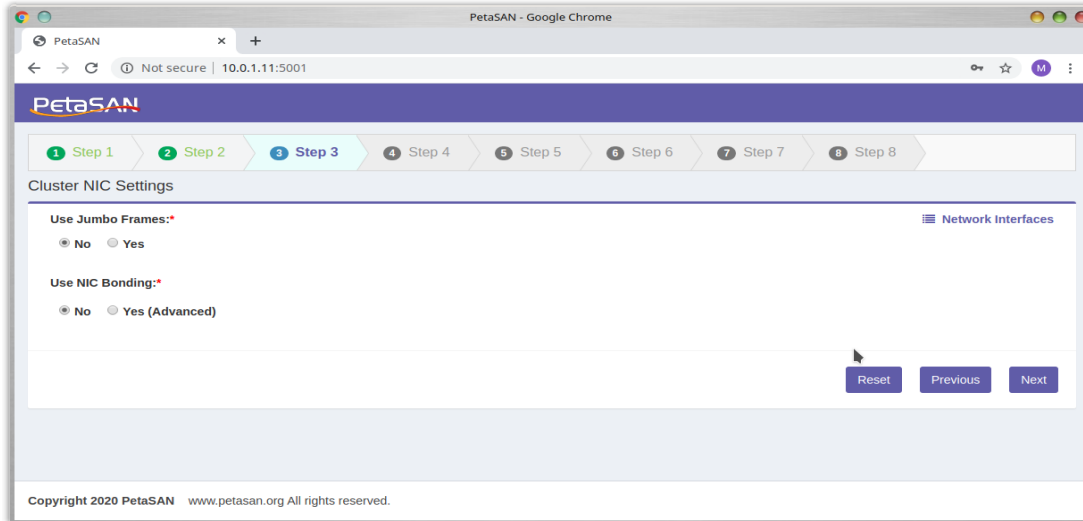
To deploy node 1, access the Deployment Wizard on port 5001 of node, in our case <http://10.0.1.11:5001> . Since this is the first node, we need to create a new cluster.



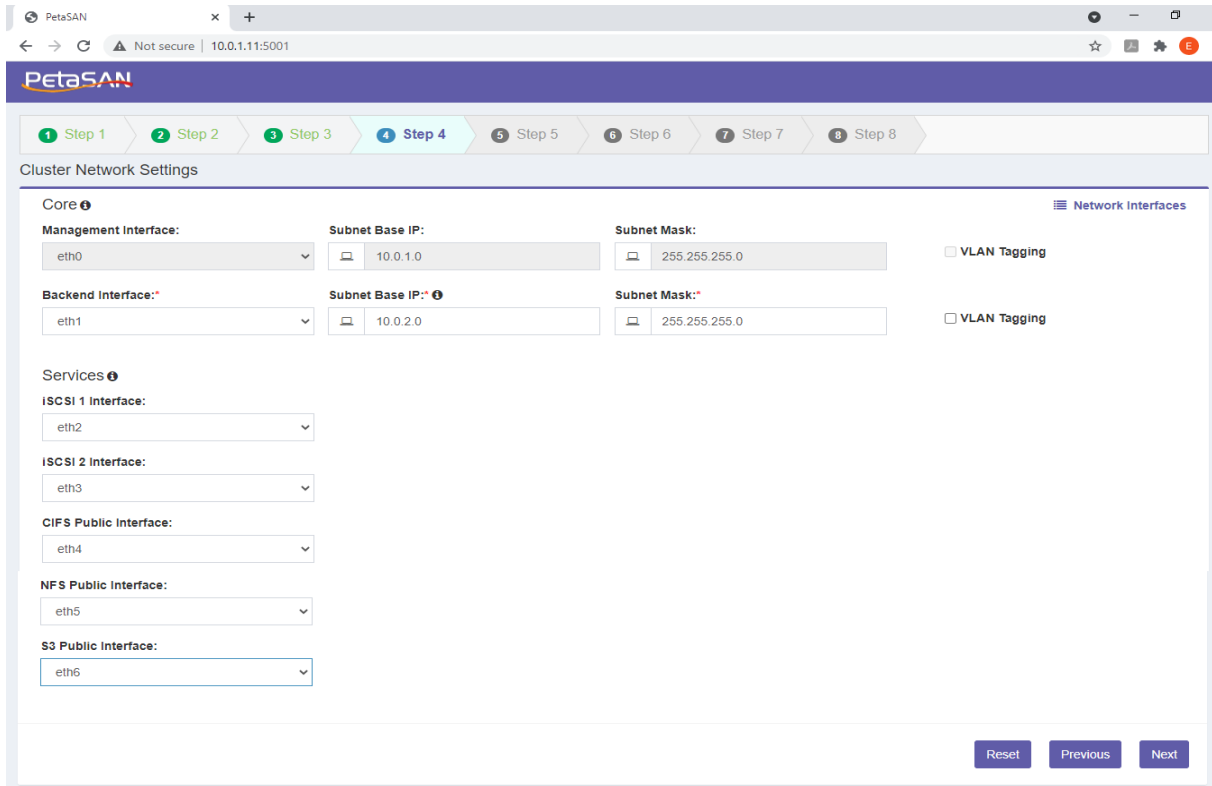
Input a name for the cluster and assign a password, this will be the root password for all cluster nodes. It is possible to ssh to all nodes using this password if desired.



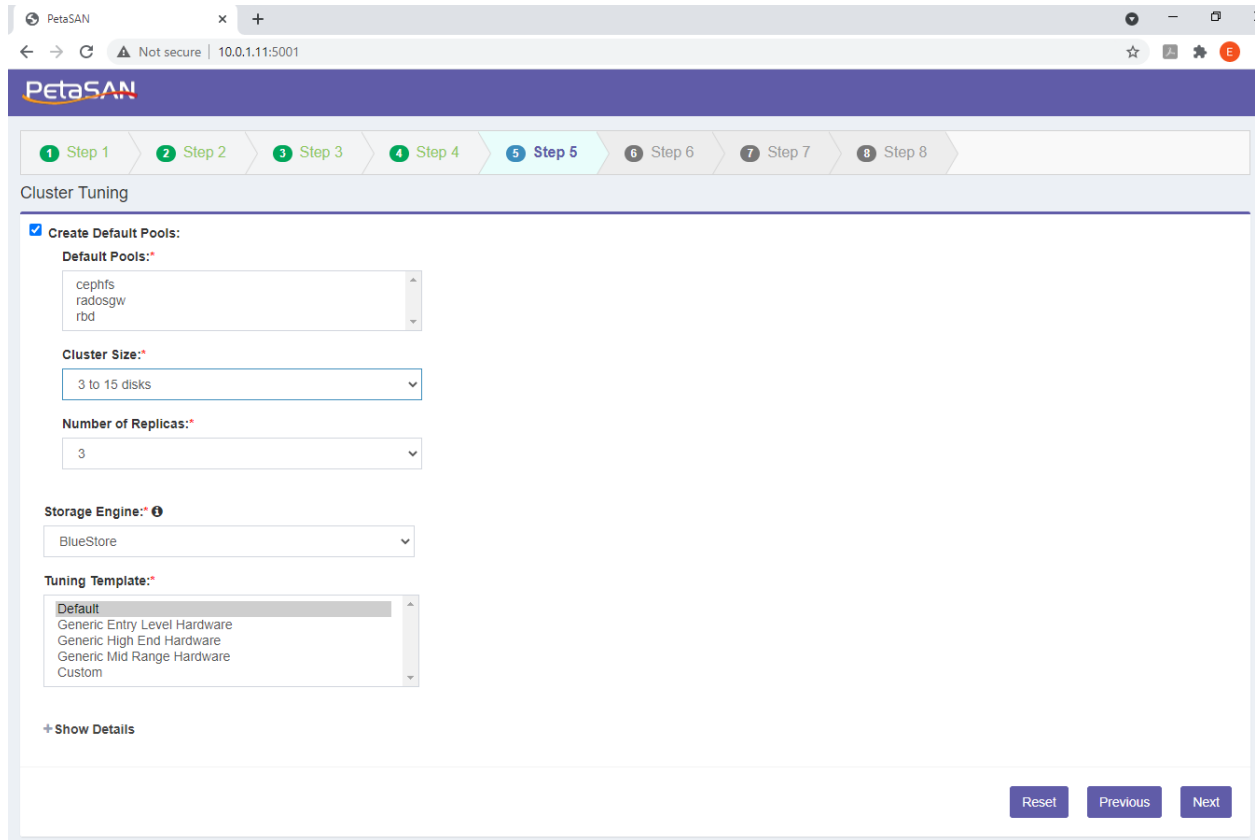
Next in the Cluster NIC Settings we need to define network interface settings such as jumbo frames and interface bonding. Note that these settings apply to the entire cluster and not specific to the node being deployed, the specific network configuration for each node will depend on the role/service that will be assigned to the node as will be shown later.



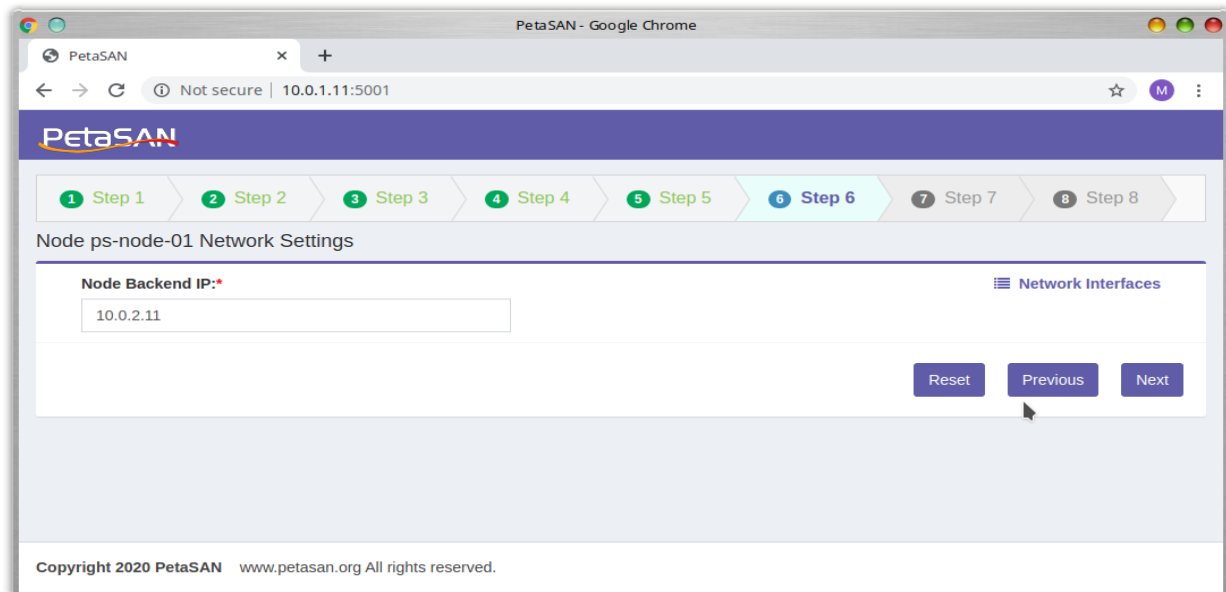
In our case we leave the defaults. Next in the Cluster Network Settings we need to map our subnets to the cluster interfaces and bonds. Be careful with the Backend network settings as they are not easy to change once the cluster is operational. The Services subnet interfaces can be changed later from the ui in an operational cluster.



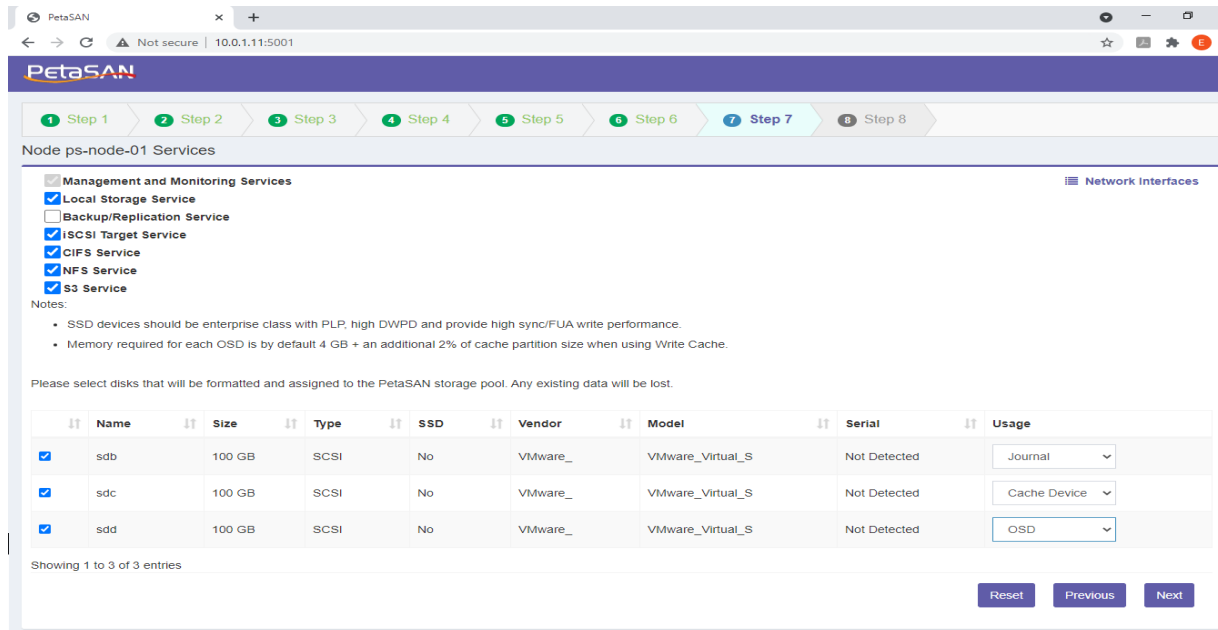
Next is the Cluster Tuning page, select the pools to create (cephfs, radosgw, rbd) according to your requirements, input the expected cluster disk size range. This applies to the storage pool that will be created by the Deployment Wizard, later in production you can create/delete pools as you wish.



Next input the Backend IP address of the current node being deployed, in our case 10.0.2.11



Next in the Node Services page, define the roles/services the current node will serve

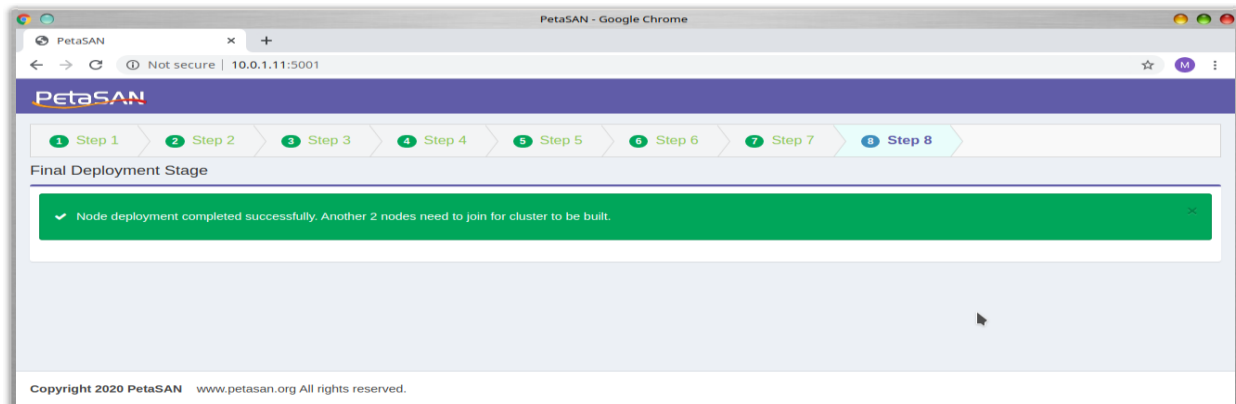


In our case we assign the Local Storage, iSCSI, CIFS, NFS and S3. If the node is assigned the Local Storage Service you can/should select the disks to act as storage OSDs/disks. In case you have a mix of HDD and SDDs, it is possible to setup the SSDs as journals and assign HDDs as OSDs (recommended ratio is 1:4), else choose them as OSDs. It is also possible to add/delete OSDs in a deployed running cluster. Note that node roles can be changed later in a deployed running cluster.

As indicated, the network configuration that will be applied to a node depends on its roles/services. It is sometimes useful to:

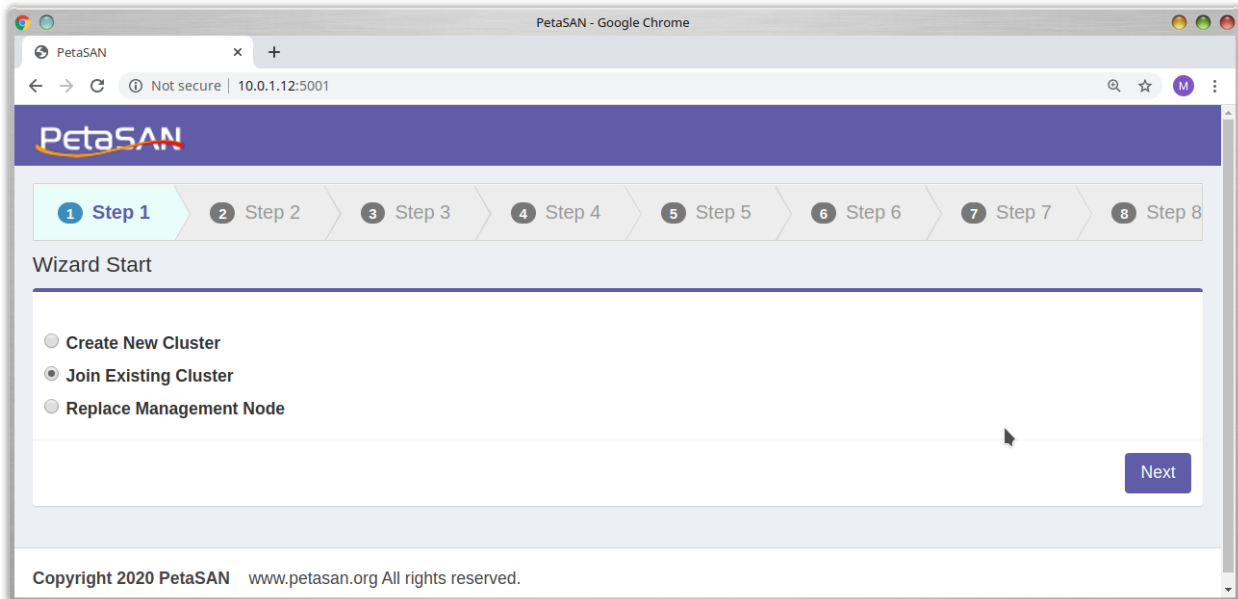
- View the current node hardware interfaces detail and compare them to the cluster network definitions for each service. This can be viewed from the Network Interfaces link located at the top right.
- In some cases, such as having different interface models across your nodes, it is useful to rename or re-arrange the interface order, this can be done from the blue node console menu.

Clicking next will show that node 1 deployment has been completed.

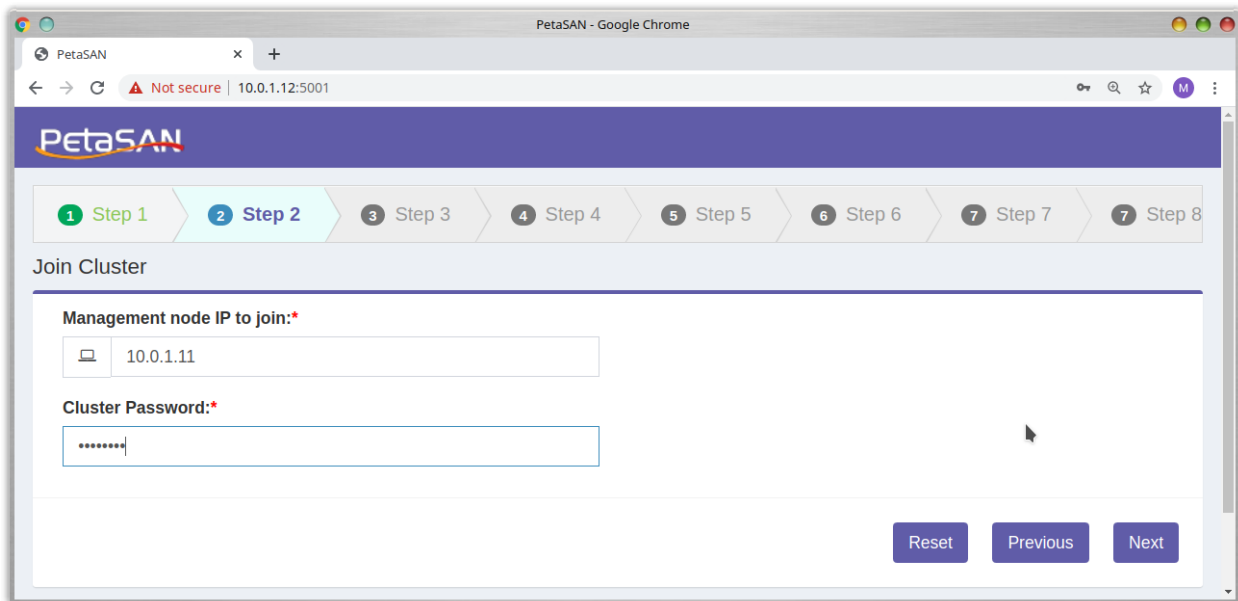


Node 2

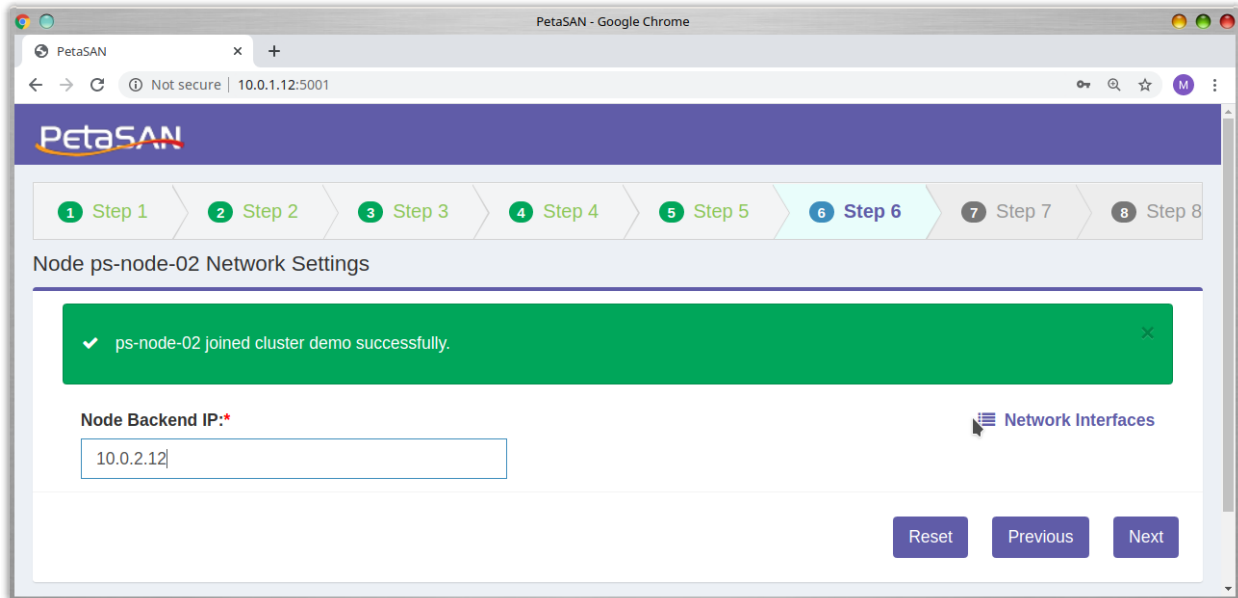
To deploy node 2, access the Deployment Wizard on port 5001 of node, In our case <http://10.0.1.12:5001>. Select to Join the cluster.



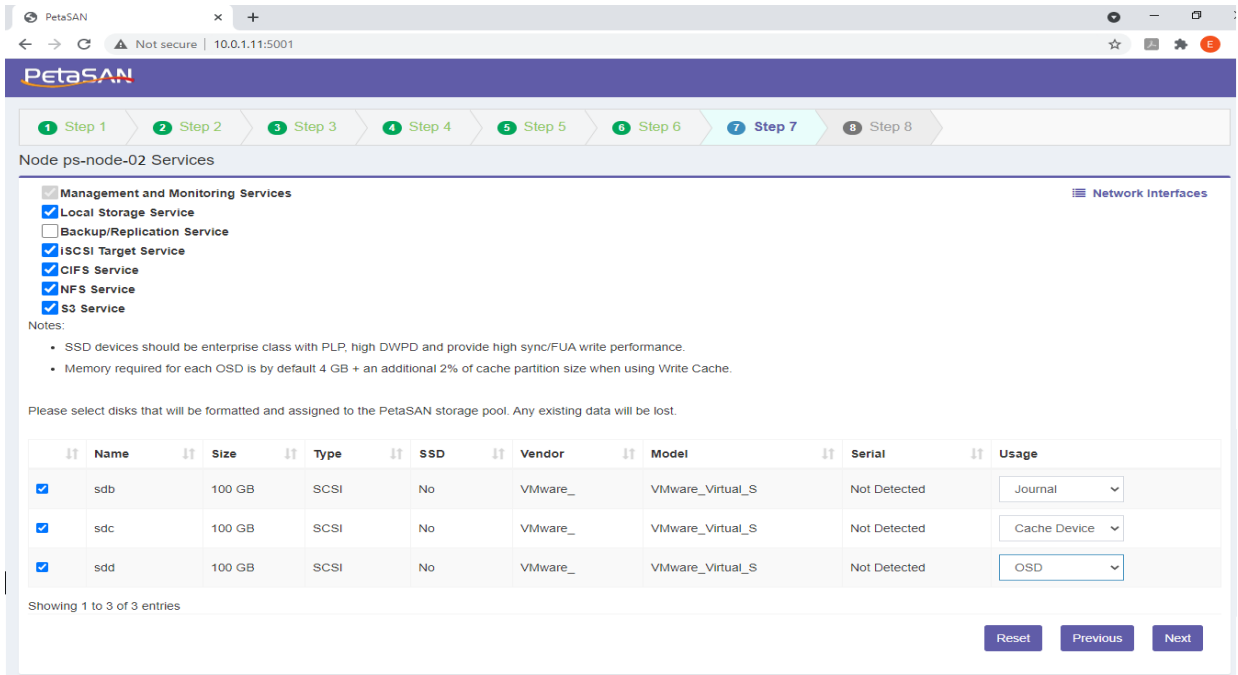
Type in the IP address of the first node we are joining, 10.0.1.11 and the password



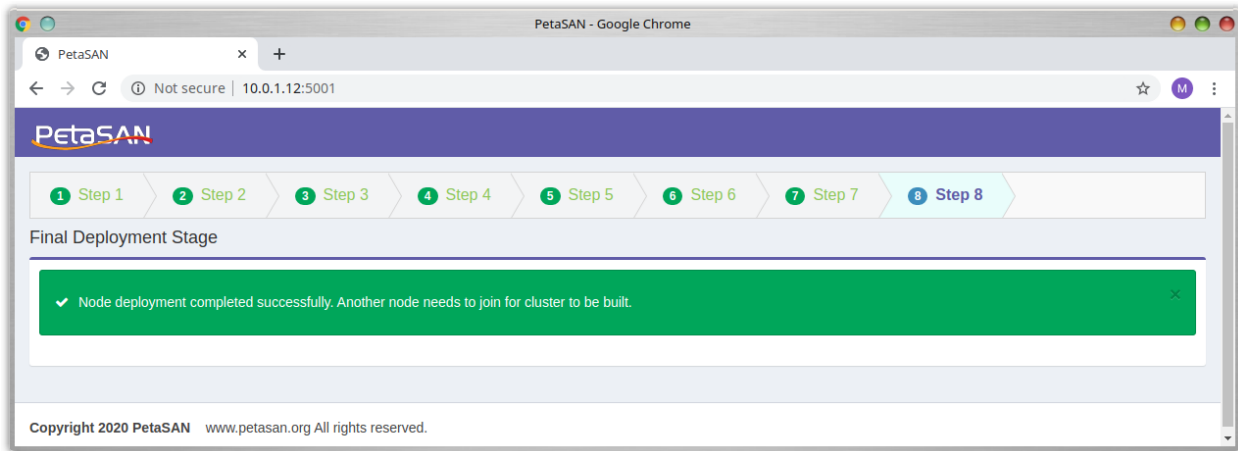
Enter IP for backend interface, in our case 10.0.2.12



Select Local Storage, iSCSI, CIFS, NFS and S3 as roles. Select desired OSDs as we did with node 1.

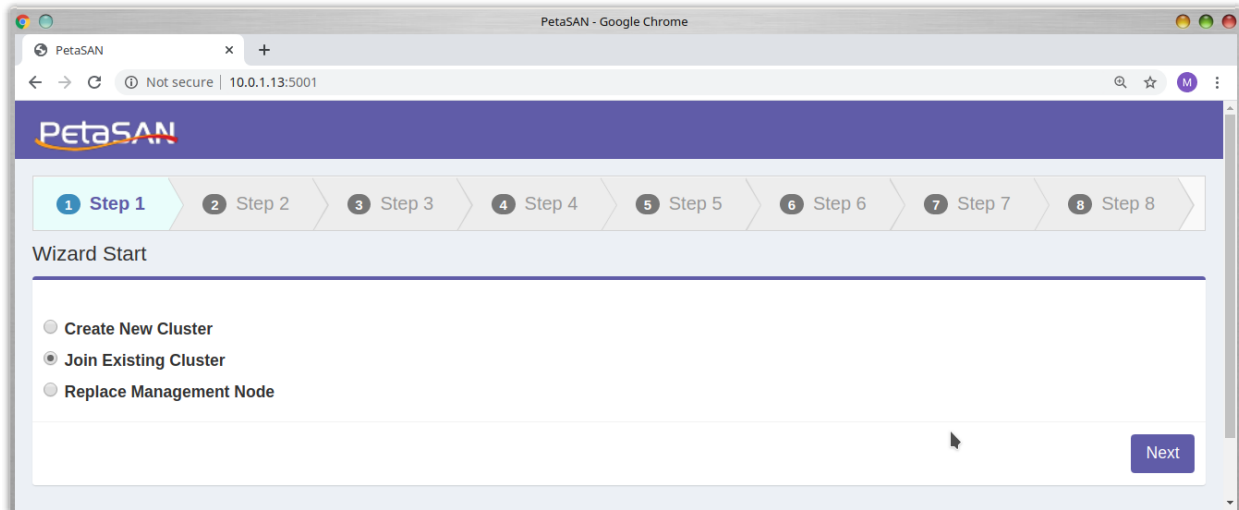


Clicking next will show that node 2 deployment has been completed.

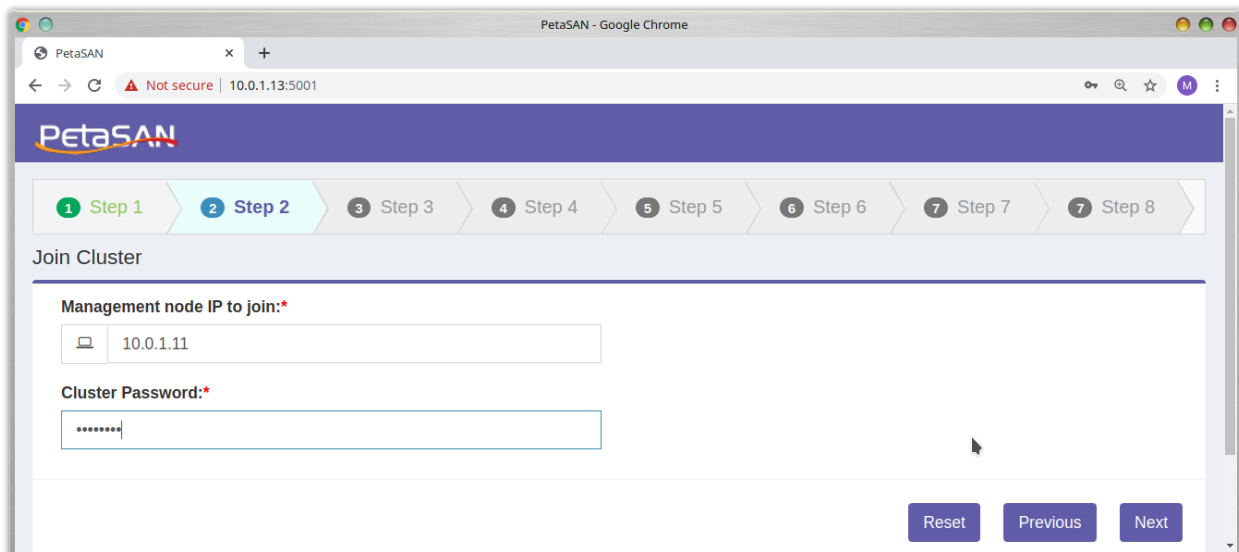


Node 3

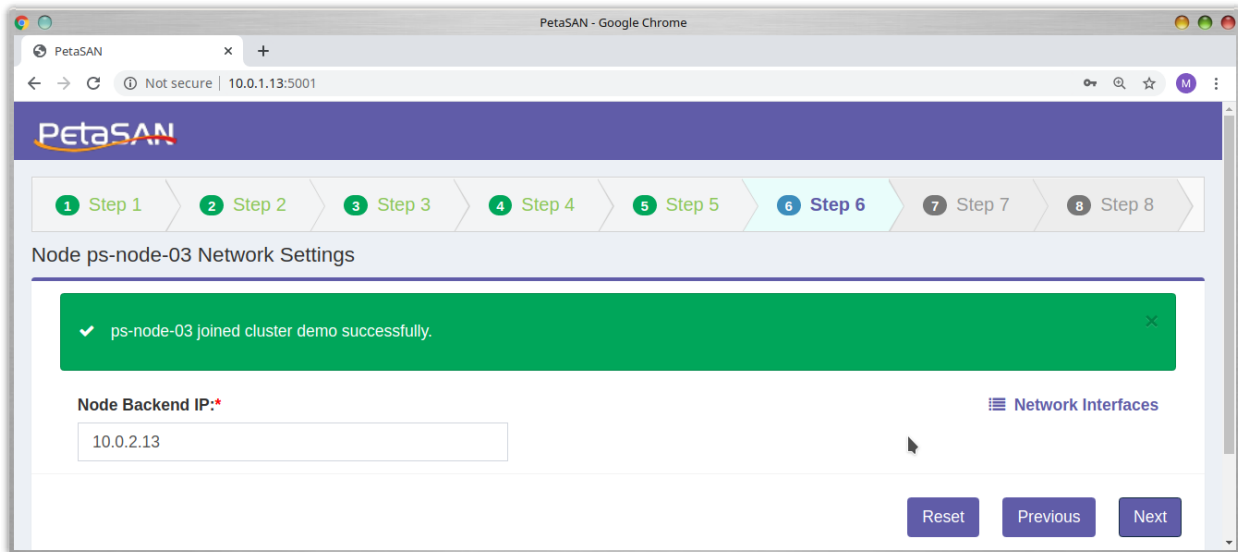
To deploy node 3, access the Deployment Wizard on port 5001 of node, In our case <http://10.0.1.13:5001>. Select to Join the cluster.



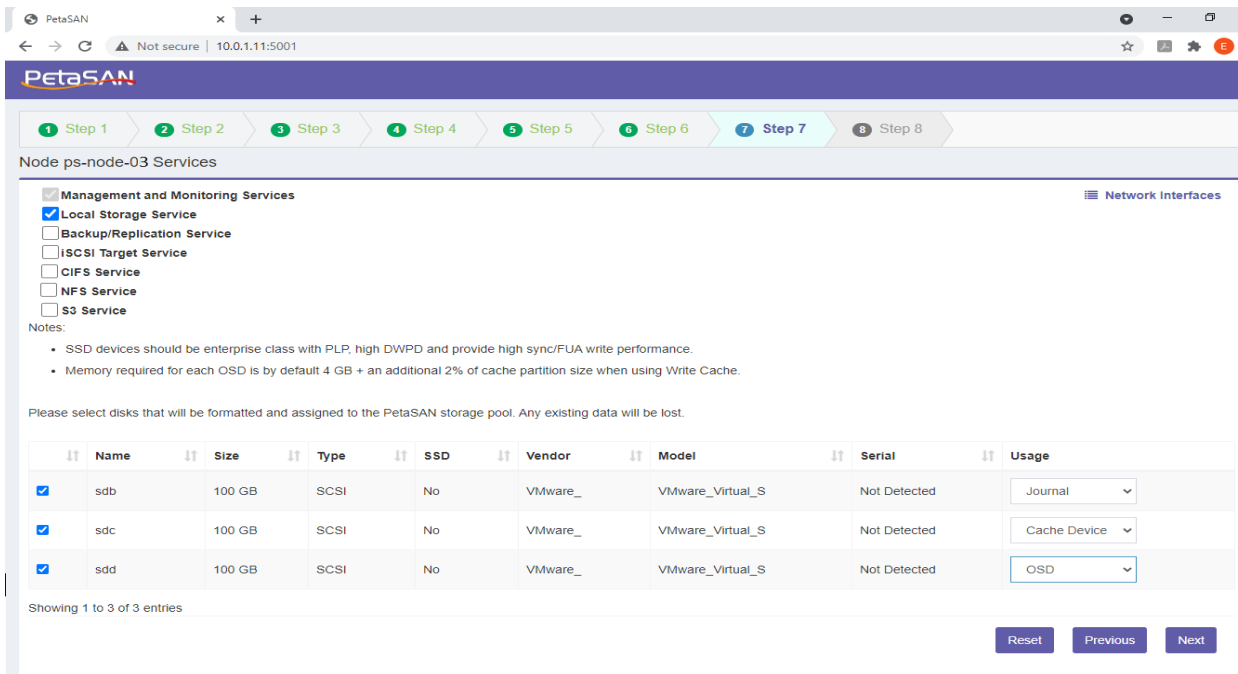
Type in the IP address of the first node we are joining, 10.0.1.11 and the password



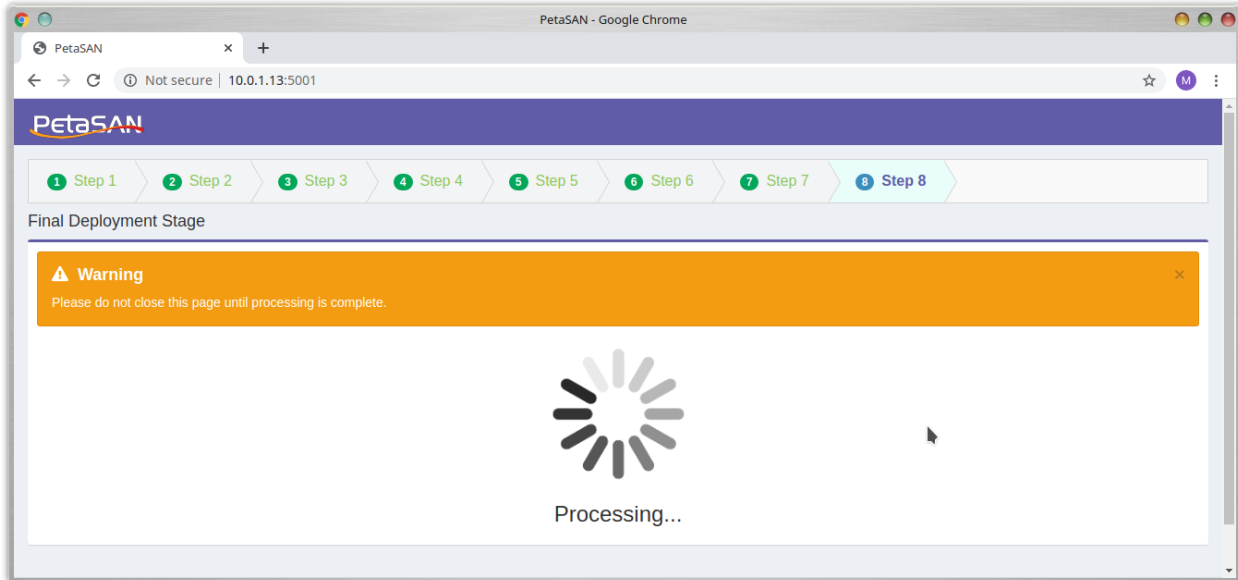
Input our backend IP, 10.0.2.13



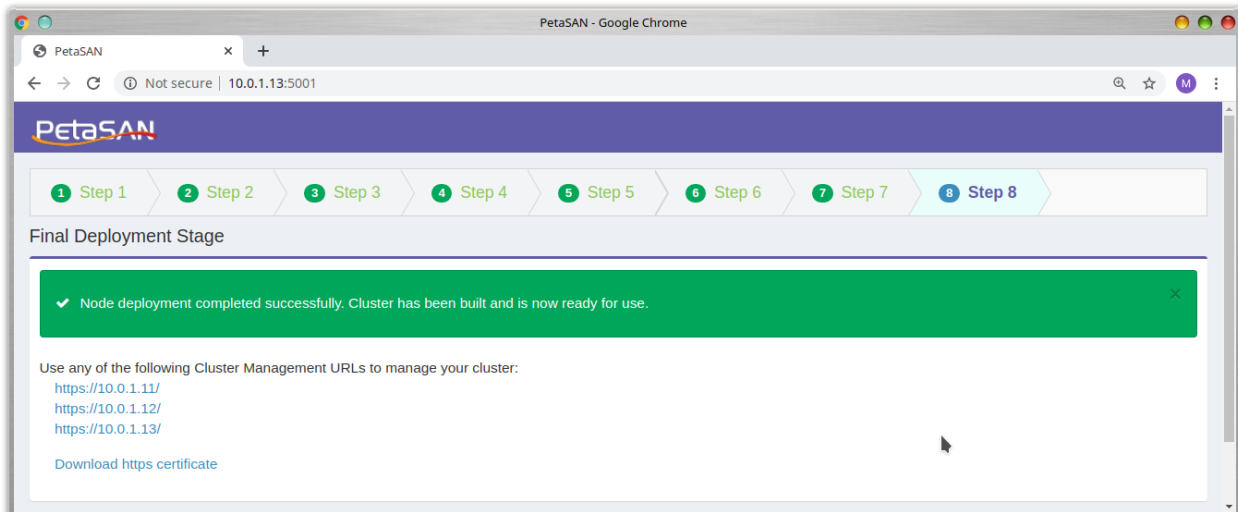
Select Local Storage role and any storage disks/OSDs.



At the end of 3rd node deployment, the cluster will be built. This may take time depending on the number of drives you have, leave it a good 30 min.



On successful completion we get a success message.



Please download and install the https certificate from the link presented. Install in your browser under trusted root certificate authorities.

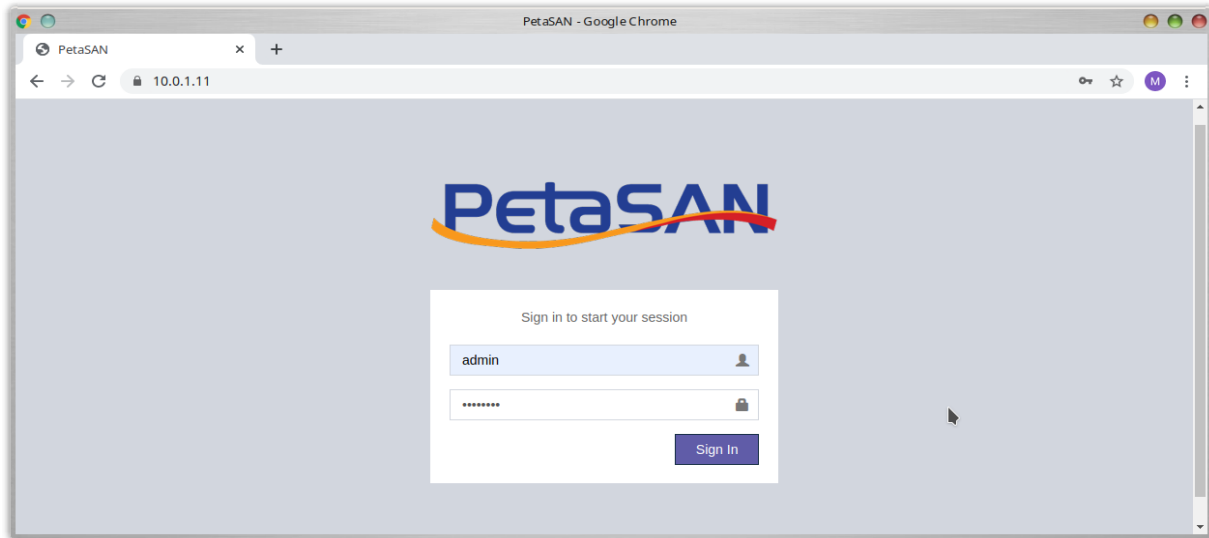
Nodes 4+

All nodes starting from 4 and above join an already built cluster, the join steps are the same as before.

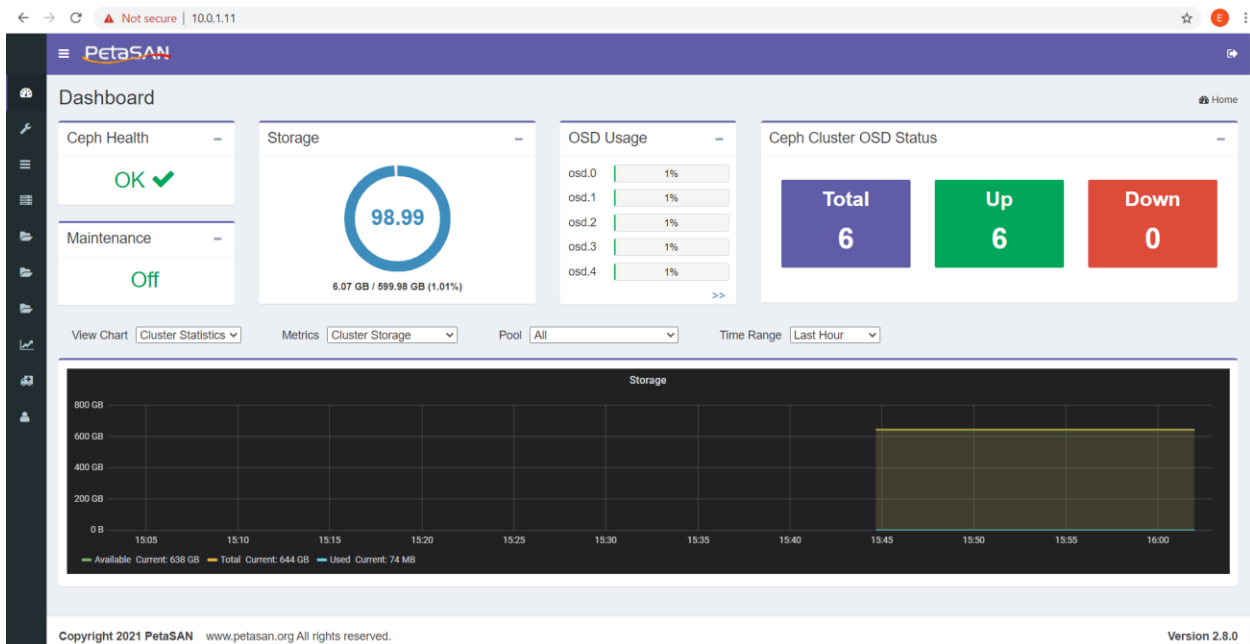
5. iSCSI Setup

Access the Management ui by accessing https on any of the first 3 nodes, for example <https://10.0.1.11>

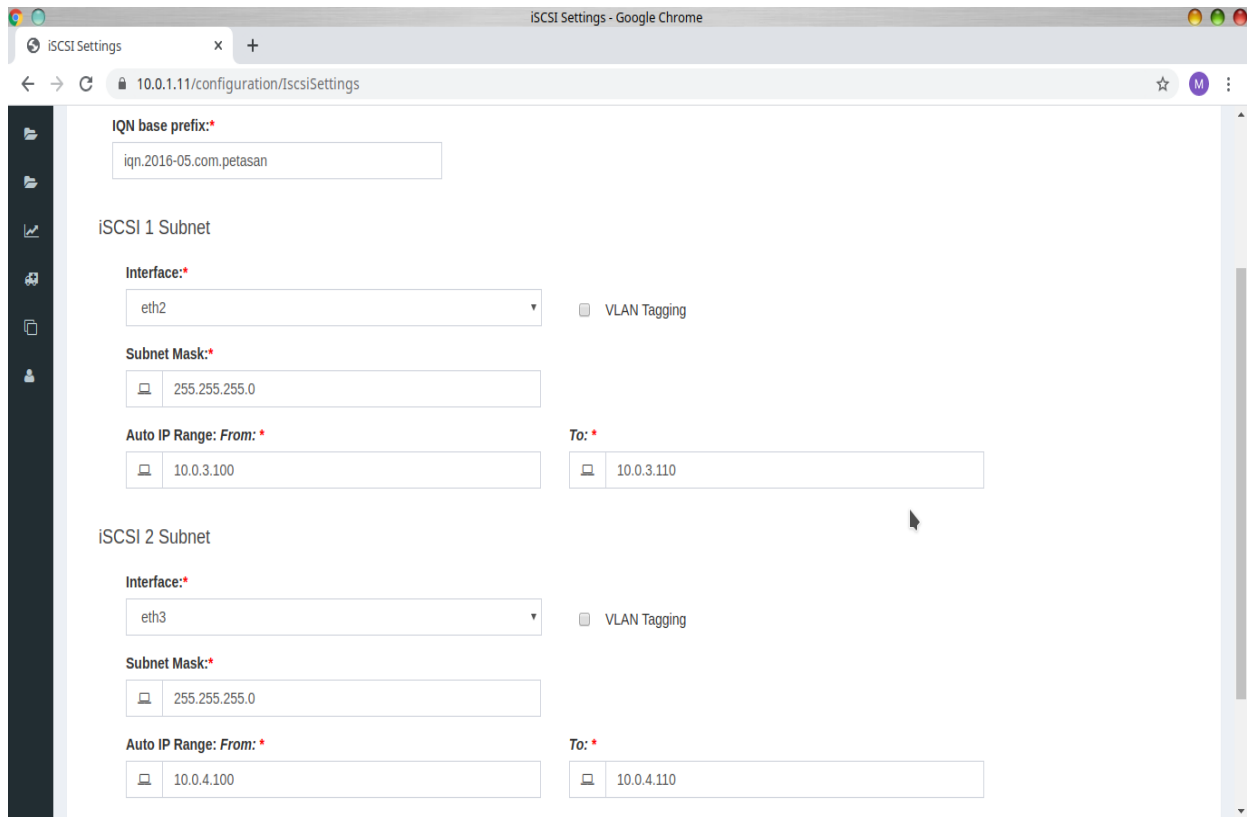
Log in as user “admin” with initial password “password”. If you had not downloaded and installed the https certificate earlier, you can still login in an un-secure session then download the certificate from the Configuration menu.



This will log you to the Management web interface.



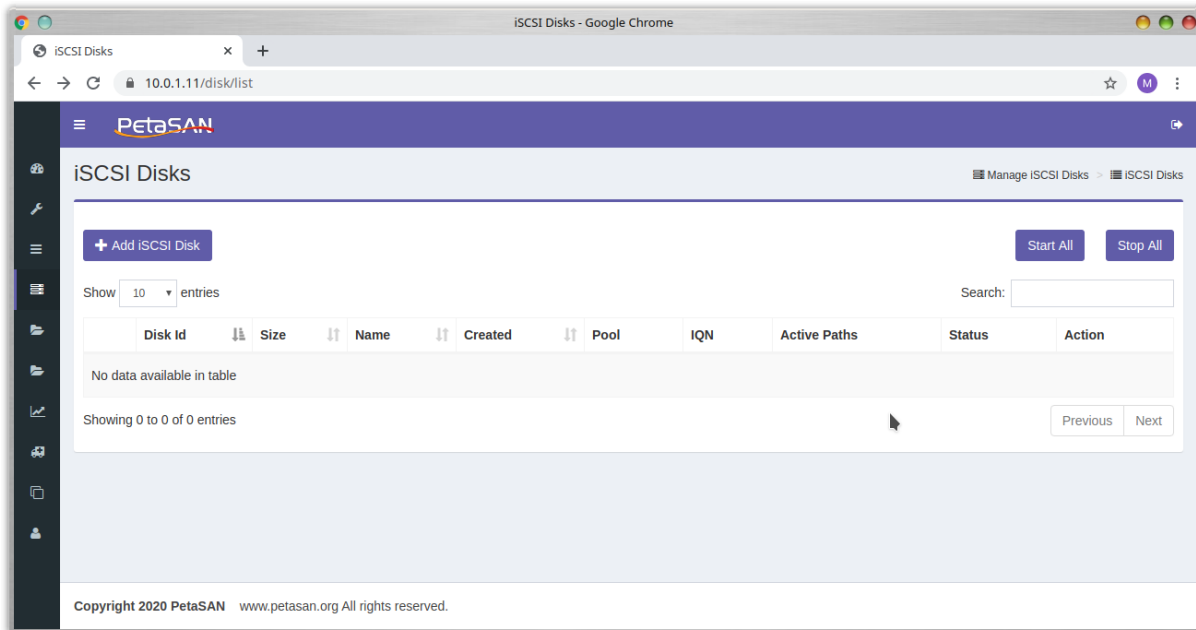
From the Configuration menu, choose iSCSI Settings



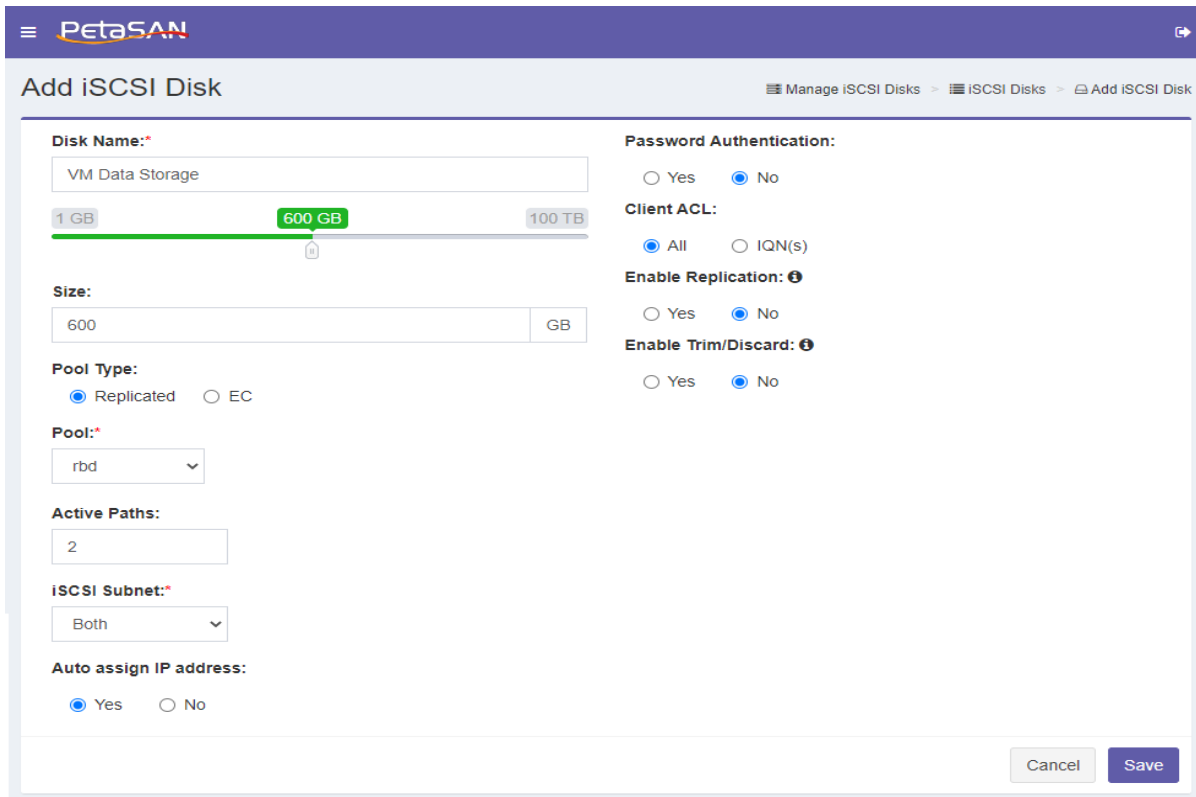
Define iSCSI 1 and 2 subnets, in our case 10.0.3.0/255.255.255.0 and 10.0.4.0/255.255.255.0.

Also define an IP from and To range for automatic IP assignment for iSCSI disks, in our case we define 10.0.3.100 to 10.0.3.110 for iSCSI 1 and 10.0.4.100 to 10.0.4.110 for iSCSI 2.

From the Manage iSCSI Disks menu select iSCSI Disks



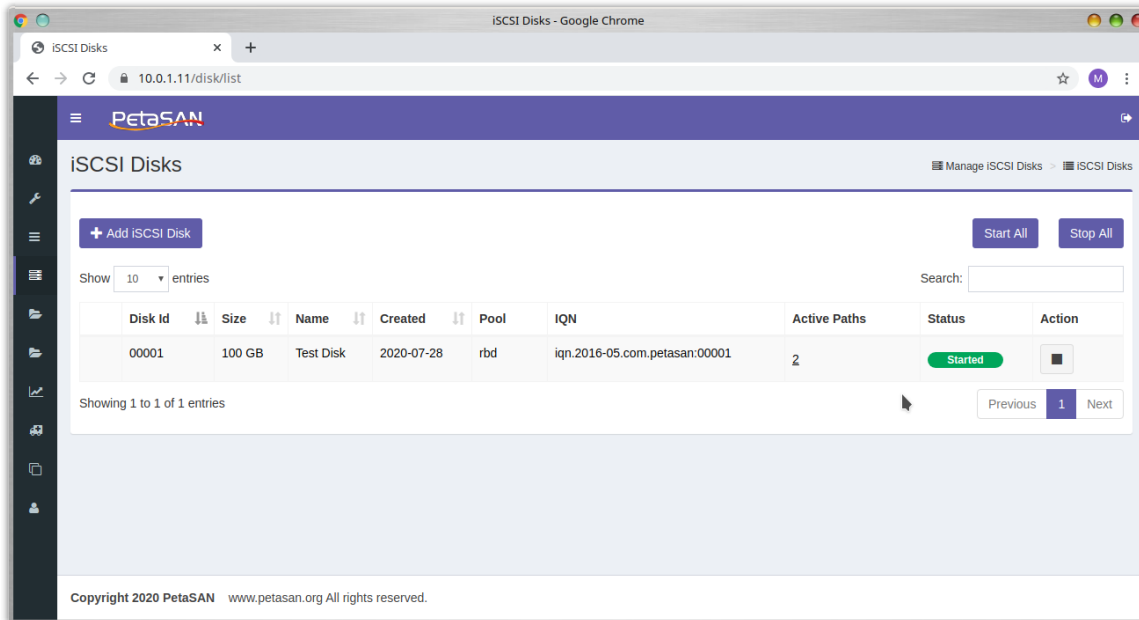
Click on the Add iSCSI Disk button



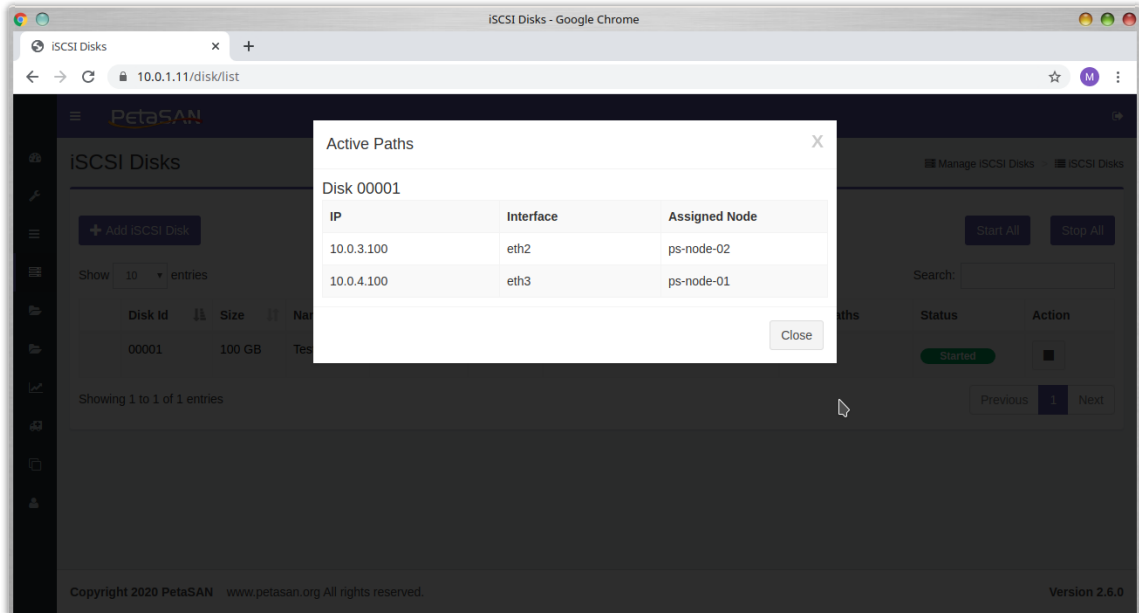
Enter a disk name and specify the desired size. We chose the default of auto-ip assignment but we could also manually enter specific IPs. It is also possible to increase the number of active paths from default

value of 2, this is useful if you have a large number of iSCSI server nodes and would like the disk to actively load balance among the different servers for better performance.

On successful creation you should see the disk as started.



Click on the Active Paths link

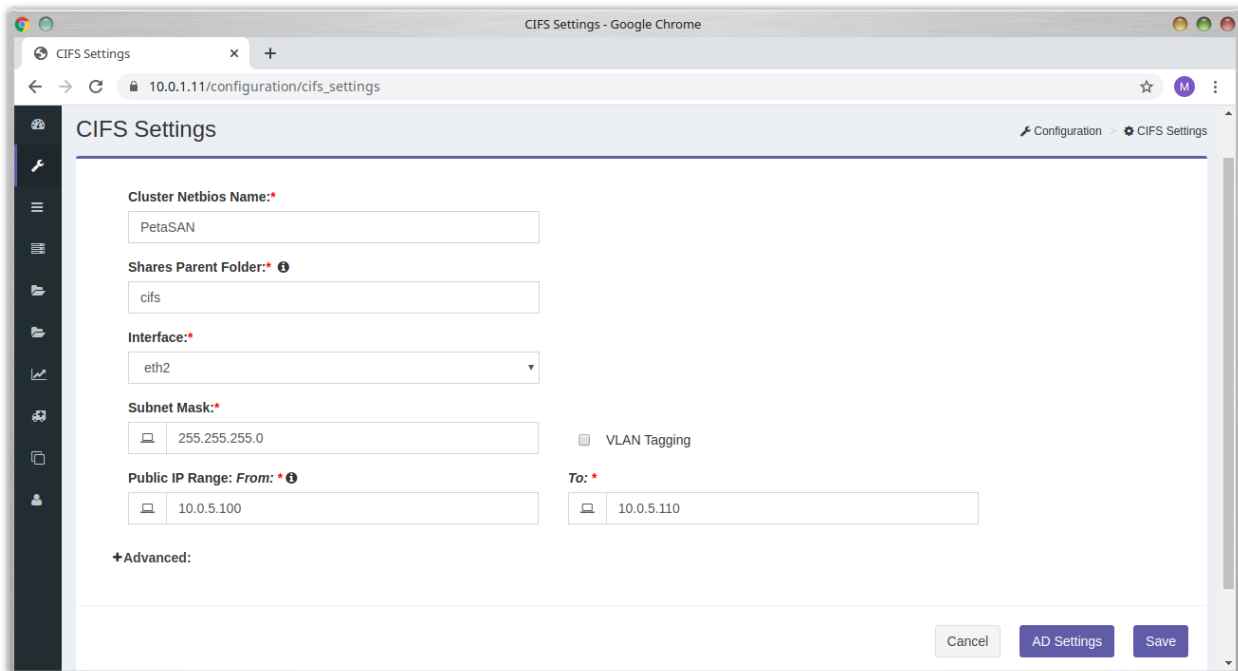


Since we chose earlier to have auto-ip assignment, please take note of the assigned IPs for the new disk.

Connecting iSCSI client depends on the client OS, please refer to specific PetaSAN guides for setting up Windows and VMWare iSCSI connections in an active/active setup to a PetaSAN cluster.

6. CIFS/SMB Setup

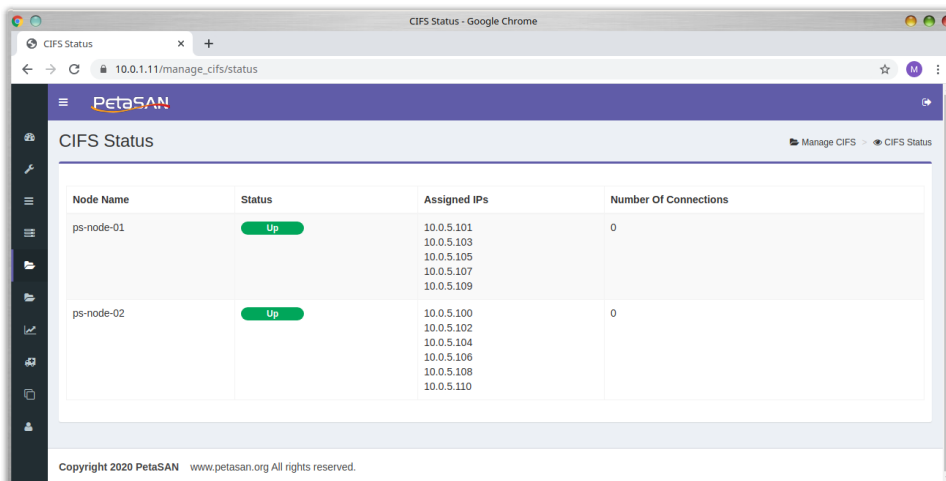
From the Configuration menu, choose CIFS Settings



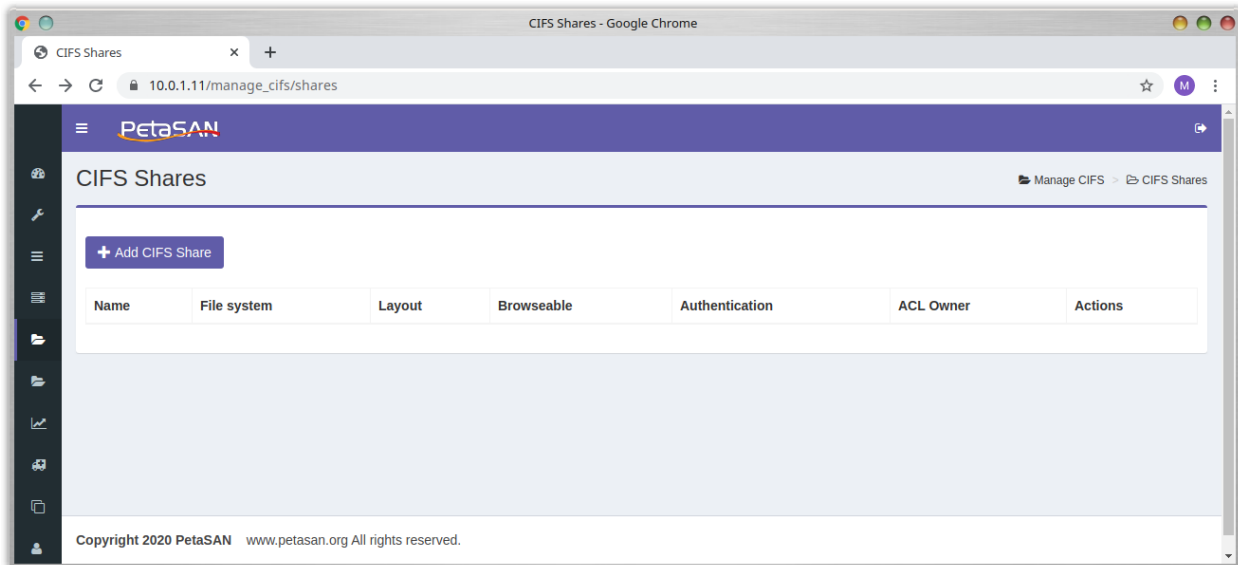
Select the public subnet for CIFS connections and range of IPs to assign to PetaSAN CIFS servers. In our case we assign a range from 10.0.5.100 to 10.0.5.110. It is also possible to join an Active Directory domain to provide secure shares from this page.

Note it is recommended to assign more than one IP per CIFS server, this is useful in a multi node CIFS environment, if a server fails, its ips will be distributed to several other active servers rather than failing all its traffic to another single server.

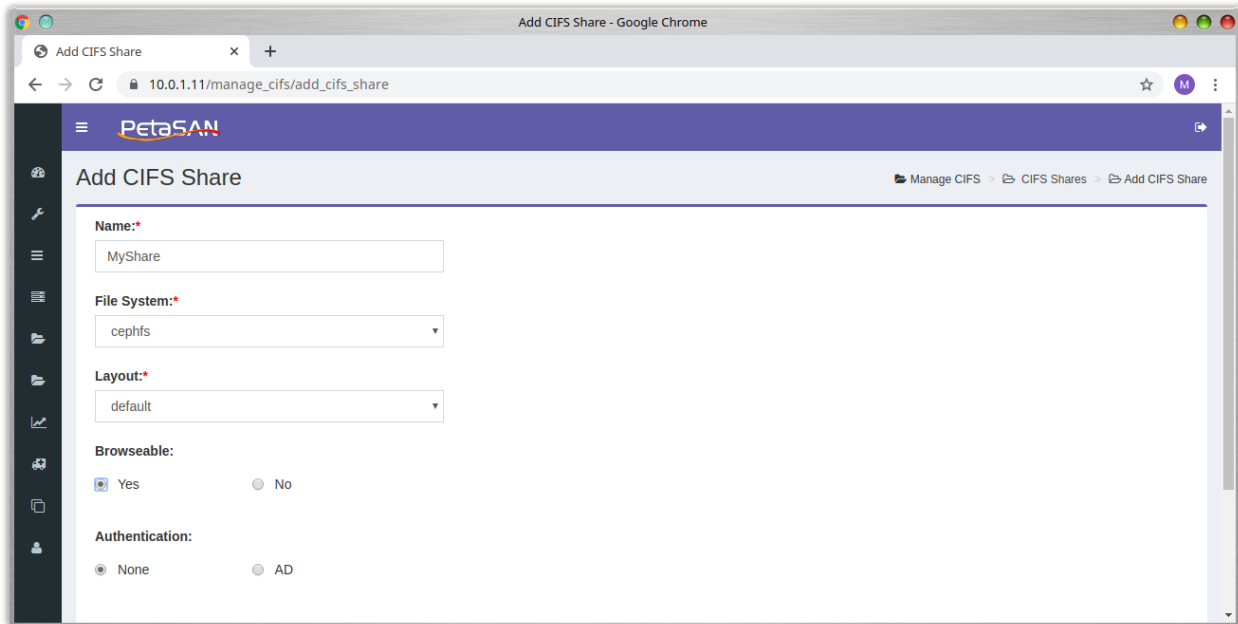
From the Manage CIFS, CIFS Status page you can view the active CIFS servers and their IPs.



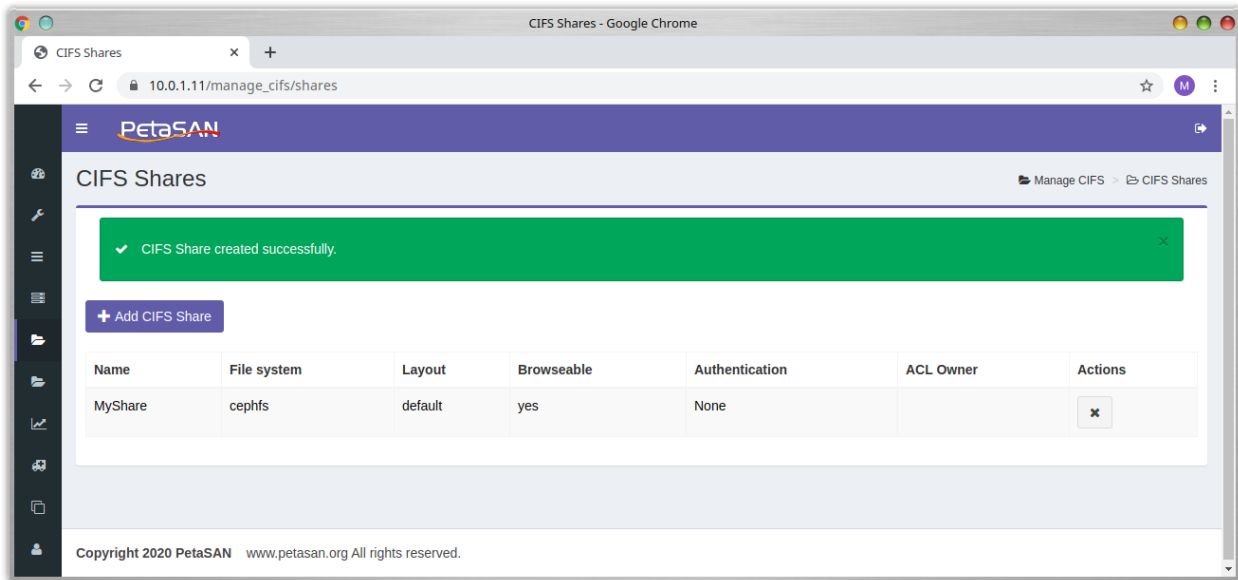
From the CIFS Shares page, click Add CIFS Share



Specify the share name to add



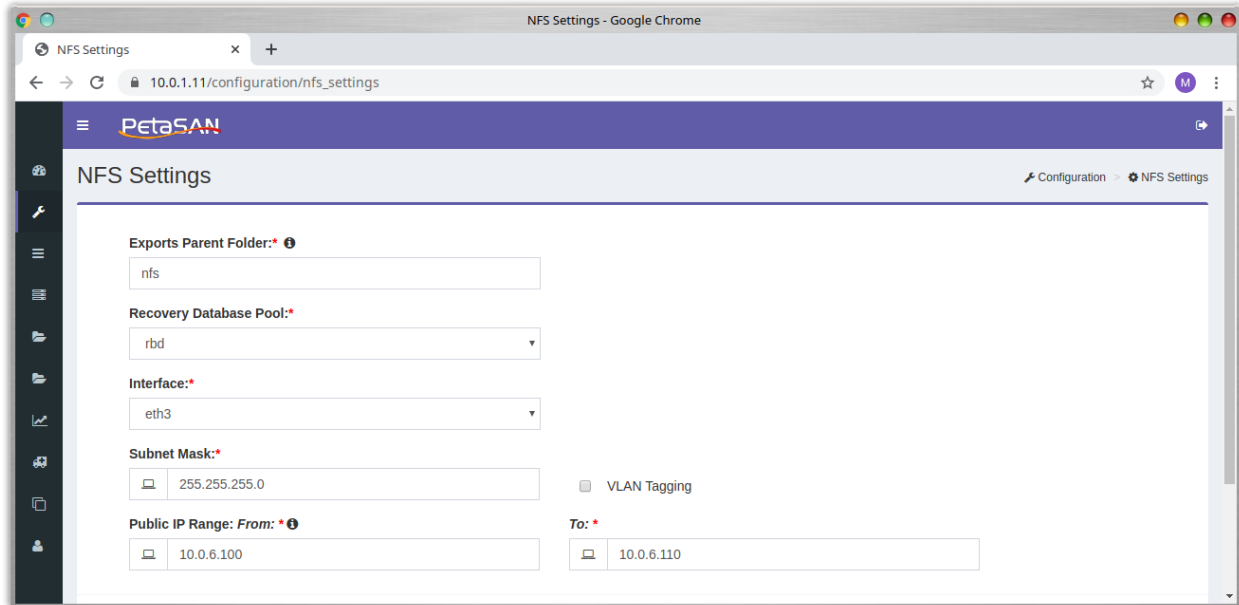
The new share is now created.



Note: CIFS/SMB clients access the share via [\\IP_Address\share](#) such as [\\10.0.5.100\MyShare](#). For load balancing it is possible to distribute the range of IPs across the clients in a manual way or better setup a round-robin DNS to serve these IPs and have all clients access the cluster using a Netbios name such as [\\PetaSAN\MyShare](#).

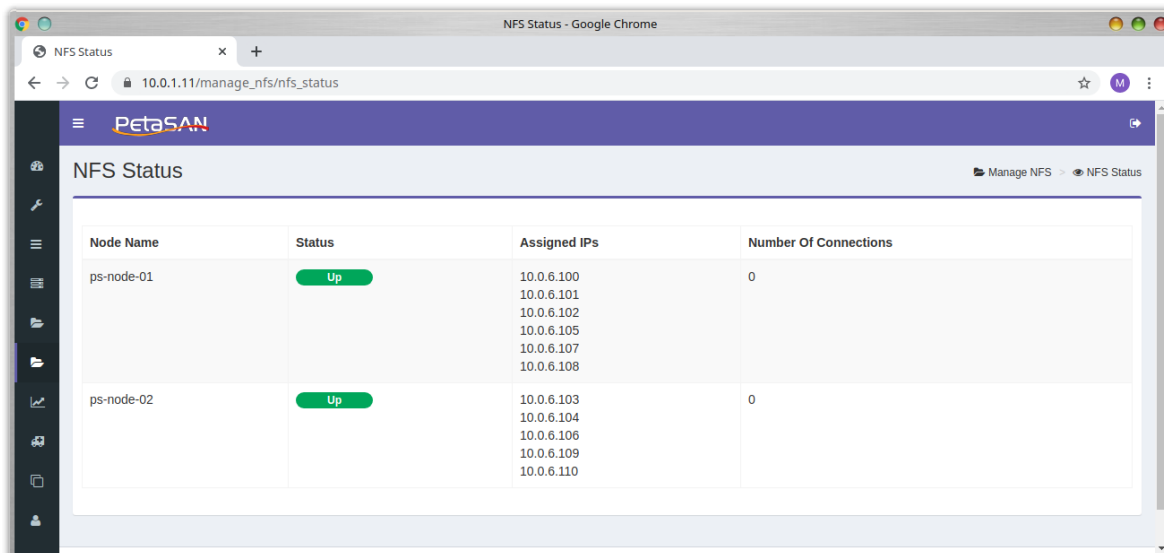
7. NFS Setup

From the Configuration menu, choose NFS Settings

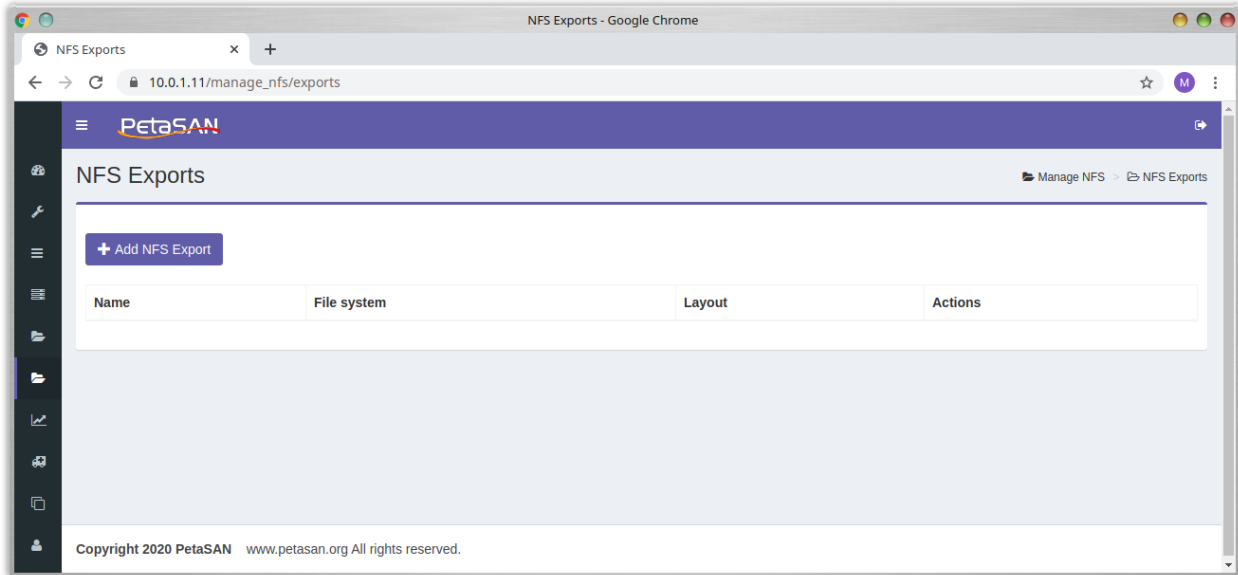


Select the public subnet for NFS connections and range of IPs to assign to PetaSAN NFS servers. In our case we assign a range from 10.0.6.100 to 10.0.6.110. Note it is recommended to assign more than one IP per NFS server, this is useful in a multi node NFS environment, if a server fails, its ips will be distributed to several other active servers rather than failing all its traffic to another single server. We also need to select a recovery database pool used by the NFS servers, select the “rbd” pool.

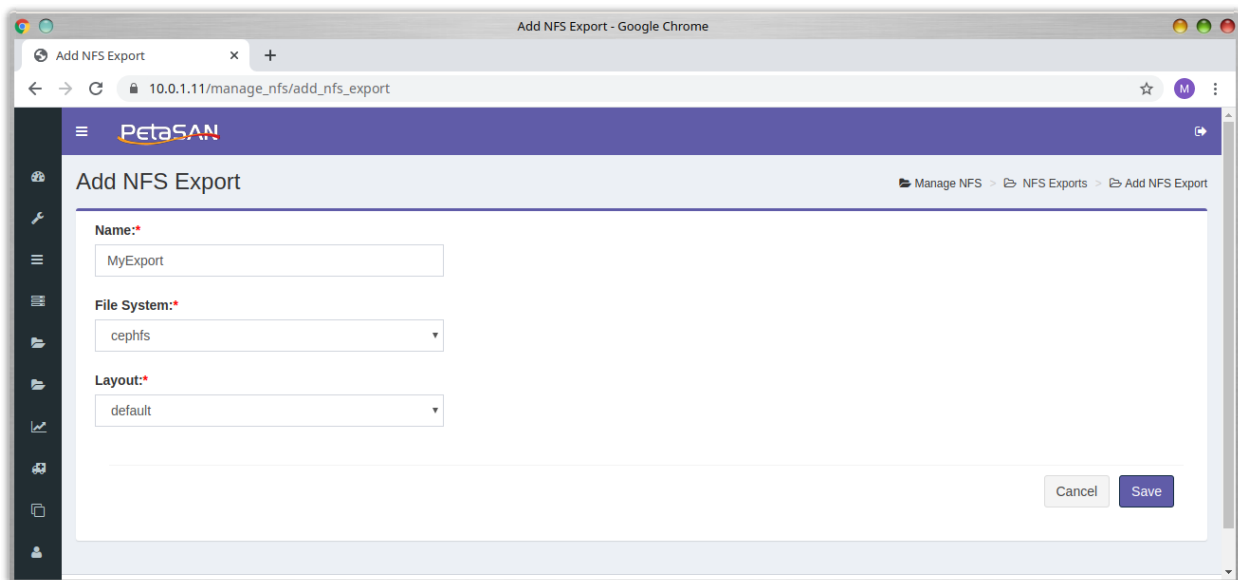
From the Manage NFS, NFS Status page you can view the active NFS servers and their IPs.



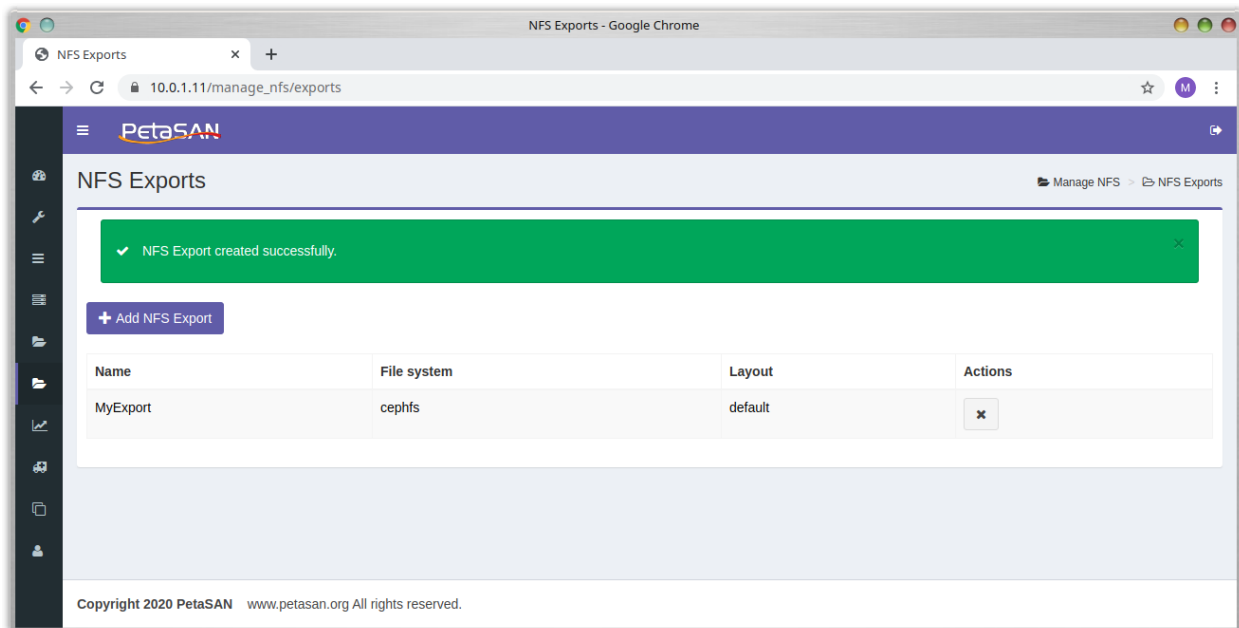
From the NFS Exports page, click Add NFS Export



Specify the export name to add



The new export is now created.



For load balancing it is possible to distribute the range of IPs across the clients in a manual way or better setup a round-robin DNS to serve these IPs and have all clients access the NFS cluster by name.

8. S3 Setup

From the Configuration menu, choose S3 Settings

RadosGW http port:
7480

Load Balancer

Port:
8000

HTTPS

Private Key:

```
-----BEGIN PRIVATE KEY-----
MIIEvgIBADANBgkqhkiG9w0BAQEFAASCCSwggkoAgEAAoIQAQCzj6hESwdtyYH
I2ZAv+wyzmm7KAAbpFZZi4MO1piXq+OkXmVhgUHlqJMeaScLQ3jFdfW91N2pqSo0
ZGLSPJZG1Rr1WU9BGSo20lq4+IG2HjUCdN+Jp1QOar2rT7QX80hMkHST12H1ZA
CETix1dlbN5wRkDyhP7sXmjVN0Ll8akRROIQXPjIUPMYZw0zJcWt6/ZIW7RXLGBx
0EJ2G2/MoHI4HLFLUuwxEs3hTcLQGAt0kNodVY8SkaWdGn/Z05YQdq6iUr9mRXmc
-----
```

Generate

Certificate:

```
-----BEGIN CERTIFICATE-----
MIIEvzCCAqegAwIBAgIUeomjwARhfJy3IZFfbMJWh0y7YHcwDQYJKoZIhvcNAQEL
BQAwDzENMAsGA1UEAwEZW1hbjAeFw0yMTA1MjMyMTI4MTJaFw00MTA1MTcyMTI4
MTJJaA8xDALBgNVBAMMBGVtYW4wgglMA0GCSqGSIb3DQEBAQUAA4ICDwAwggIk
AoICAQCzj6hESwdtyYHI2ZAv+wyzmm7KAAbpFZZi4MO1piXq+OkXmVhgUHlqJMe
aScLQ3jFdfW91N2pqSo0ZGLSPJZG1Rr1WU9BGSo20lq4+IG2HjUCdN+Jp1QOar2
-----
```

Interface:
eth6

Subnet Mask:
255.255.255.0 VLAN Tagging

Public IP Range: From: 10.0.7.100 **To:** 10.0.7.110

Gateway:
 Default Custom

Buttons: Cancel, Download Certificate, Save

Enter the Load Balancer port number and in case you want to use https then click on Generate button to create SSL certificate.

Select the public subnet for S3 connections and range of public IPs to assign to PetaSAN S3 servers. In our case we assign a range from 10.0.7.100 to 10.0.7.110. Note it is recommended to assign more than one IP per S3 server, this is useful in a multi node S3 environment, if a server fails, its ips will be distributed to several other active servers rather than failing all its traffic to another single server.

You need to download and install the certificate on your clients.

From the Manage S3, S3 Load Balancer Status page you can view the active S3 servers and their IPs.

S3 Load Balancers Status

Show 10 entries Search:

Node Name	Status	Assigned IPs
ps-node-01	Up	10.0.7.100 10.0.7.101 10.0.7.102 10.0.7.103 10.0.7.105 10.0.7.106
ps-node-03	Up	10.0.7.104 10.0.7.107 10.0.7.108 10.0.7.109 10.0.7.110

Showing 1 to 2 of 2 entries

Previous 1 Next

PetaSAN S3 supports one realm and one zonegroup with multiple zones (one local zone per site)

Here we setup a single site

From the Configuration menu, choose S3 Zoneroups to add the zonegroup

Add Zonegroup

Name: default

.rgw.root Pool

No of PGs: 16

Placement Rule: replicated_rule

Size: 3

Placement Targets

- default-placement

Storage Class Name

- STANDARD

Cancel Save

Create the local zone under the created zonegroup

From the Configuration menu, choose S3 Zones to add the zone

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="default.rgw.control"/>	<input type="text" value="16"/>	<input type="text" value="replicated_rule"/>	<input type="text" value="3"/>
Meta	<input type="text" value="default.rgw.meta"/>	<input type="text" value="16"/>	<input type="text" value="replicated_rule"/>	<input type="text" value="3"/>
Log	<input type="text" value="default.rgw.log"/>	<input type="text" value="16"/>	<input type="text" value="replicated_rule"/>	<input type="text" value="3"/>

Default Placement +

Buckets Index Pool:

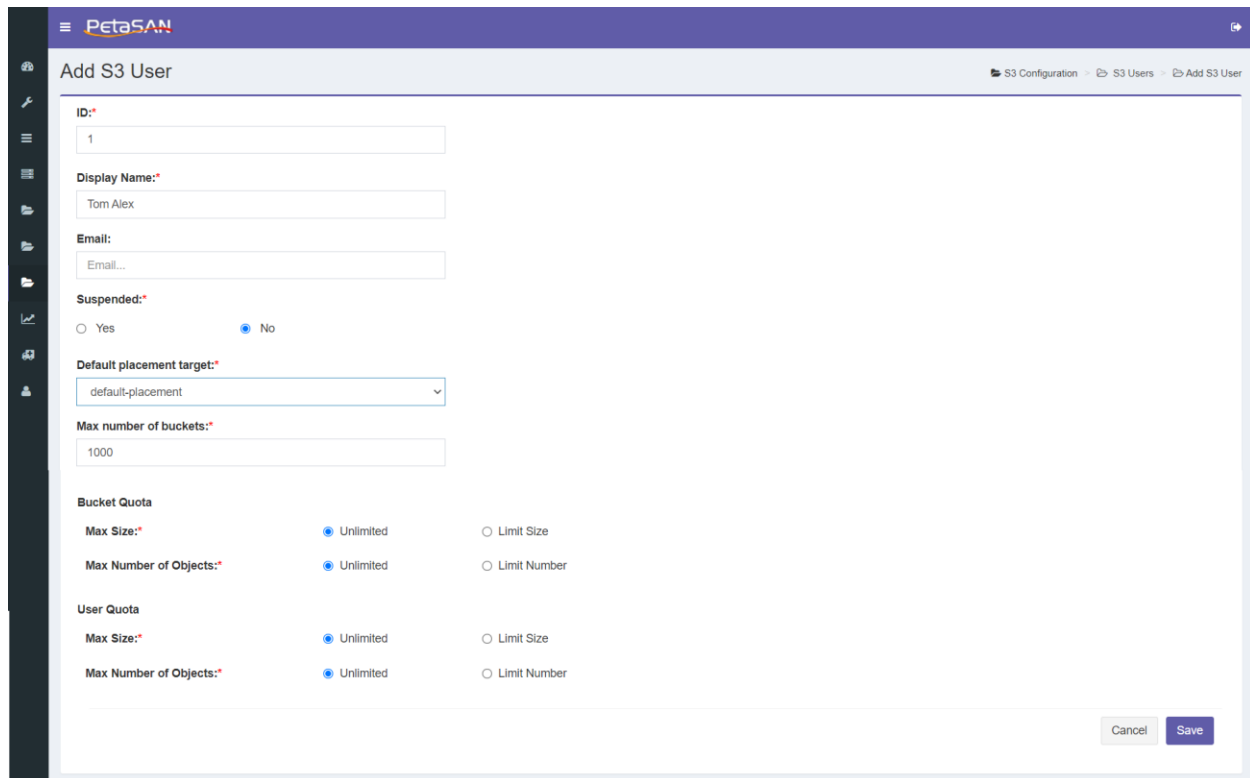
Function	Pool Name	PGs	Rule Name	Size
Buckets Index	<input type="text" value="default.rgw.buckets.index"/>	<input type="text" value="16"/>	<input type="text" value="replicated_rule"/>	<input type="text" value="3"/>

Storage Classes Buckets Data Pools: +

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

Enter the pools specifications and click save

From the Manage S3 , you need to Add a S3 User



The screenshot shows the 'Add S3 User' configuration page in the PetaSAN interface. The page includes the following fields and options:

- ID:** 1
- Display Name:** Tom Alex
- Email:** Email...
- 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
- User Quota:**
 - Max Size:** Unlimited Limit Size
 - Max Number of Objects:** Unlimited Limit Number

Buttons: Cancel, Save

You can enter the user id and name and specify its default placement target then click save

After saving you will find that the user's Access Key Id and the user's Secret Key Id has been created

Edit S3 User

✓ User saved successfully.

ID:
1

Display Name:
Tom Alex

Email:

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

User Quota

Max Size: Unlimited Limit Size

Max Number of Objects: Unlimited Limit Number

Access Key ID:
5J85BXHN88EN4L0RSK18

Secret Access Key:
0suysOmWg6mwCNgMvF3EpPFxc8QoBqHXoMvXQJYT

Cancel Regenerate Keys Save

Use the access key and secret key to connect to one of the load balancer using any of the S3 Clients. It is recommended to setup a round robin DNS to serve public ips to clients, for proper https validation you should specify the CN common name of the https certificate as the name to resolve by the DNS.