Setup S3 Object Storage using PetaSAN

Version 1.0

Revision History

Date	Version	Description
26/5/2022	1.0	Initial version

Contents

1. Purpose 4
2. Pre-requisites 4
3. Single Site Installation 5
3.1. Configuring S3 5
3.1.1. S3 Settings 5
3.1.2. Assign S3 Role to nodes 7
3.1.3. Add Zonegroup 8
3.1.4. Add Local Zone 9
3.1.5. Add S3 user 11
4. Clients Connectivity 14
4.1. S3 Browser 14
4.1.1. Define service ip in hosts file 14
4.1.2. Import s3-service certificate 14
4.1.3. Create S3 Browser user account 21
4.2. Cyberduck 25
4.2.1. Define certificate in hosts file 25
4.2.2. Import s3-service certificate 25
4.2.3. Connect using cyberduck 25
4.3. Amazon CLI Tool 27
4.3.1. Define certificate in hosts file 27
4.3.2. Import s3-service certificate 27
4.3.3. Configure the aws using configure command 27
4.3.4. Get bucket List 28
4.3.5. Create new Bucket 28
4.3.6. Upload file 28
4.3.7. List bucket content 28
5. Multi Site Installation 29
5.1. Configuring S3 29
5.1.1. S3 Settings 29

5.1.2. Assign S3 Role to nodes 29
5.1.3. Add End Point to the Master Zone 29
5.1.4. Import Peer Certificate 30
5.1.5. Define the service names in hosts files 31
5.1.6. Import Peer Certificate 33
5.1.7. Pull First Cluster S3 Configuration 33
5.1.8. Add local zone 34
6. Add S3 User 36
7. Client Connectivity 36
7.1. S3 Browser 36
7.1.1. Define certificate in hosts file 36
7.1.2. Import s3-service certificate 37
7.1.3. Create S3 Browser user account 37
7.2. Cyberduck 39
7.3. Amazon CLI Tool 39
8. Promote Zone 39

1. Purpose

The purpose of this document is to describe how to setup S3 Object storage using PetaSAN.

2. Pre-requisites

This document assumes the user has already setup a PetaSAN cluster and optionally a second cluster in case a multisite S3 setup is required.

The examples used in this guide assume the following sample configuration:

Single Site Installation

- One PetaSAN cluster with Release 3.0.0 or higher installed, each cluster consists of 3 nodes with no pools created during deployment.
- Each node has 3 interfaces:
  - **Management** uses subnet ip 10.0.1.0 and subnet mask 255.255.255.0
  - **Backend** uses subnet ip 10.0.2.0 and subnet mask 255.255.255.0
  - **S3** uses subnet ip 10.0.3.0 and subnet mask 255.255.255.0
    (Defined post installation as will be shown)
- Nodes have the following IPs:
  - **Node1**
    - Management uses subnet ip 10.0.1.10
    - Backend uses subnet ip 10.0.2.10
  - **Node2**
    - Management uses subnet ip 10.0.1.11
    - Backend uses subnet ip 10.0.2.11
  - **Node3**
    - Management uses subnet ip 10.0.1.12
    - Backend uses subnet ip 10.0.2.12
- An EC rule is created based on template ec-by-host-hdd (assuming HDD drives)

Multi Site Installation

- A second PetaSAN cluster with Release 3.0.0 or higher installed, cluster consists of 3 nodes with no pools created during deployment.
- Each node has 3 interfaces:
  - **Management** uses subnet ip 10.0.1.0 and subnet mask 255.255.255.0
  - **Backend** uses subnet ip 10.0.2.0 and subnet mask 255.255.255.0
  - **S3 Public** uses subnet ip 10.0.3.0 and subnet mask 255.255.255.0
    (Defined post installation as will be shown).
- Nodes have the following IPs:
  - **Node1**
    - Management uses subnet ip 10.0.1.90
    - Backend uses subnet ip 10.0.2.90

    - **Node2**
        - Management uses subnet ip 10.0.1.91
        - Backend uses subnet ip 10.0.2.91
    - **Node3**
        - Management uses subnet ip 10.0.1.92
        - Backend uses subnet ip 10.0.2.92
- An EC rule is created based on template ec-by-host-hdd hdd (assuming HDD drives).

3. Single Site Installation

3.1. Configuring S3

3.1.1. S3 Settings

- Define protocol as http or https, interface to use and the IP range for the S3 service. You can also define a custom gateway.

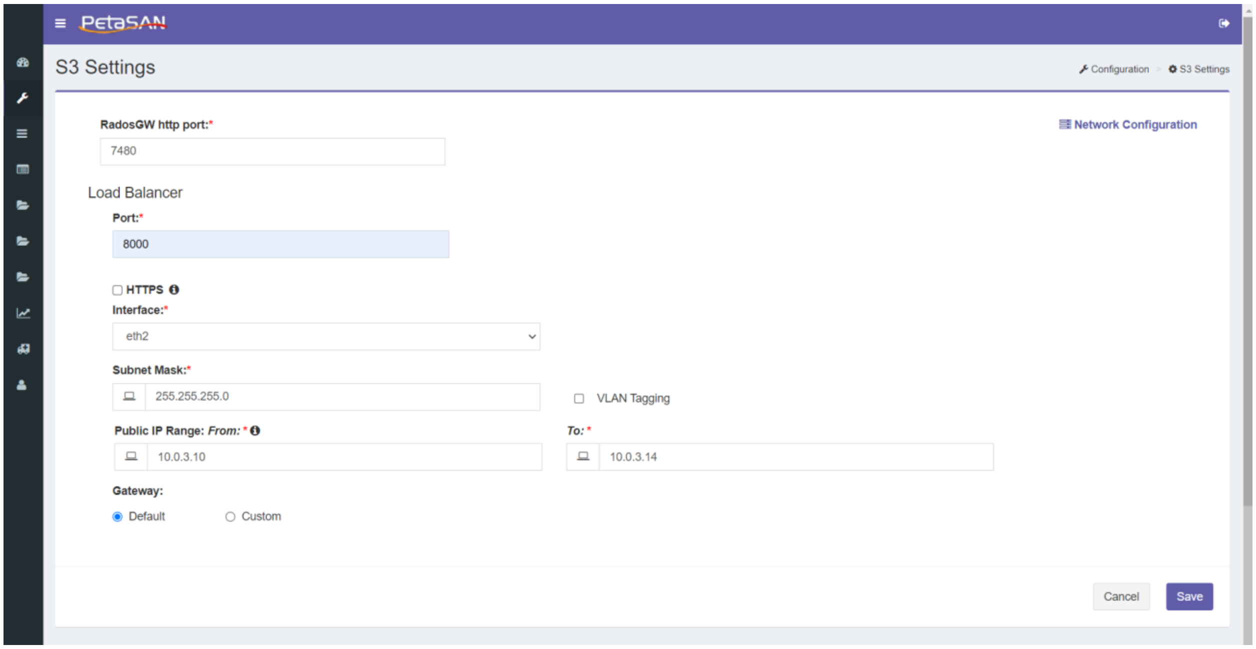

S3 Settings

RadosGW http port:* 7480

Load Balancer

Port:* 8000

□ HTTPS

Interface:* eth2

Subnet Mask:* 255.255.255.0

□ VLAN Tagging

Public IP Range: *From:** 10.0.3.10 *To:** 10.0.3.14

Gateway:

● Default ○ Custom

Note:

It is recommended to use https when using Veeam.

To use https with a self signed certificate, generate the certificate by clicking on the Generate button and specify the certificate common name. The common name needs to match the S3 service URL that will clients will use to access the service, it will be configured in the hosts file or round robin DNS as will be discussed later.

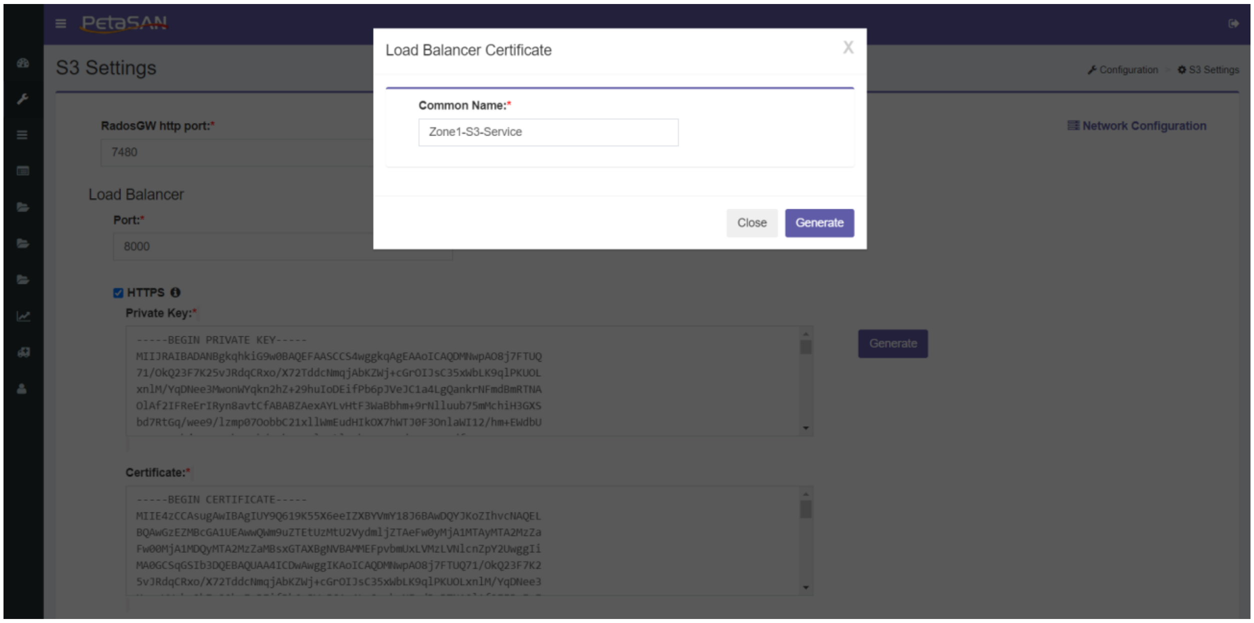

RadosGW http port:*
7480

Load Balancer

Port:*
8000

☑ HTTPS

Private Key:*

```
-----BEGIN PRIVATE KEY-----
MIIJRAIBADANBgkqhkiG9w0BAQEFAASCCS4wggkqAgEAAoICAQDMNwpAO8j7FTUQ
71/OkQ23F7K25vJRdqCRxo/X72TddcNmqjAbKZWj+cGrOIJsC35xWbLK9qlPKUOL
xnlM/YqDNee3MwonWYqkn2hZ+29huIoDEifPb6pJVeJC1a4LgQankrNFmdBmRTNA
OlAf2IFReErIRyn8avtCfABABZAexAYLvHtF3WaBbhm+9rNlluub75mMchiH3GXS
bd7RtGq/wee9/lzmp07OobbC21xllWmEudHIkOX7hWTJ0F3OnlaWI12/hm+EWdbU
```

Certificate:*

```
-----BEGIN CERTIFICATE-----
MIIE4zCCAsugAwIBAgIUY9Q619K55X6eeIZXBYVmY18J6BAwDQYJKoZIhvcNAQEL
BQAwGzEZMBcGA1UEAwwQWm9uZTEtUzMtU2VydmljZTAeFw0yMjA1MTAyMTA2MzZa
Fw00MjA1MDQyMTA2MzZaMBsxGTAXBgNVBAMMEFpvbmUxLVMzLVNlcnZpY2UwggIi
MA0GCSqGSIb3DQEBAQUAA4ICDwAwggIKAoICAQDMNwpAO8j7FTUQ71/OkQ23F7K2
5vJRdqCRxo/X72TddcNmqjAbKZWj+cGrOIJsC35xWbLK9qlPKUOLxnlM/YqDNee3
```

Common Name:*
Zone1-S3-Service

After entering the interface name and the IP range, click on save.

In the case of self signed certificate as we use in this example, download the certificate, so it can be later installed on S3 clients.

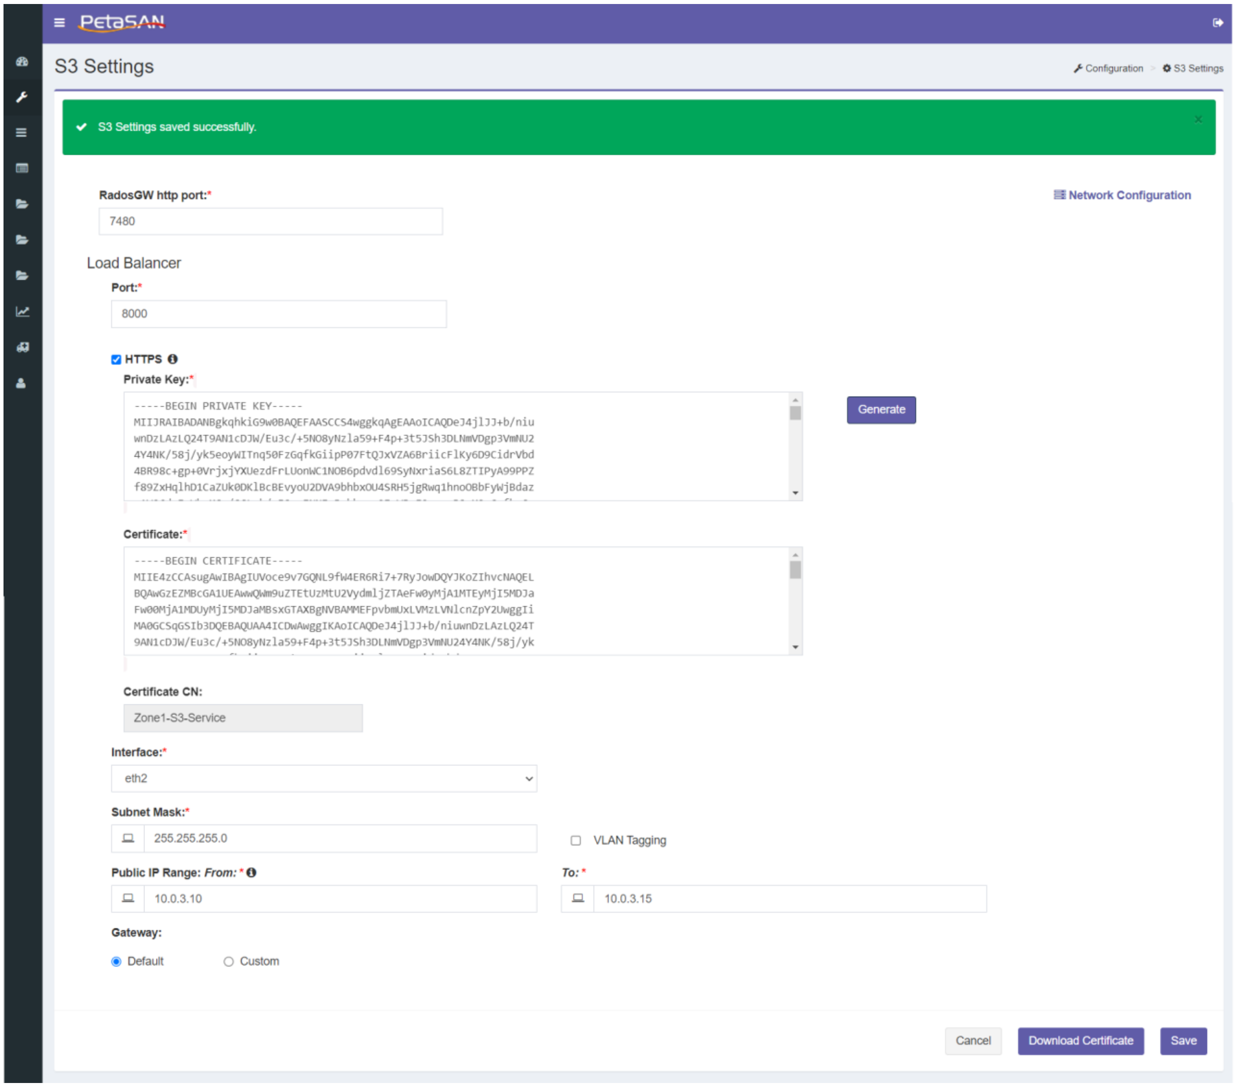

S3 Settings

✓ S3 Settings saved successfully.

RadosGW http port:* 7480

Load Balancer

Port:* 8000

☑ HTTPS

Private Key:*

```
-----BEGIN PRIVATE KEY-----
MIIJRAIBADANBgkqhkiG9w0BAQEFAASCCS4wggkqAgEAAoICAQDeJ4jlJJ+b/niu
wnDzLAzLQ24T9AN1cDJW/Eu3c/+5NO8yNzla59+F4p+3t5JSh3DLNmVDgp3VmNU2
4Y4NK/58j/yk5eoyWITnq50FzGqfkGiipP07FtQJxVZA6BriicFlKy6D9CidrVbd
4BR98c+gp+0VrjxjYXUezdFrLUonWC1NOB6pdvdl69SyNxriaS6L8ZTIPyA99PPZ
f89ZxHqlhD1CaZUk0DKlBcBEvyoU2DVA9bhbxOU4SRH5jgRwq1hnoOBbFyWjBdaz
```

Certificate:*

```
-----BEGIN CERTIFICATE-----
MIIE4zCCAsugAwIBAgIUVoce9v7GQNL9fW4ER6Ri7+7RyJowDQYJKoZIhvcNAQEL
BQAwGzEZMBcGA1UEAwwQWm9uZTEtUzMtU2VydmljZTAeFw0yMjA1MTEyMjI5MDJa
Fw00MjA1MDUyMjI5MDJaMBsxGTAXBgNVBAMMEFpvbmUxLVMzLVNlcnZpY2UwggIi
MA0GCSqGSIb3DQEBAQUAA4ICDwAwggIKAoICAQDeJ4jlJJ+b/niuwnDzLAzLQ24T
9AN1cDJW/Eu3c/+5NO8yNzla59+F4p+3t5JSh3DLNmVDgp3VmNU24Y4NK/58j/yk
```

Certificate CN: Zone1-S3-Service

Interface:* eth2

Subnet Mask:* 255.255.255.0

□ VLAN Tagging

Public IP Range: *From:** 10.0.3.10 To:* 10.0.3.15

Gateway:

● Default ○ Custom

3.1.2. Assign S3 Role to nodes

- Start by assigning the S3 Role to one or more cluster nodes, in this example we will assign the S3 role to all the 3 nodes.

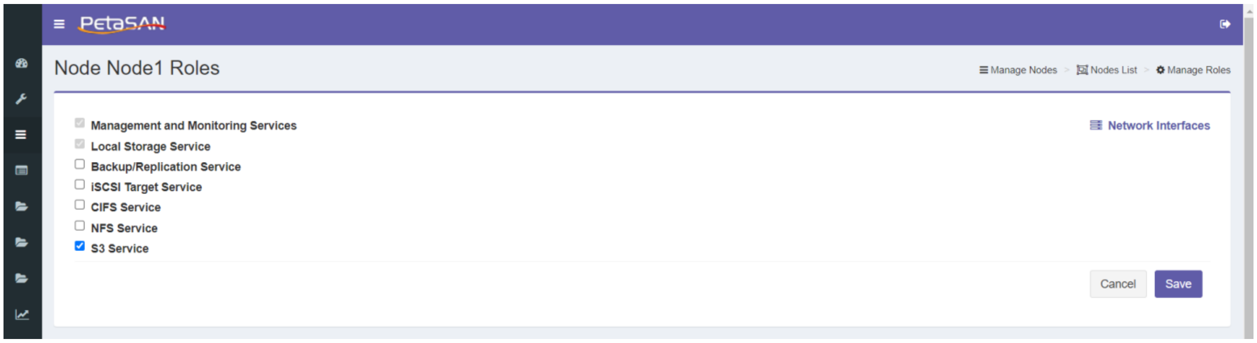

Node Node1 Roles

☑ Management and Monitoring Services

☑ Local Storage Service

□ Backup/Replication Service

□ iSCSI Target Service

□ CIFS Service

□ NFS Service

☑ S3 Service

3.1.3. Add Zonegroup

Add the zonegroup from Configuration->S3 Configuration->Zonegroups->Add Zonegroup

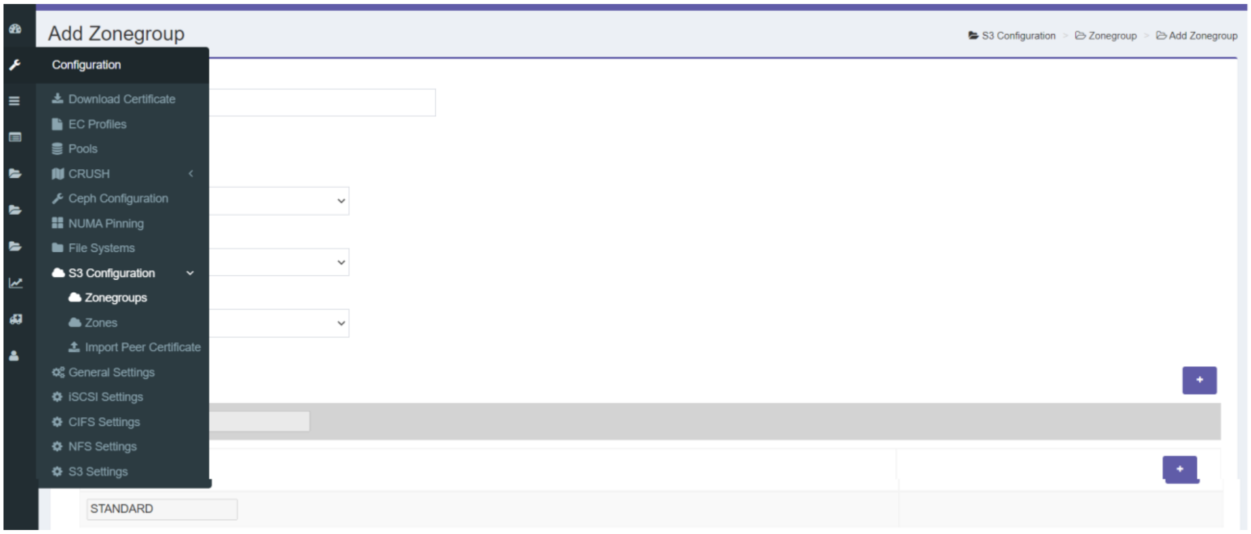

- Enter the zonegroup name, the root pool settings and the placement targets.

  In this example we will create zone named Zonegroup1 with 2 placement targets, first is named "default-placement" and the second is named "backups"

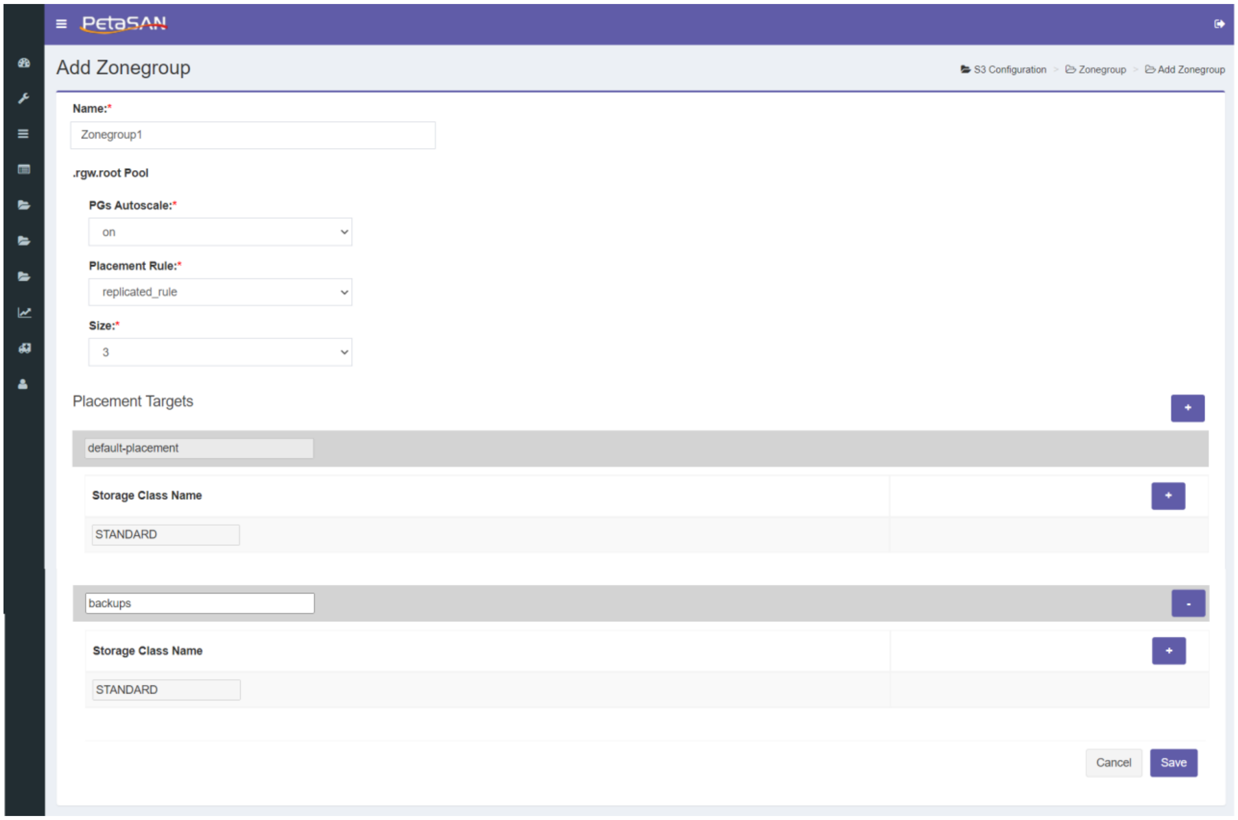

Name:* Zonegroup1

.rgw.root Pool

PGs Autoscale:* on

Placement Rule:* replicated_rule

Size:* 3

Placement Targets

default-placement

Storage Class Name
STANDARD

backups

Storage Class Name
STANDARD

- You can also add storage classes but in this example we will only use the “STANDARD” storage class.

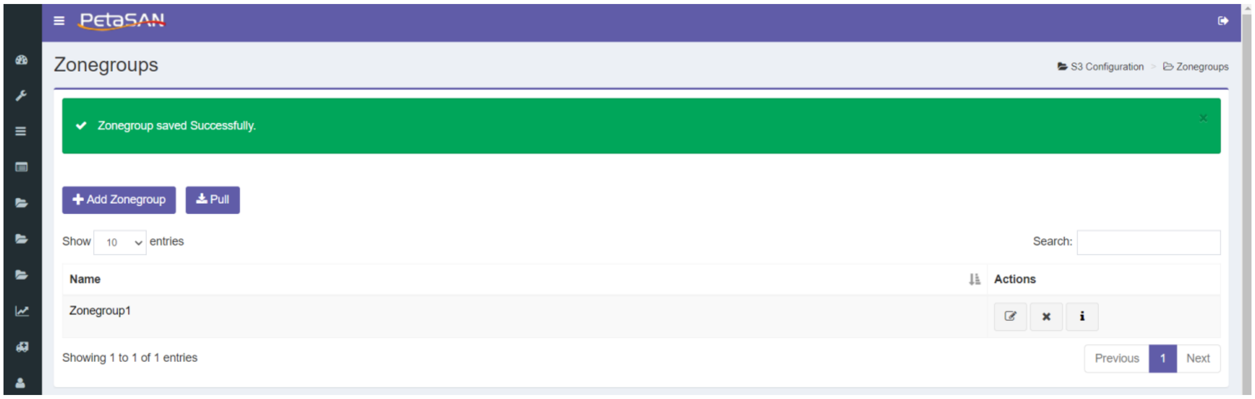

✓ Zonegroup saved Successfully.

Name	Actions
Zonegroup1	

Showing 1 to 1 of 1 entries

3.1.4. Add Local Zone

- Add zone from Configuration->S3 Configuration->Zones->Add Zone.
- In the zone form, The system will display the main pools that will be used for internal operations (Control,Meta and log pools). You can optionally define their crush placement rule and replica count if desired. In this example we will use the defaults.
- We configure the pools that will be created for each placement target (Bucket Index and Storage class bucket data pools).We need to define their crush placement rule and replica count.
- In this example we will create the main pools and the default-placement pools using the replicated rule which is selected by default.
- For the backups placement target we will create a data pool with EC rule “ec-by-host-hdd” and profile “ec-21-profile” for testing purpose but in production you should use a higher profile like “ec-42-profile”

  Note:
  Using an EC data pool is for ideal for backups due to the storage efficiency.

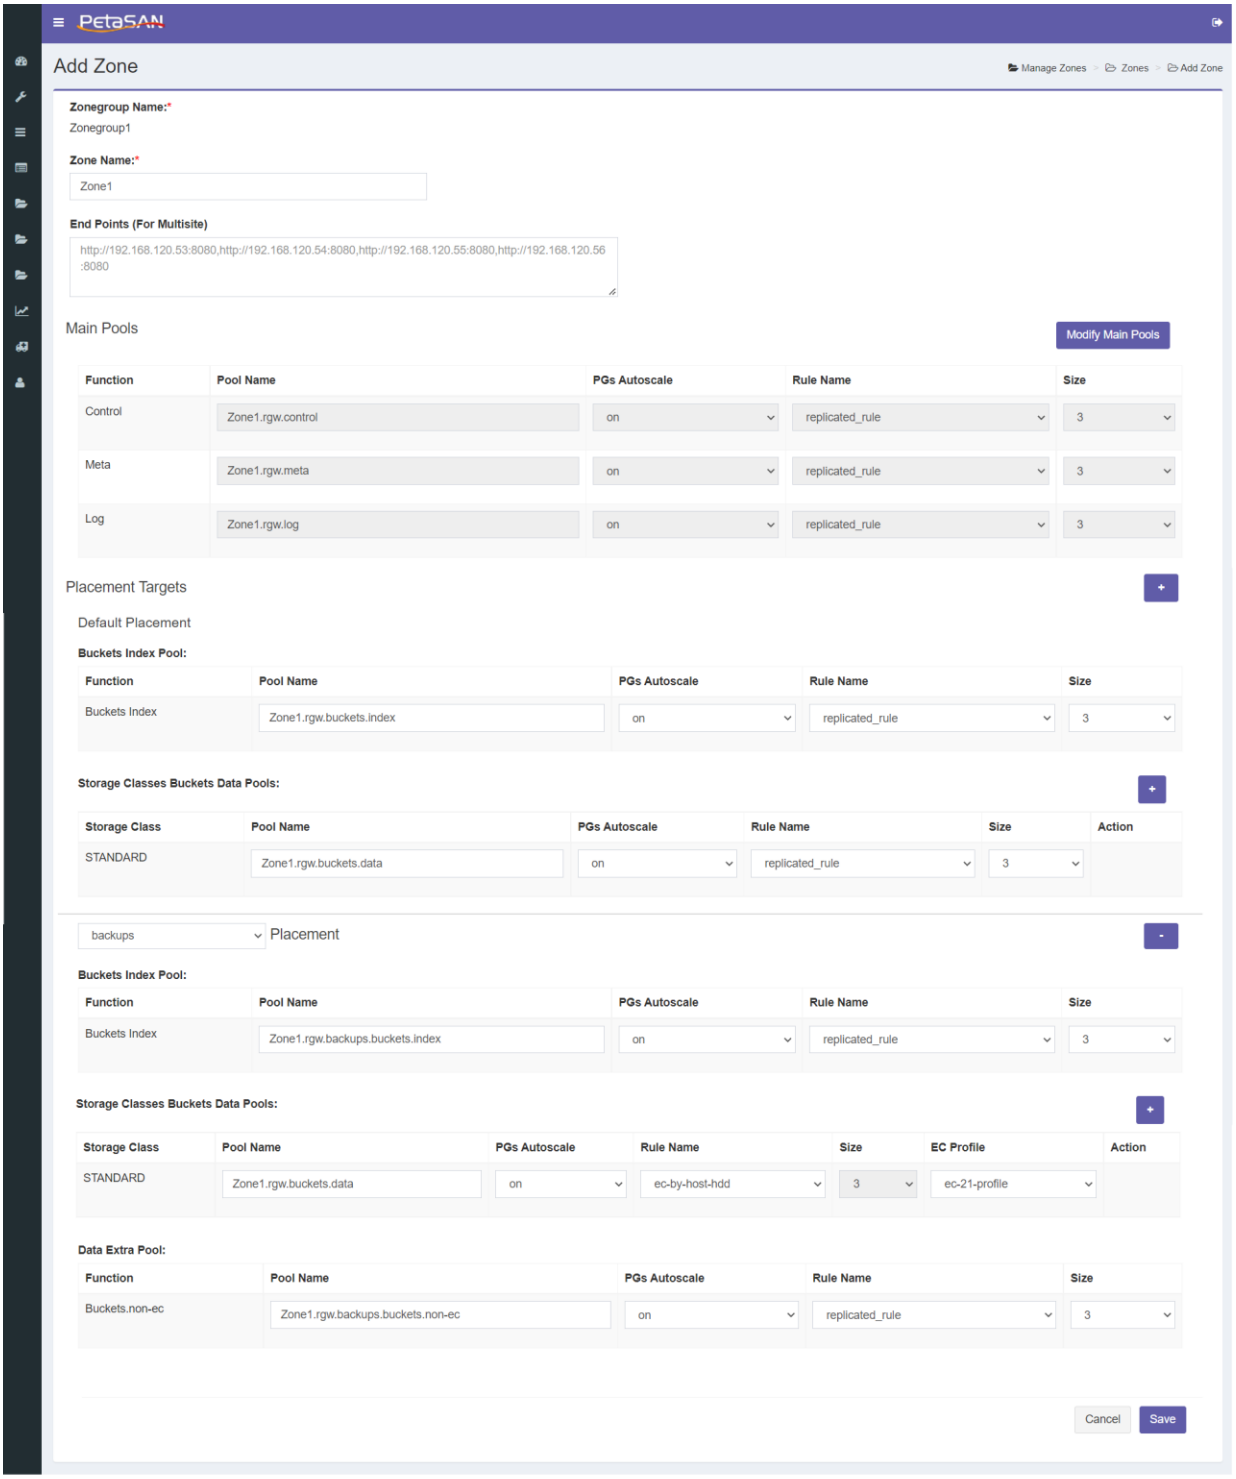

Add Zone

Zonegroup Name:*
Zonegroup1

Zone Name:*
Zone1

End Points (For Multisite)
http://192.168.120.53:8080,http://192.168.120.54:8080,http://192.168.120.55:8080,http://192.168.120.56:8080

Main Pools

Function	Pool Name	PGs Autoscale	Rule Name	Size
Control	Zone1.rgw.control	on	replicated_rule	3
Meta	Zone1.rgw.meta	on	replicated_rule	3
Log	Zone1.rgw.log	on	replicated_rule	3

Placement Targets

Default Placement

Buckets Index Pool:

Function	Pool Name	PGs Autoscale	Rule Name	Size
Buckets Index	Zone1.rgw.buckets.index	on	replicated_rule	3

Storage Classes Buckets Data Pools:

Storage Class	Pool Name	PGs Autoscale	Rule Name	Size	Action
STANDARD	Zone1.rgw.buckets.data	on	replicated_rule	3	

backups Placement

Buckets Index Pool:

Function	Pool Name	PGs Autoscale	Rule Name	Size
Buckets Index	Zone1.rgw.backups.buckets.index	on	replicated_rule	3

Storage Classes Buckets Data Pools:

Storage Class	Pool Name	PGs Autoscale	Rule Name	Size	EC Profile	Action
STANDARD	Zone1.rgw.buckets.data	on	ec-by-host-hdd	3	ec-21-profile	

Data Extra Pool:

Function	Pool Name	PGs Autoscale	Rule Name	Size
Buckets.non-ec	Zone1.rgw.backups.buckets.non-ec	on	replicated_rule	3

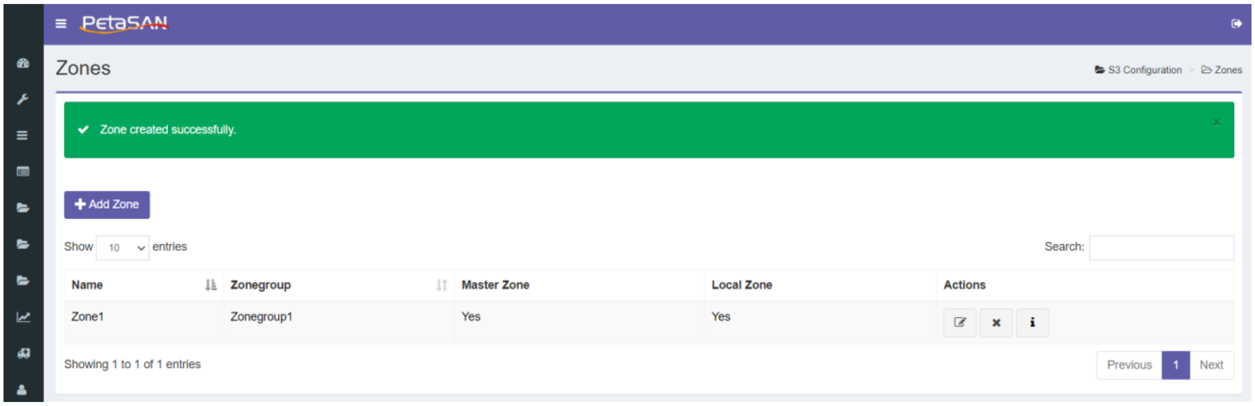

Zones

S3 Configuration > Zones

✔ Zone created successfully.

Show 10 entries

Search:

Name	Zonegroup	Master Zone	Local Zone	Actions
Zone1	Zonegroup1	Yes	Yes	

Showing 1 to 1 of 1 entries

After creating the zone you could notice that the new pools have been created

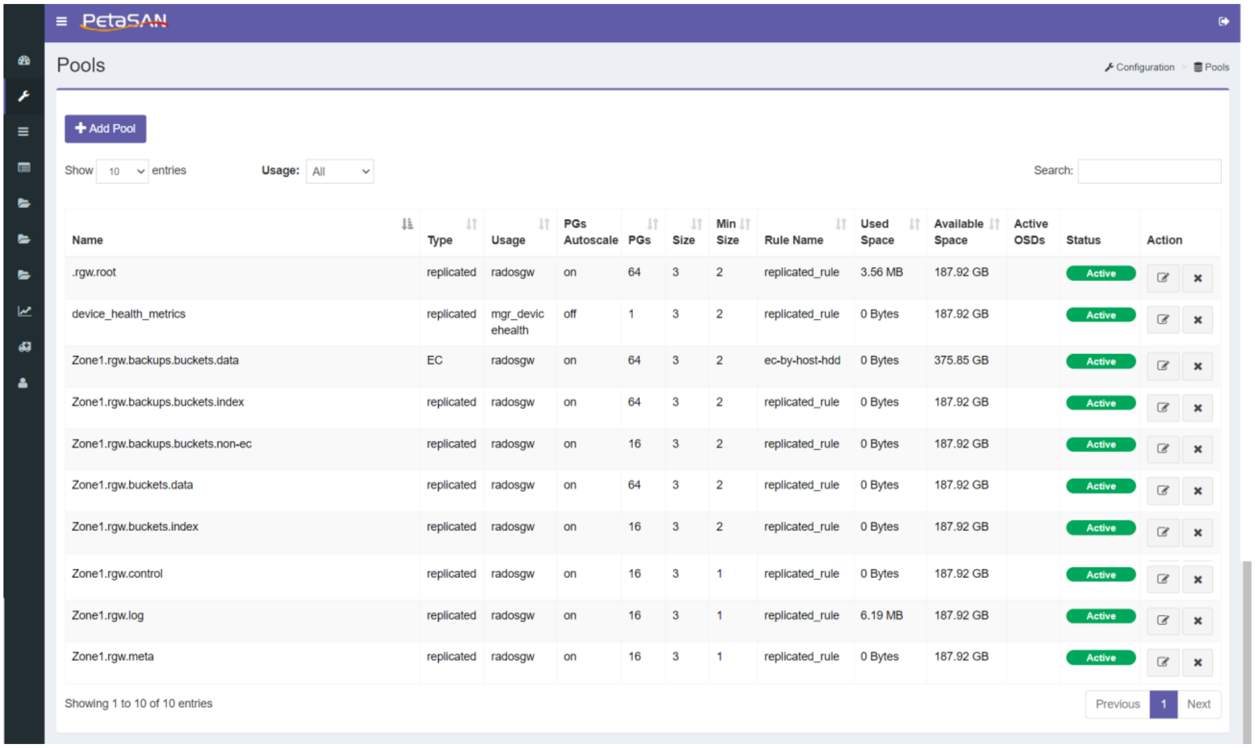

Pools

Configuration > Pools

Show 10 entries

Usage: All

Search:

Name	Type	Usage	PGs Autoscale	PGs	Size	Min Size	Rule Name	Used Space	Available Space	Active OSDs	Status	Action
.rgw.root	replicated	radosgw	on	64	3	2	replicated_rule	3.56 MB	187.92 GB		Active	
device_health_metrics	replicated	mgr_devicehealth	off	1	3	2	replicated_rule	0 Bytes	187.92 GB		Active	
Zone1.rgw.backups.buckets.data	EC	radosgw	on	64	3	2	ec-by-host-hdd	0 Bytes	375.85 GB		Active	
Zone1.rgw.backups.buckets.index	replicated	radosgw	on	64	3	2	replicated_rule	0 Bytes	187.92 GB		Active	
Zone1.rgw.backups.buckets.non-ec	replicated	radosgw	on	16	3	2	replicated_rule	0 Bytes	187.92 GB		Active	
Zone1.rgw.buckets.data	replicated	radosgw	on	64	3	2	replicated_rule	0 Bytes	187.92 GB		Active	
Zone1.rgw.buckets.index	replicated	radosgw	on	16	3	2	replicated_rule	0 Bytes	187.92 GB		Active	
Zone1.rgw.control	replicated	radosgw	on	16	3	1	replicated_rule	0 Bytes	187.92 GB		Active	
Zone1.rgw.log	replicated	radosgw	on	16	3	1	replicated_rule	6.19 MB	187.92 GB		Active	
Zone1.rgw.meta	replicated	radosgw	on	16	3	1	replicated_rule	0 Bytes	187.92 GB		Active	

Showing 1 to 10 of 10 entries

3.1.5. Add S3 user

- In this example we will create a new S3 user named S3-User1 that stores its data in the backups placement target which we created earlier.
- This can be done from Manage S3->S3 Users Menu item

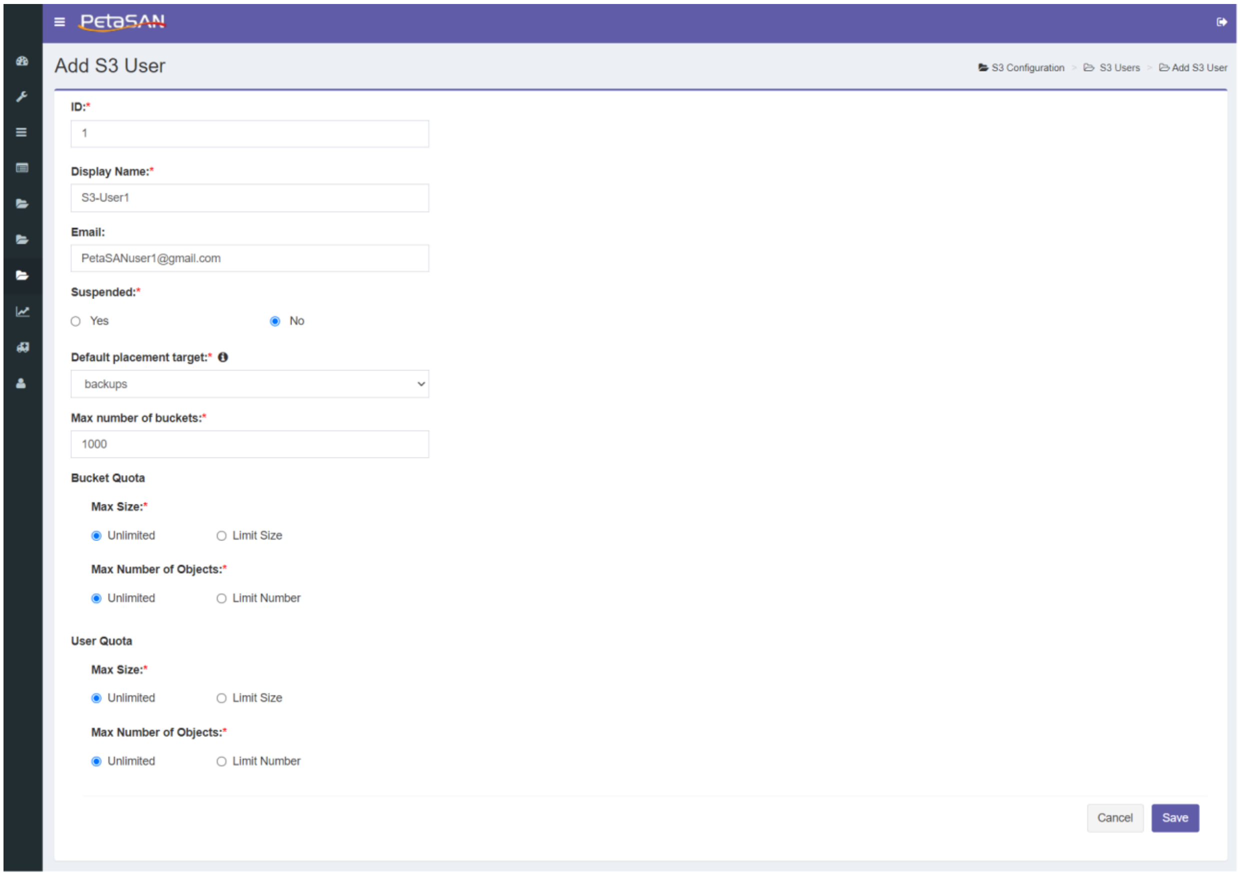

Add S3 User

ID:* 1

Display Name:* S3-User1

Email: PetaSANuser1@gmail.com

Suspended:* ○ Yes ● No

Default placement target:* backups

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

After saving the S3 user you can view the user's generated access and secret keys which will be used later while connecting with S3 clients.

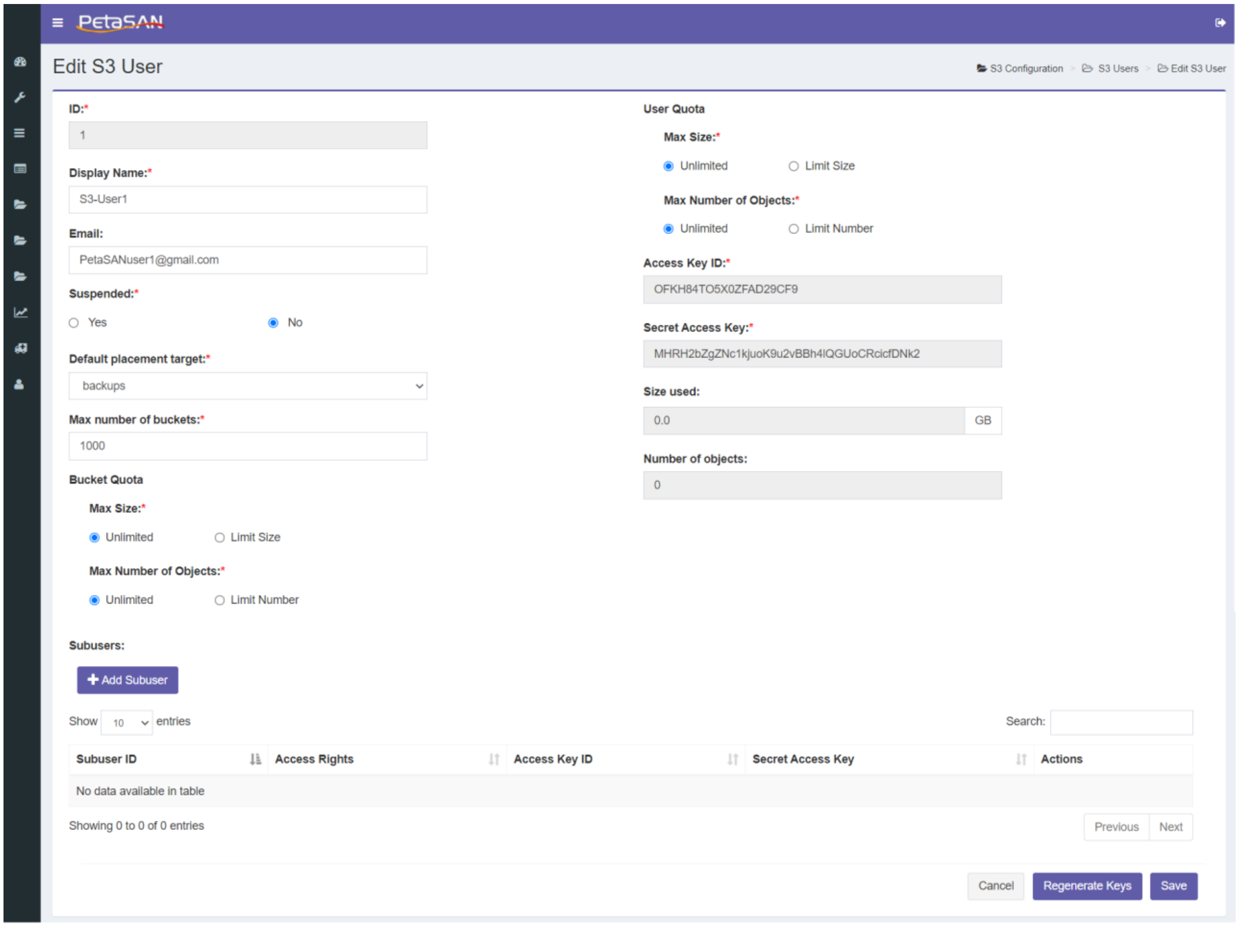

Edit S3 User

ID:* 1

Display Name:* S3-User1

Email: PetaSANuser1@gmail.com

Suspended:* ○ Yes ● No

Default placement target:* backups

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:* OFKH84TO5X0ZFAD29CF9

Secret Access Key:* MHRH2bZgZNc1kjuoK9u2vBBh4IQGUoCRcicfDNk2

Size used: 0.0 GB

Number of objects: 0

Subusers:

Show 10 entries

Subuser ID	Access Rights	Access Key ID	Secret Access Key	Actions
No data available in table				

Showing 0 to 0 of 0 entries

Note:

- You can set maximum number of buckets the user can create or set the max size or maximum number of objects the user can upload for any bucket.
- You can add one or more Sub User under the user you created, for each sub user you need to set the sub user id and its Access Right, but in this example we will not create any sub users.

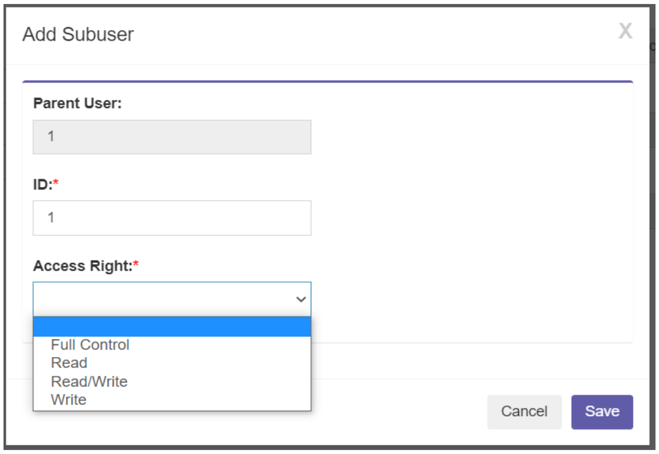

Add Subuser

Parent User: 1

ID:* 1

Access Right:*

4. Clients Connectivity

4.1. S3 Browser

- One of the client applications that is widely used is S3 Browser.
- We will be using the self-signed certificate previously created from the S3 Settings form.

4.1.1. Define service ip in hosts file

- If not using a DNS, define the ip address corresponding to the S3 service. The name of the service should match the common name of the certificate. The IP address should be one within the range of IPs defined for the service.
- If using a DNS, setup a round robin configuration with all the range of IPs defined. In this example we will use the hosts file method.

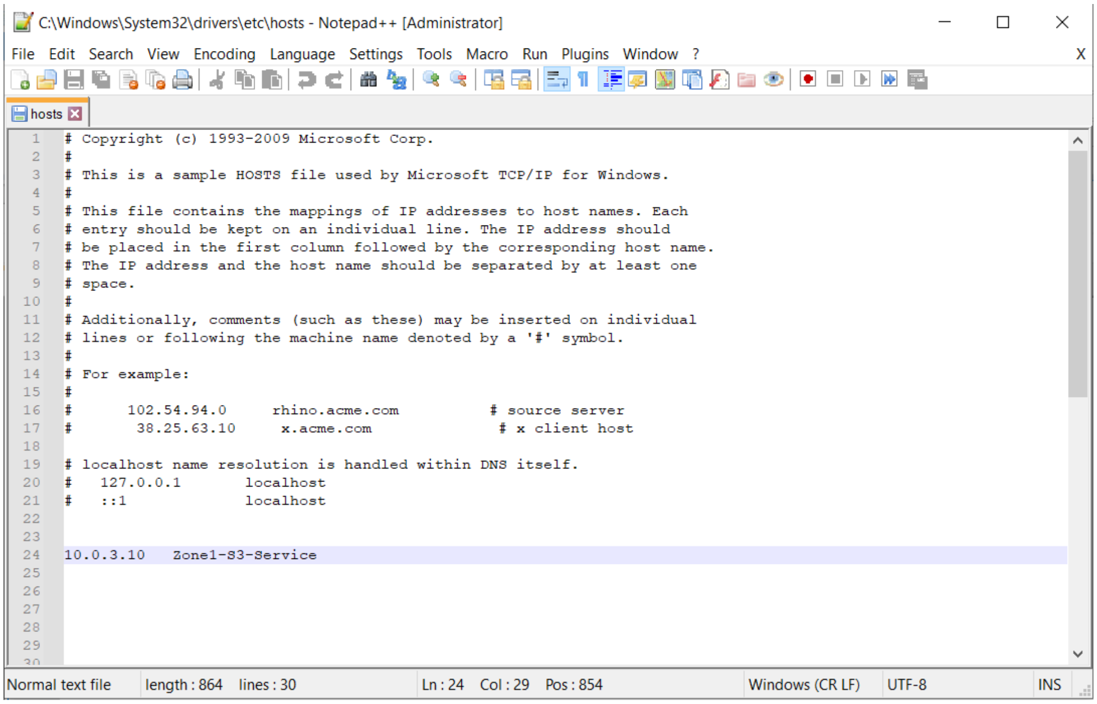
C:\Windows\System32\drivers\etc\hosts - Notepad++ [Administrator]
File Edit Search View Encoding Language Settings Tools Macro Run Plugins Window ?
hosts
Normal text file
length : 864 lines : 30
Ln : 24 Col : 29 Pos : 854
Windows (CR LF)
UTF-8
INS

```
# Copyright (c) 1993-2009 Microsoft Corp.
#
# This is a sample HOSTS file used by Microsoft TCP/IP for Windows.
#
# This file contains the mappings of IP addresses to host names. Each
# entry should be kept on an individual line. The IP address should
# be placed in the first column followed by the corresponding host name.
# The IP address and the host name should be separated by at least one
# space.
#
# Additionally, comments (such as these) may be inserted on individual
# lines or following the machine name denoted by a '#' symbol.
#
# For example:
#
#      102.54.94.0     rhino.acme.com          # source server
#       38.25.63.10     x.acme.com              # x client host

# localhost name resolution is handled within DNS itself.
#	127.0.0.1       localhost
#	::1             localhost


10.0.3.10   Zone1-S3-Service
```

4.1.2. Import s3-service certificate

- In Chrome settings, select privacy and security then Security.

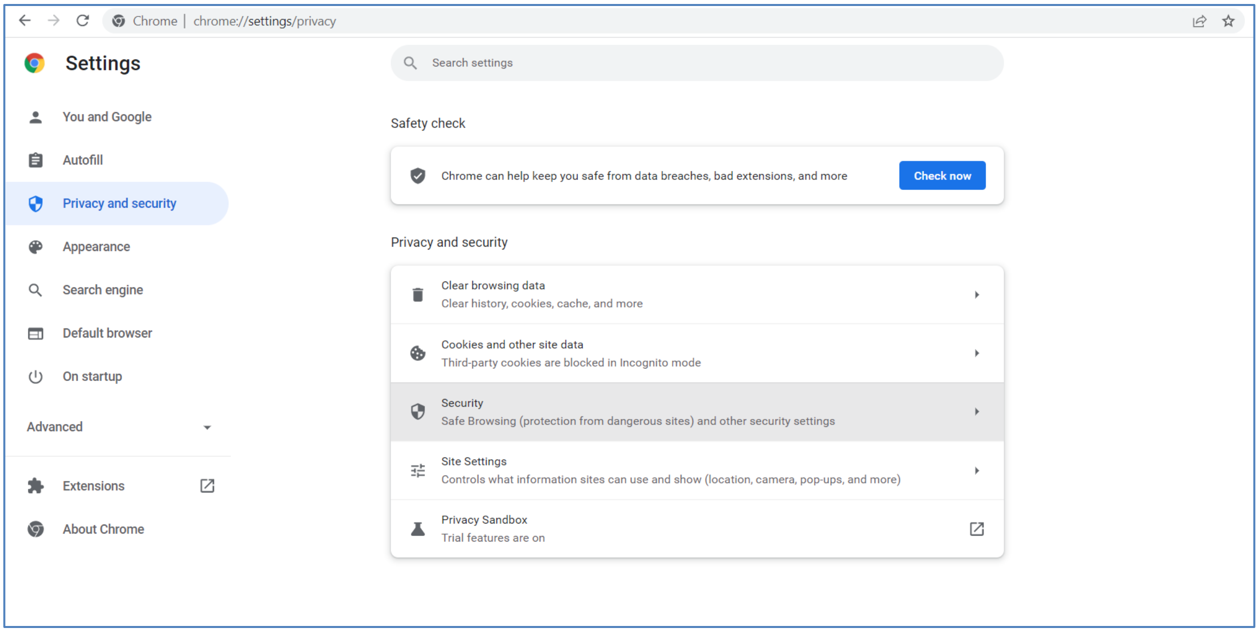

Safety check

Chrome can help keep you safe from data breaches, bad extensions, and more

Privacy and security

Clear browsing data
Clear history, cookies, cache, and more

Cookies and other site data
Third-party cookies are blocked in Incognito mode

Security
Safe Browsing (protection from dangerous sites) and other security settings

Site Settings
Controls what information sites can use and show (location, camera, pop-ups, and more)

Privacy Sandbox
Trial features are on

- Then select Manage Certificates

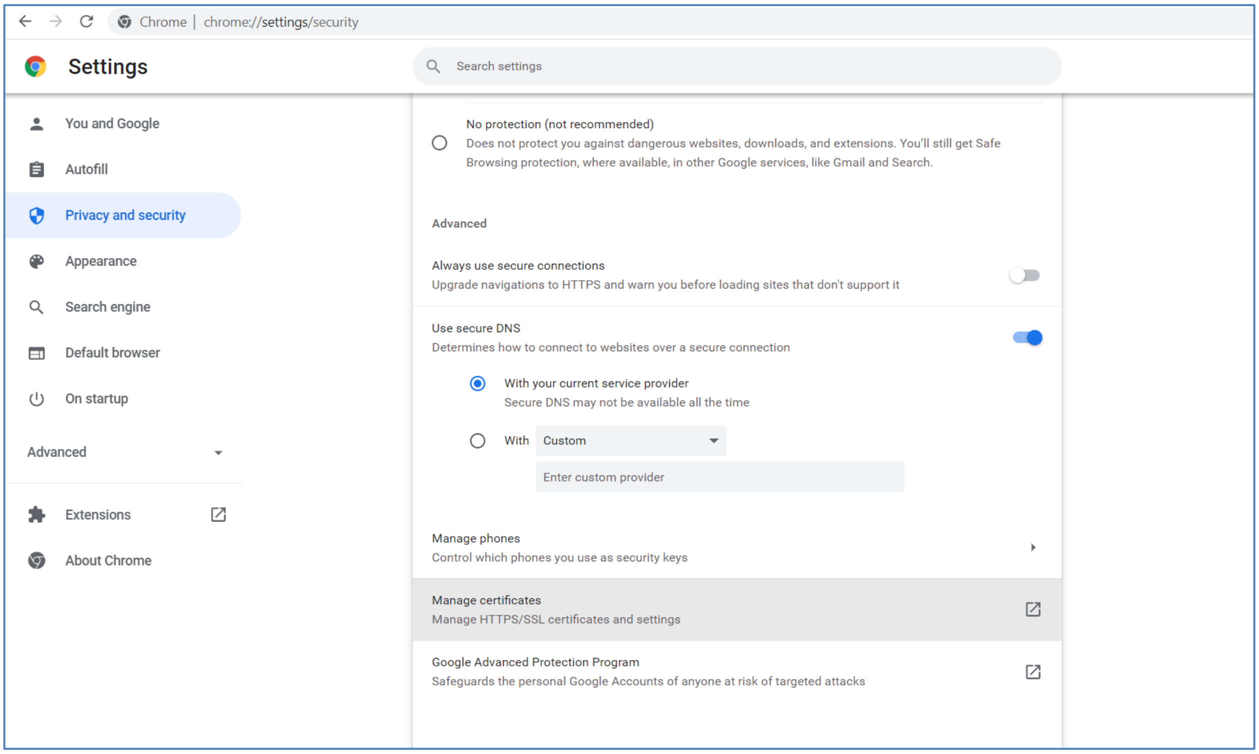

○ No protection (not recommended)
Does not protect you against dangerous websites, downloads, and extensions. You'll still get Safe Browsing protection, where available, in other Google services, like Gmail and Search.

Advanced

Always use secure connections
Upgrade navigations to HTTPS and warn you before loading sites that don't support it

Use secure DNS
Determines how to connect to websites over a secure connection

● With your current service provider
Secure DNS may not be available all the time

○ With Custom
Enter custom provider

Manage phones
Control which phones you use as security keys

Manage certificates
Manage HTTPS/SSL certificates and settings

Google Advanced Protection Program
Safeguards the personal Google Accounts of anyone at risk of targeted attacks

- Under Trusted Root Certification Authorities tab select import button

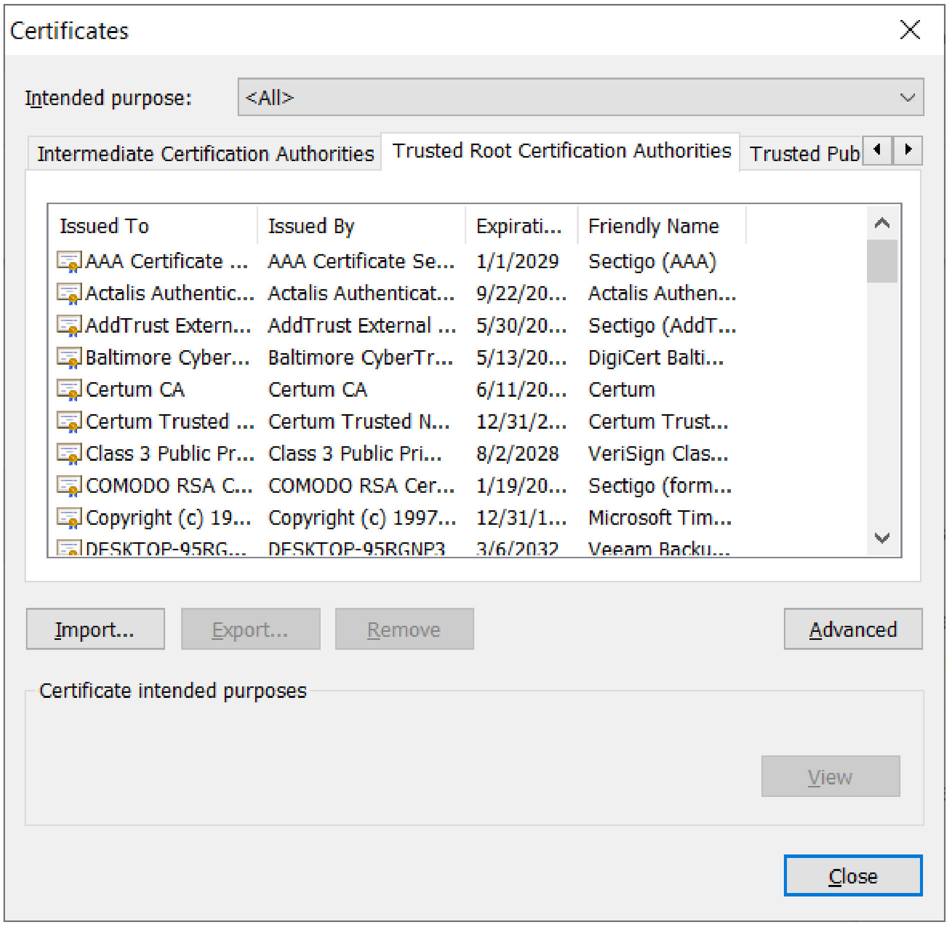
Certificates
Trusted Root Certification Authorities
Intermediate Certification Authorities
Trusted Pub
Import...
Export...
Remove
Advanced
Certificate intended purposes
View
Close

Intended purpose: <All>

Issued To	Issued By	Expirati...	Friendly Name
AAA Certificate ...	AAA Certificate Se...	1/1/2029	Sectigo (AAA)
Actalis Authentic...	Actalis Authenticat...	9/22/20...	Actalis Authen...
AddTrust Extern...	AddTrust External ...	5/30/20...	Sectigo (AddT...
Baltimore Cyber...	Baltimore CyberTr...	5/13/20...	DigiCert Balti...
Certum CA	Certum CA	6/11/20...	Certum
Certum Trusted ...	Certum Trusted N...	12/31/2...	Certum Trust...
Class 3 Public Pr...	Class 3 Public Pri...	8/2/2028	VeriSign Clas...
COMODO RSA C...	COMODO RSA Cer...	1/19/20...	Sectigo (form...
Copyright (c) 19...	Copyright (c) 1997...	12/31/1...	Microsoft Tim...
DESKTOP-95RG...	DESKTOP-95RGNP3	3/6/2032	Veeam Backu...

- Continue with the steps of importing the certificate as follows:

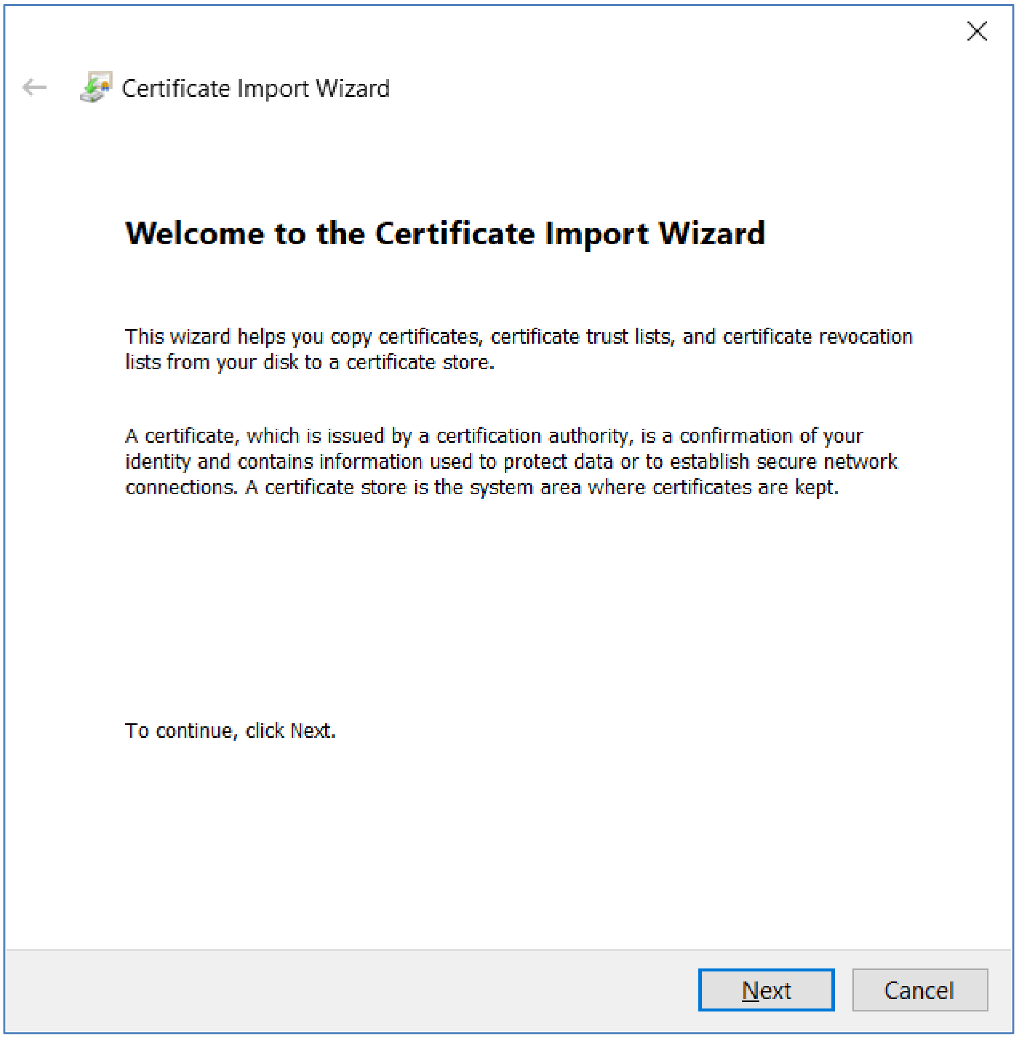

Welcome to the Certificate Import Wizard

This wizard helps you copy certificates, certificate trust lists, and certificate revocation lists from your disk to a certificate store.

A certificate, which is issued by a certification authority, is a confirmation of your identity and contains information used to protect data or to establish secure network connections. A certificate store is the system area where certificates are kept.

To continue, click Next.

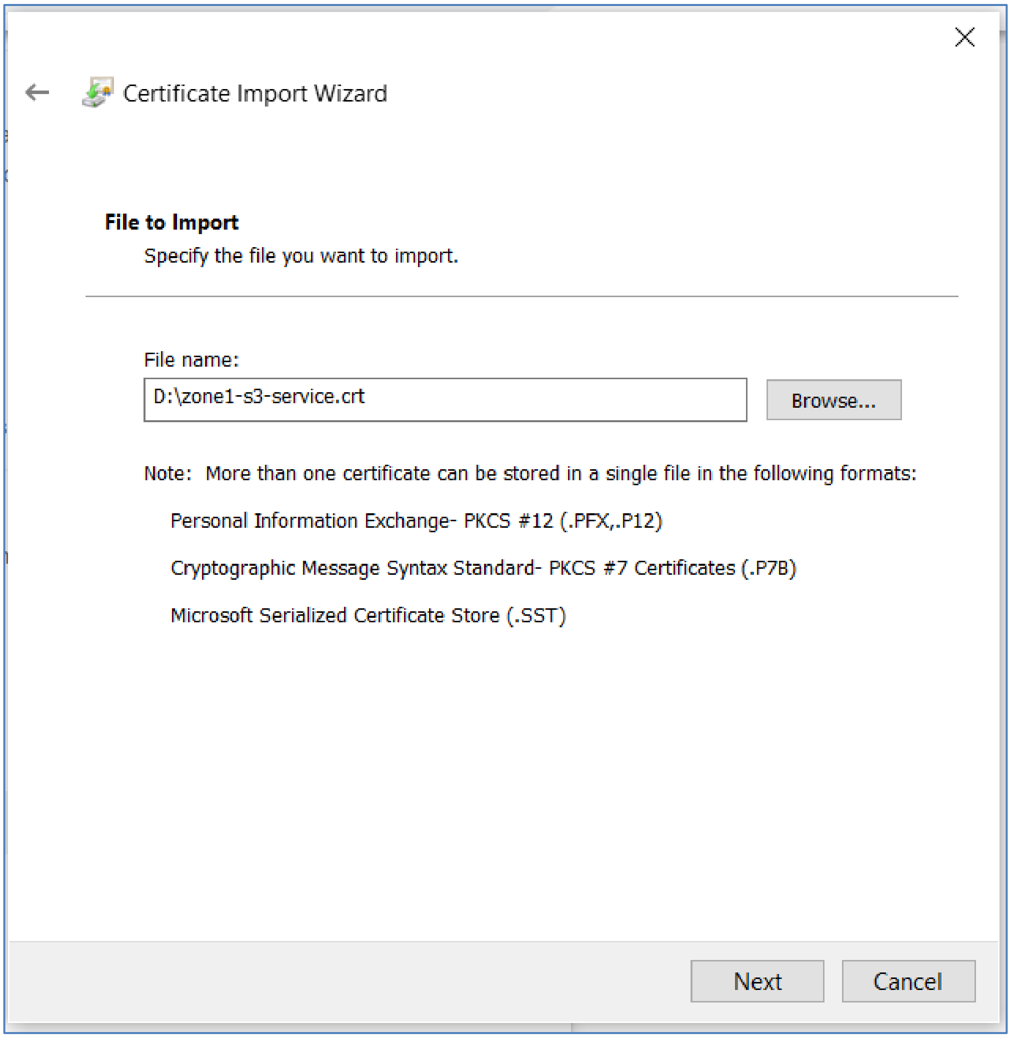

Certificate Import Wizard

File to Import

Specify the file you want to import.

File name:

D:\zone1-s3-service.crt

Note: More than one certificate can be stored in a single file in the following formats:

Personal Information Exchange- PKCS #12 (.PFX,.P12)

Cryptographic Message Syntax Standard- PKCS #7 Certificates (.P7B)

Microsoft Serialized Certificate Store (.SST)

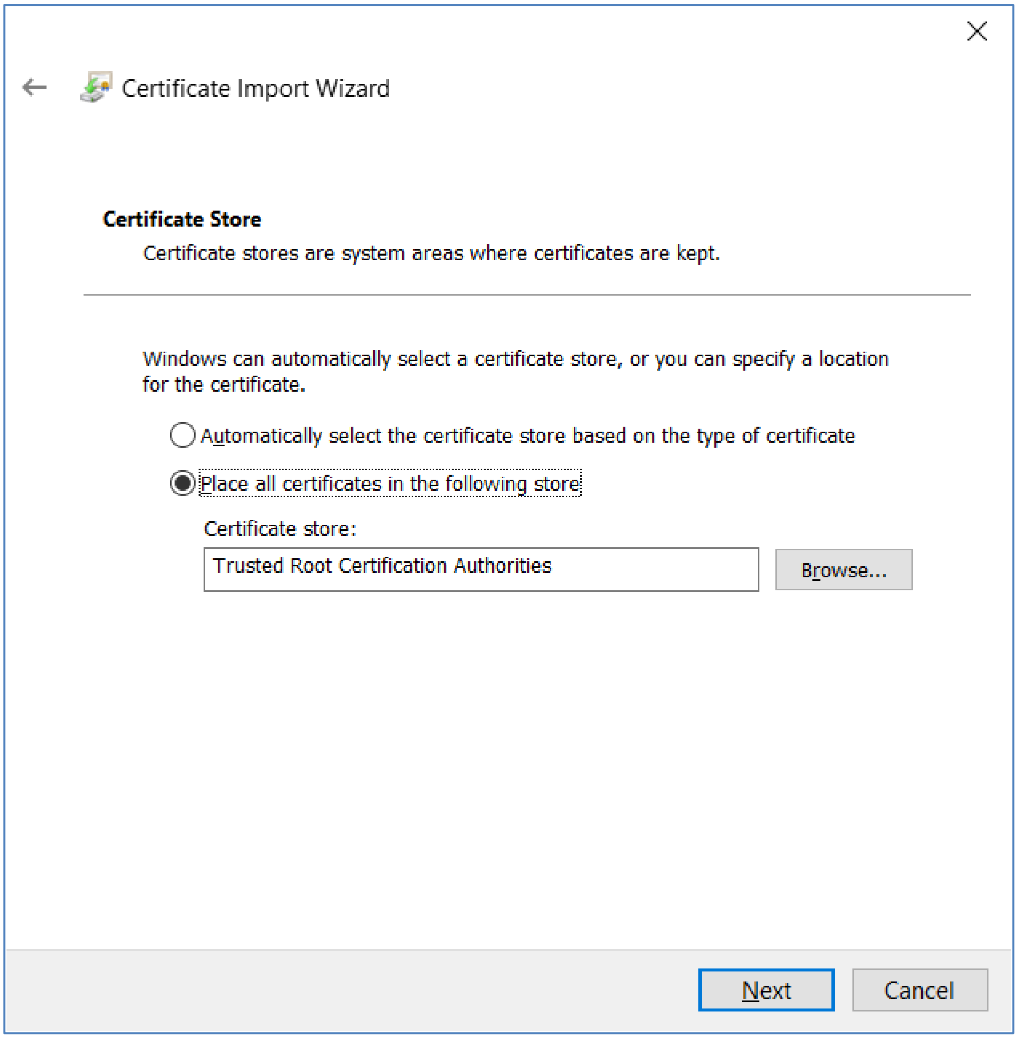

Certificate Store

Certificate stores are system areas where certificates are kept.

Windows can automatically select a certificate store, or you can specify a location for the certificate.

○ Automatically select the certificate store based on the type of certificate

● Place all certificates in the following store

Certificate store: Trusted Root Certification Authorities

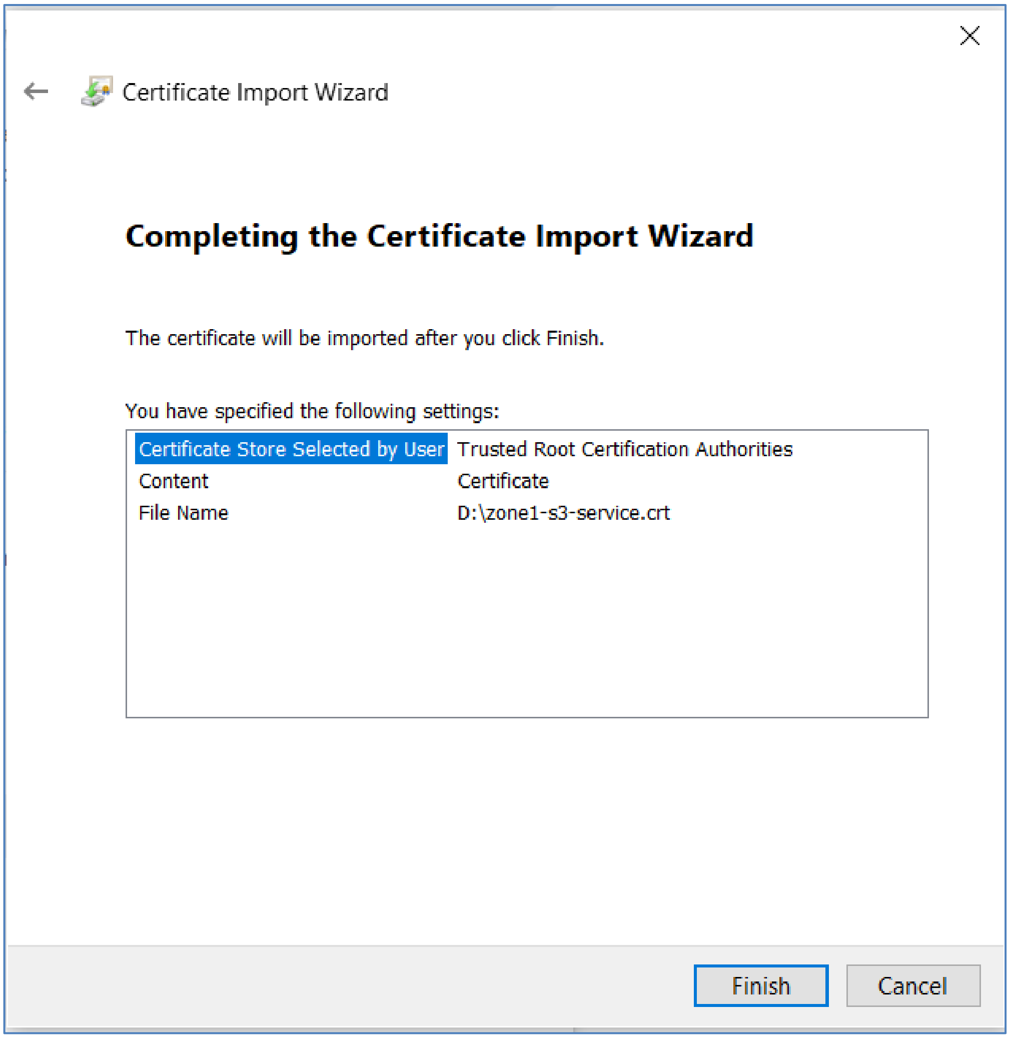

Completing the Certificate Import Wizard

The certificate will be imported after you click Finish.

You have specified the following settings:

Certificate Store Selected by User	Trusted Root Certification Authorities
Content	Certificate
File Name	D:\zone1-s3-service.crt

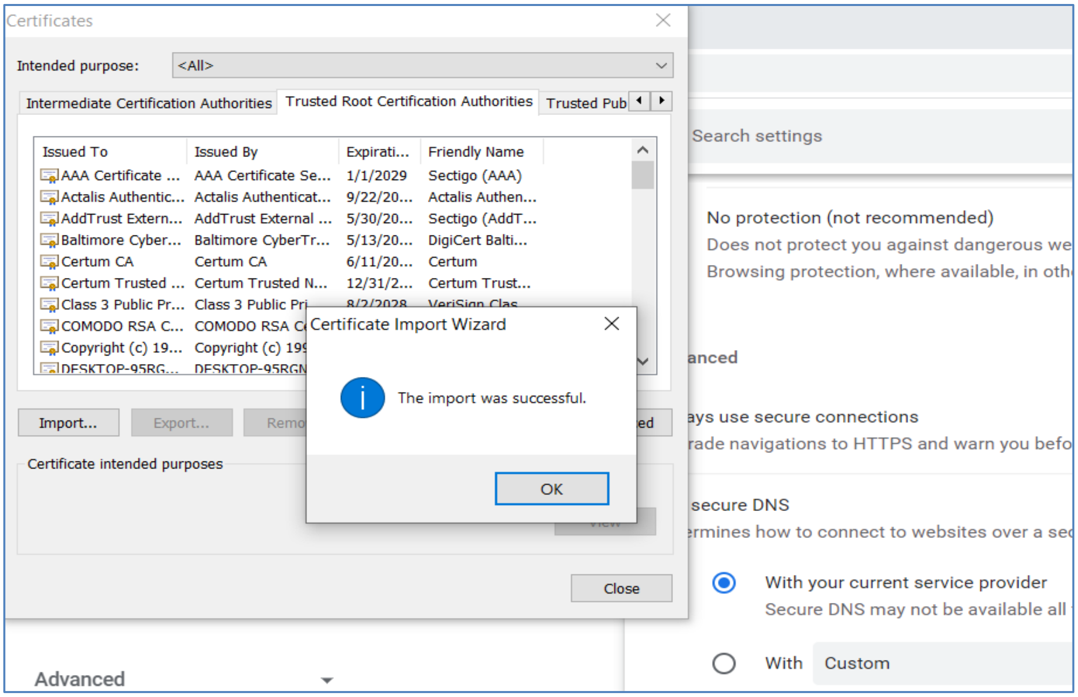
Certificates

Intended purpose: <All>

Intermediate Certification Authorities | Trusted Root Certification Authorities | Trusted Pub

Issued To	Issued By	Expirati...	Friendly Name
AAA Certificate ...	AAA Certificate Se...	1/1/2029	Sectigo (AAA)
Actalis Authentic...	Actalis Authenticat...	9/22/20...	Actalis Authen...
AddTrust Extern...	AddTrust External ...	5/30/20...	Sectigo (AddT...
Baltimore Cyber...	Baltimore CyberTr...	5/13/20...	DigiCert Balti...
Certum CA	Certum CA	6/11/20...	Certum
Certum Trusted ...	Certum Trusted N...	12/31/2...	Certum Trust...
Class 3 Public Pr...	Class 3 Public Pri	8/2/2028	VeriSign Clas
COMODO RSA C...	COMODO RSA C		
Copyright (c) 19...	Copyright (c) 199		
DESKTOP-95RG...	DESKTOP-95RGN		

Import... Export... Remo

Certificate intended purposes

Close

Certificate Import Wizard

The import was successful.

OK

Search settings

No protection (not recommended)

Does not protect you against dangerous we

Browsing protection, where available, in oth

anced

ays use secure connections

rade navigations to HTTPS and warn you befo

secure DNS

ermines how to connect to websites over a sec

● With your current service provider

Secure DNS may not be available all

○ With Custom

Advanced

4.1.3. Create S3 Browser user account

- We will create a S3 user account using “S3-User1” Access and Secret keys previously created in PetaSAN
- From S3 Browser select Accounts->Manage Accounts->Add Account

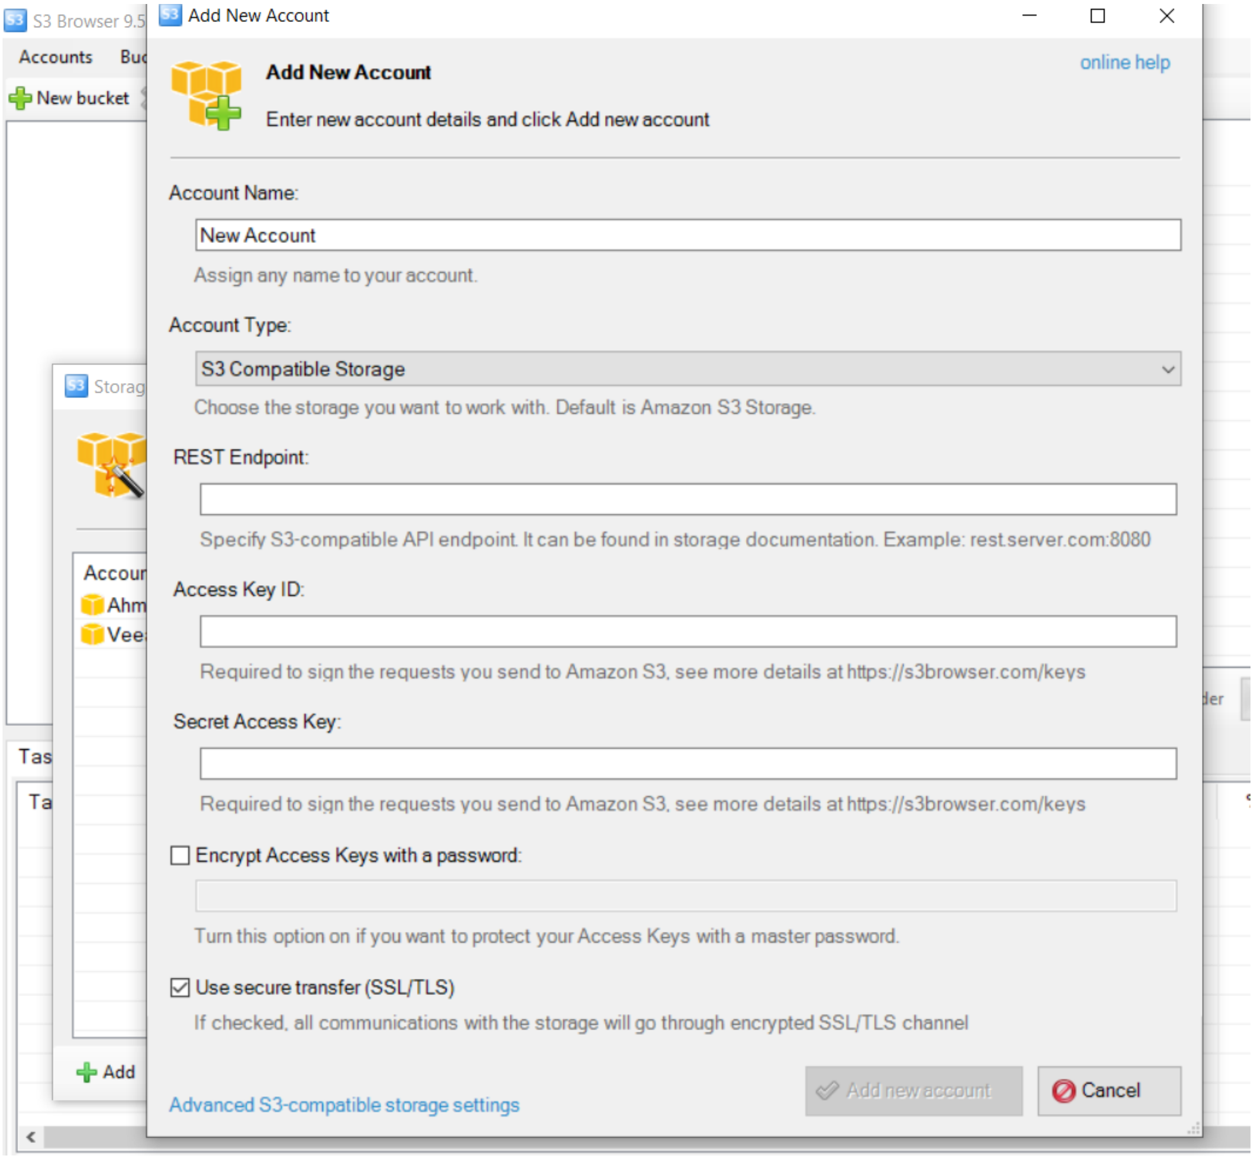

Add New Account

Enter new account details and click Add new account

Account Name:

New Account

Assign any name to your account.

Account Type:

S3 Compatible Storage

Choose the storage you want to work with. Default is Amazon S3 Storage.

REST Endpoint:

Specify S3-compatible API endpoint. It can be found in storage documentation. Example: rest.server.com:8080

Access Key ID:

Required to sign the requests you send to Amazon S3, see more details at https://s3browser.com/keys

Secret Access Key:

Required to sign the requests you send to Amazon S3, see more details at https://s3browser.com/keys

□ Encrypt Access Keys with a password:

Turn this option on if you want to protect your Access Keys with a master password.

☑ Use secure transfer (SSL/TLS)

If checked, all communications with the storage will go through encrypted SSL/TLS channel

- Enter the account name, select S3 Compatible Storage then enter the Endpoint (Service name: port number).
- Then enter the S3-User1 Access and Secret Keys.

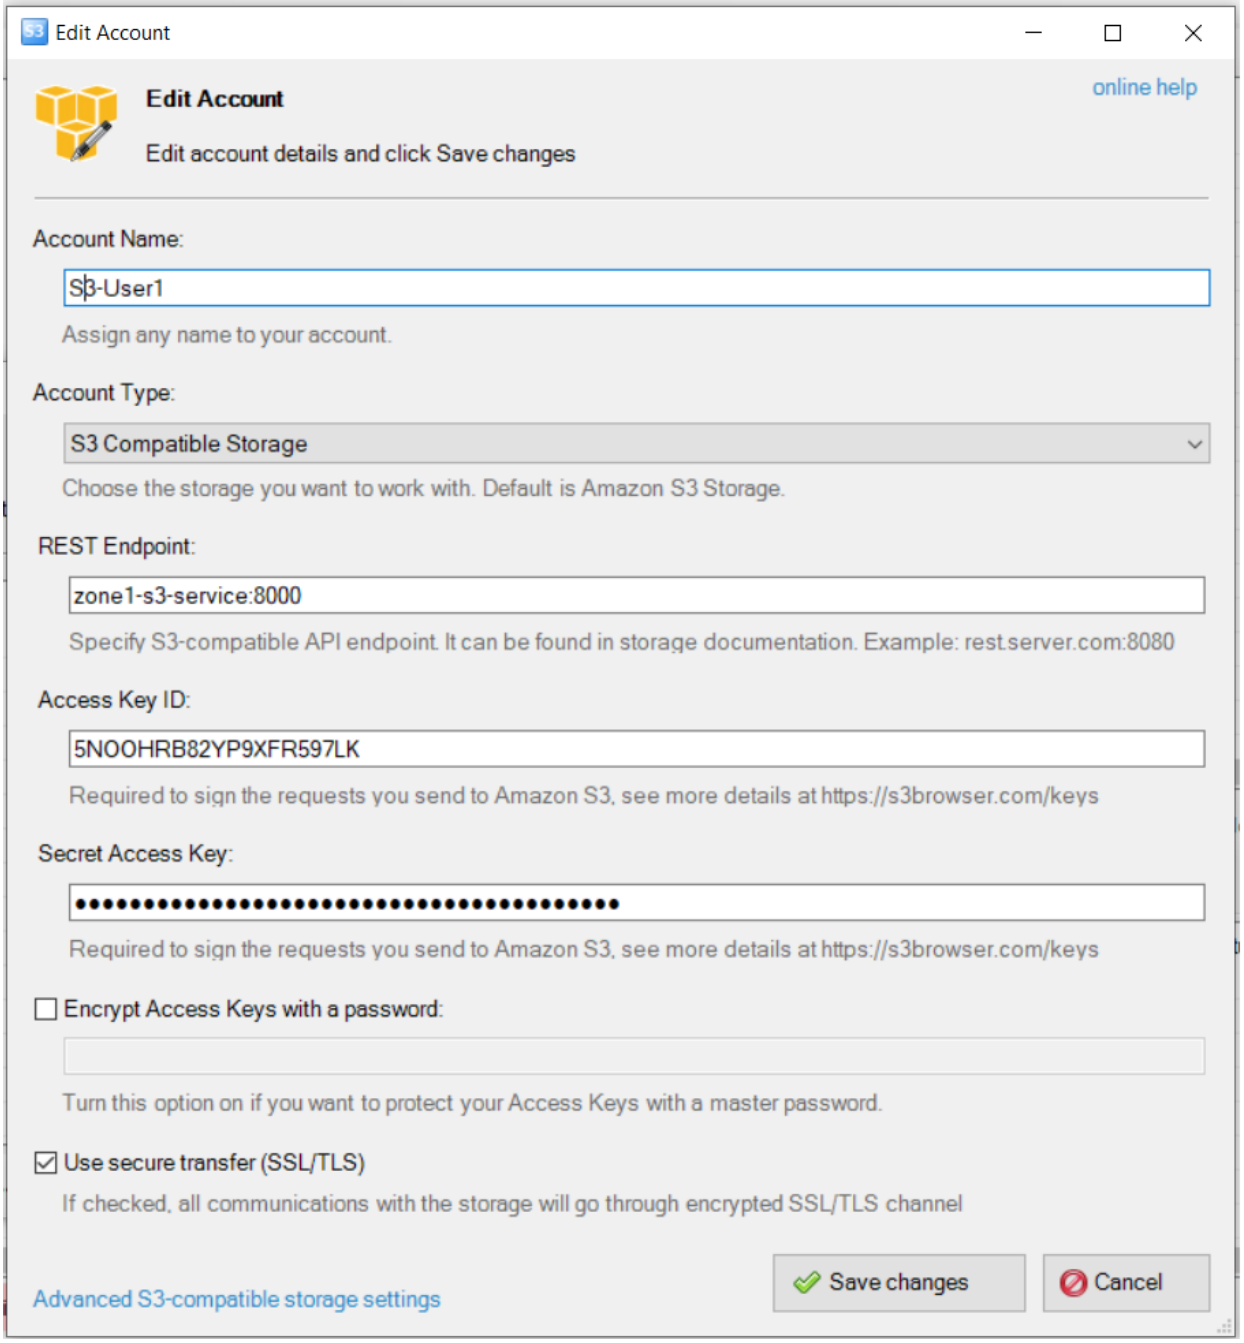

Edit Account

Edit Account

Edit account details and click Save changes

online help

Account Name:

S3-User1

Assign any name to your account.

Account Type:

S3 Compatible Storage

Choose the storage you want to work with. Default is Amazon S3 Storage.

REST Endpoint:

zone1-s3-service:8000

Specify S3-compatible API endpoint. It can be found in storage documentation. Example: rest.server.com:8080

Access Key ID:

5NOOHRB82YP9XFR597LK

Required to sign the requests you send to Amazon S3, see more details at https://s3browser.com/keys

Secret Access Key:

●●

Required to sign the requests you send to Amazon S3, see more details at https://s3browser.com/keys

□ Encrypt Access Keys with a password:

Turn this option on if you want to protect your Access Keys with a master password.

☑ Use secure transfer (SSL/TLS)

If checked, all communications with the storage will go through encrypted SSL/TLS channel

- After saving the account you will be able to view the existing buckets , create new buckets and upload your files.

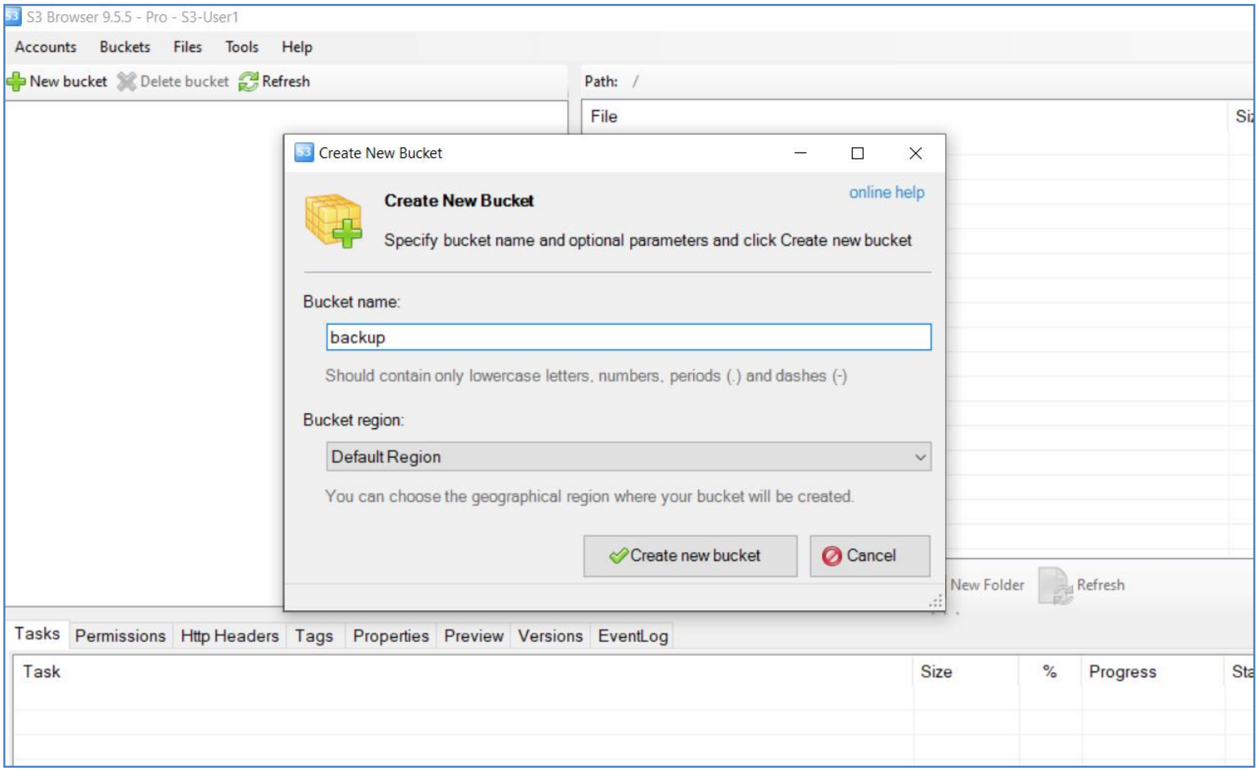

Create New Bucket

Specify bucket name and optional parameters and click Create new bucket

Bucket name: backup

Should contain only lowercase letters, numbers, periods (.) and dashes (-)

Bucket region: Default Region

You can choose the geographical region where your bucket will be created.

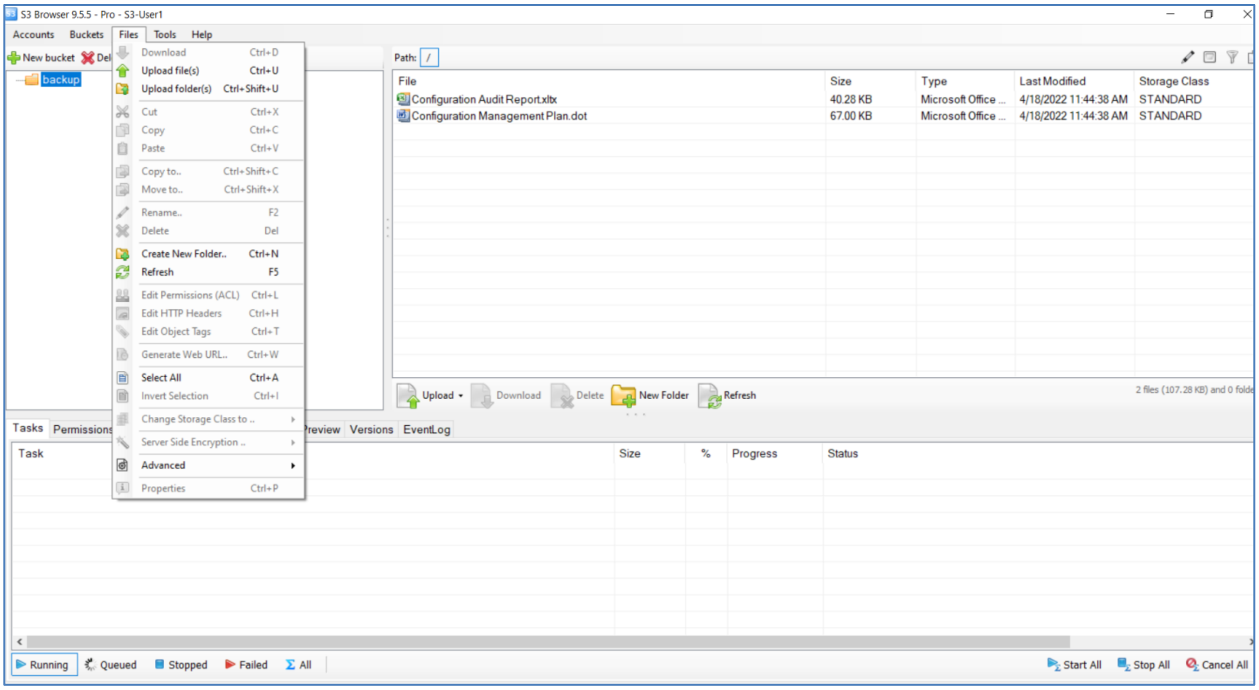

File	Size	Type	Last Modified	Storage Class
Configuration Audit Report.xltx	40.28 KB	Microsoft Office ...	4/18/2022 11:44:38 AM	STANDARD
Configuration Management Plan.dot	67.00 KB	Microsoft Office ...	4/18/2022 11:44:38 AM	STANDARD

2 files (107.28 KB) and 0 folde

- Data will be stored in the Backups placement target data pool

.

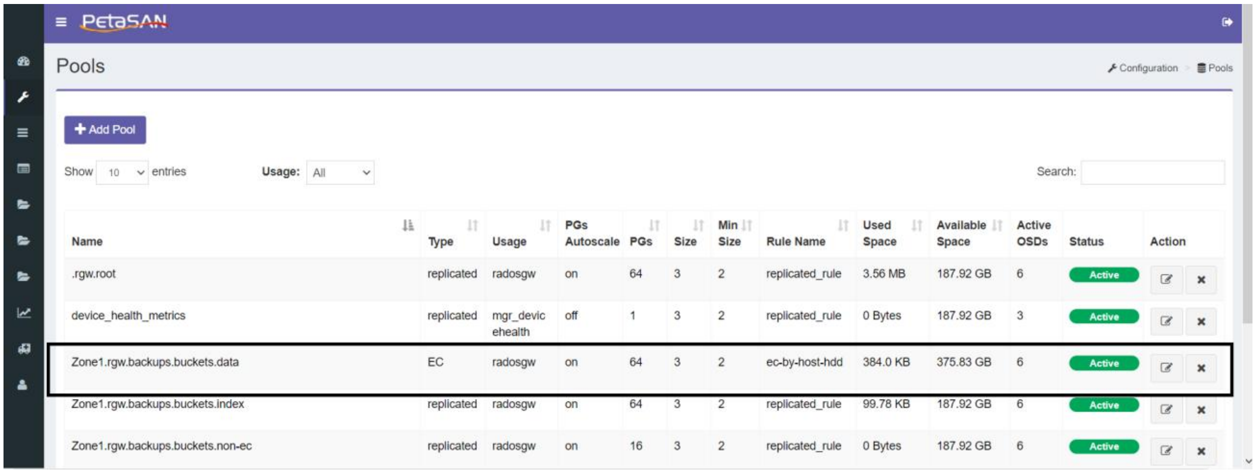

Name	Type	Usage	PGs Autoscale	PGs	Size	Min Size	Rule Name	Used Space	Available Space	Active OSDs	Status	Action
.rgw.root	replicated	radosgw	on	64	3	2	replicated_rule	3.56 MB	187.92 GB	6	Active	
device_health_metrics	replicated	mgr_devicehealth	off	1	3	2	replicated_rule	0 Bytes	187.92 GB	3	Active	
Zone1.rgw.backups.buckets.data	EC	radosgw	on	64	3	2	ec-by-host-hdd	384.0 KB	375.83 GB	6	Active	
Zone1.rgw.backups.buckets.index	replicated	radosgw	on	64	3	2	replicated_rule	99.78 KB	187.92 GB	6	Active	
Zone1.rgw.backups.buckets.non-ec	replicated	radosgw	on	16	3	2	replicated_rule	0 Bytes	187.92 GB	6	Active	

4.2. Cyberduck

4.2.1. Define certificate in hosts file

- Same as mentioned in step 4.1.1, if done before no need to repeat it

4.2.2. Import s3-service certificate

- Same as mentioned in step 4.1.2, if done before no need to repeat it

4.2.3. Connect using cyberduck

- Open new connection by entering the service name and port number ,S3-User1 Access and Secret Keys

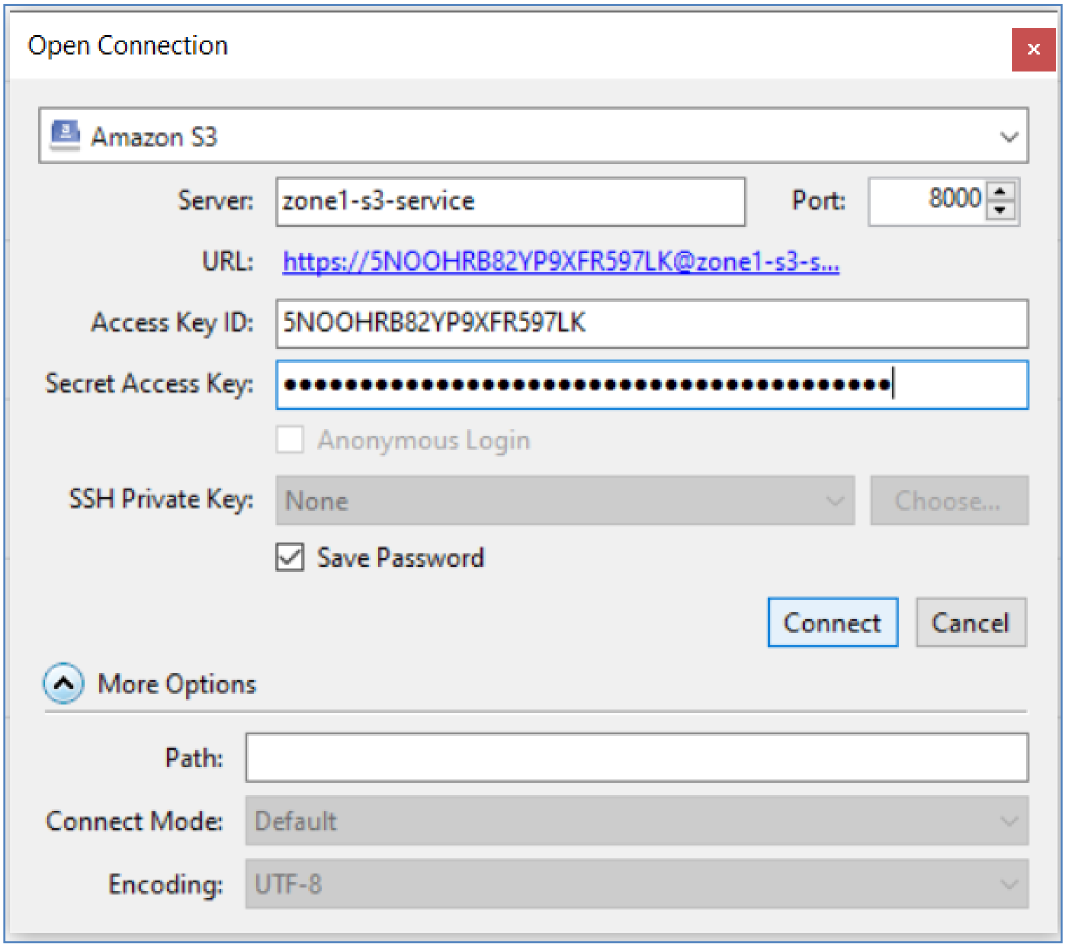

Open Connection

Amazon S3

Server: zone1-s3-service

Port: 8000

URL: https://5NOOHRB82YP9XFR597LK@zone1-s3-s...

Access Key ID: 5NOOHRB82YP9XFR597LK

Secret Access Key: ••

□ Anonymous Login

SSH Private Key: None

☑ Save Password

Path:

Connect Mode: Default

Encoding: UTF-8

- You will be able to view the user’s bucket list

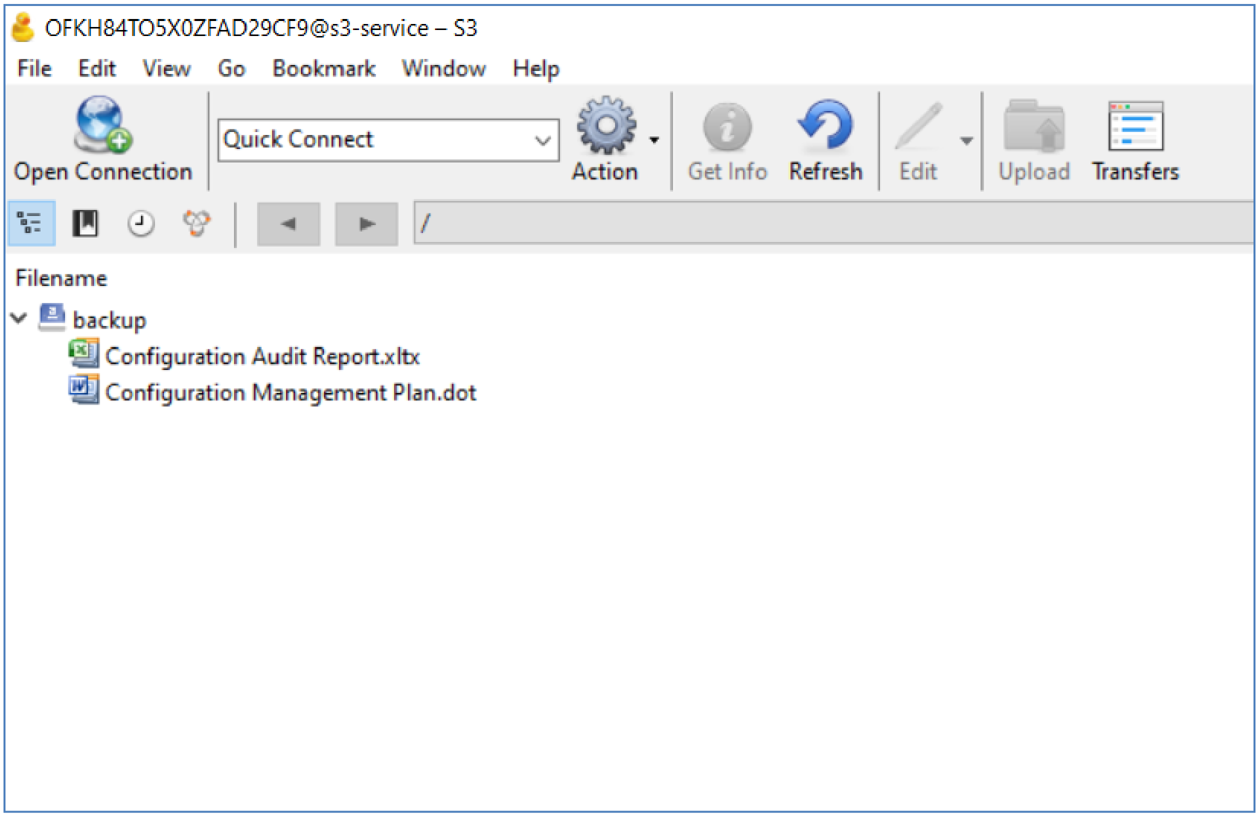

backup
Configuration Audit Report.xltx
Configuration Management Plan.dot

4.3. Amazon CLI Tool

4.3.1. Define certificate in hosts file

- Same as mentioned in step 4.1.1, if done before no need to repeat it

4.3.2. Import s3-service certificate

- setup the s3 service certificate using command line, example if the cert file placed on the D: drive

```
aws configure set default.ca_bundle "D:\zone1-s3-service.crt”
```

4.3.3. Configure the aws using configure command

- Configure AWS using command

```
aws configure
```

- Enter the S3-User1 Access and Secret Keys and enter the zonegroup name

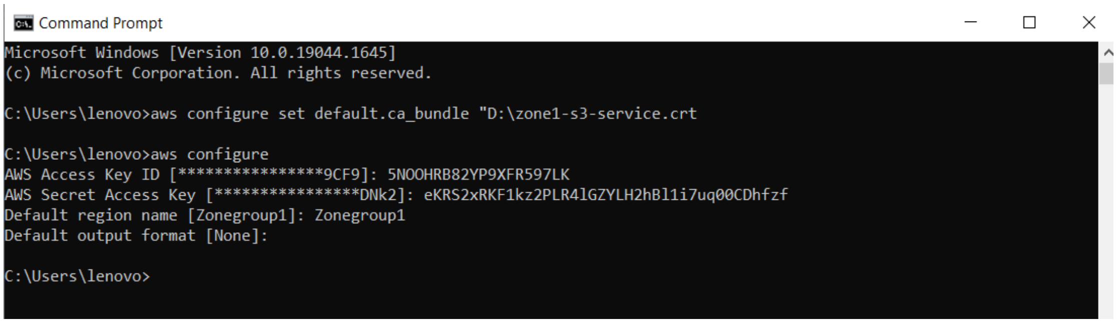

```
Microsoft Windows [Version 10.0.19044.1645]
(c) Microsoft Corporation. All rights reserved.

C:\Users\lenovo>aws configure set default.ca_bundle "D:\zone1-s3-service.crt

C:\Users\lenovo>aws configure
AWS Access Key ID [****************9CF9]: 5NOOHRB82YP9XFR597LK
AWS Secret Access Key [****************DNk2]: eKRS2xRKF1kz2PLR4lGZYLH2hBl1i7uq00CDhfzf
Default region name [Zonegroup1]: Zonegroup1
Default output format [None]:

C:\Users\lenovo>
```

4.3.4. Get bucket List

- You can get the bucket list using command
  aws s3 ls --endpoint-url https://zone1-s3-service:8000

4.3.5. Create new Bucket

- You can create new bucket named "bucket1" using command:
  aws s3api create-bucket --bucket bucket1 --endpoint-url https://zone1-s3-service:8000

4.3.6. Upload file

- You can upload a file named Notes in my desktop in bucket1 using the following command:
  aws s3 cp Desktop\Notes.docx s3://bucket1/ --endpoint-url https://zone1-s3-service:8000

4.3.7. List bucket content

- You can list the content in a specifc bucket using the following command:
  aws s3 ls s3://bucket1/ --endpoint-url https://zone1-s3-service:8000

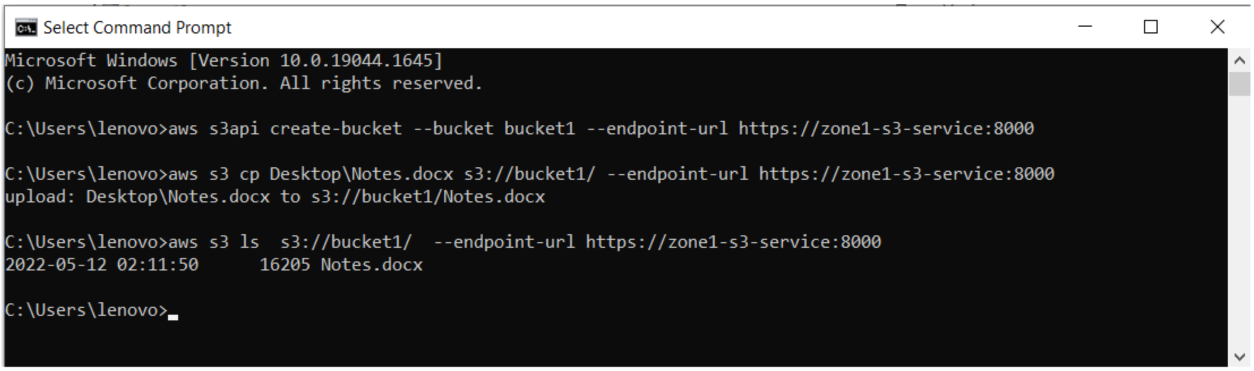

```
Microsoft Windows [Version 10.0.19044.1645]
(c) Microsoft Corporation. All rights reserved.

C:\Users\lenovo>aws s3api create-bucket --bucket bucket1 --endpoint-url https://zone1-s3-service:8000

C:\Users\lenovo>aws s3 cp Desktop\Notes.docx s3://bucket1/ --endpoint-url https://zone1-s3-service:8000
upload: Desktop\Notes.docx to s3://bucket1/Notes.docx

C:\Users\lenovo>aws s3 ls  s3://bucket1/  --endpoint-url https://zone1-s3-service:8000
2022-05-12 02:11:50      16205 Notes.docx

C:\Users\lenovo>_
```

5. Multi Site Installation

- You can setup a multi site by doing the following :

5.1. Configuring S3

5.1.1. S3 Settings

- Follow the same steps done for the first cluster.

5.1.2. Assign S3 Role to nodes

- Follow the same steps done for the first cluster.

5.1.3. Add End Point to the Master Zone

- Go to the first cluster, open the zones view list and select to edit the master zone, in this example it is Zone1.

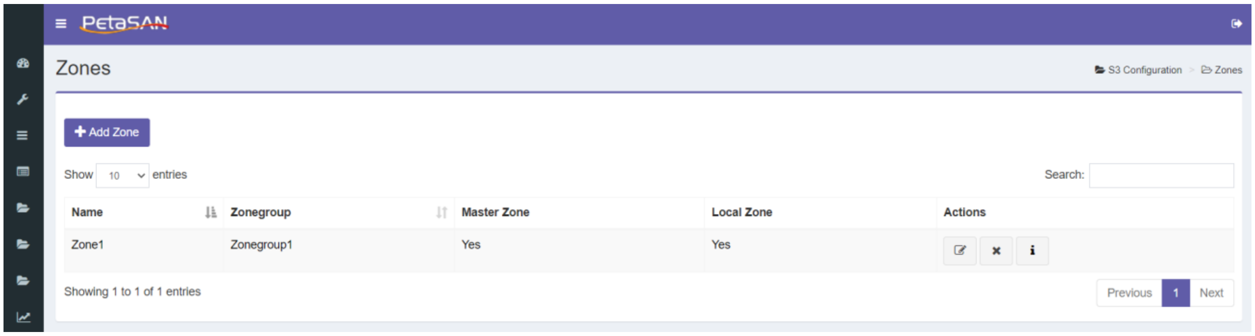

Name	Zonegroup	Master Zone	Local Zone	Actions
Zone1	Zonegroup1	Yes	Yes	

Showing 1 to 1 of 1 entries

- Enter the endpoint of the zone.

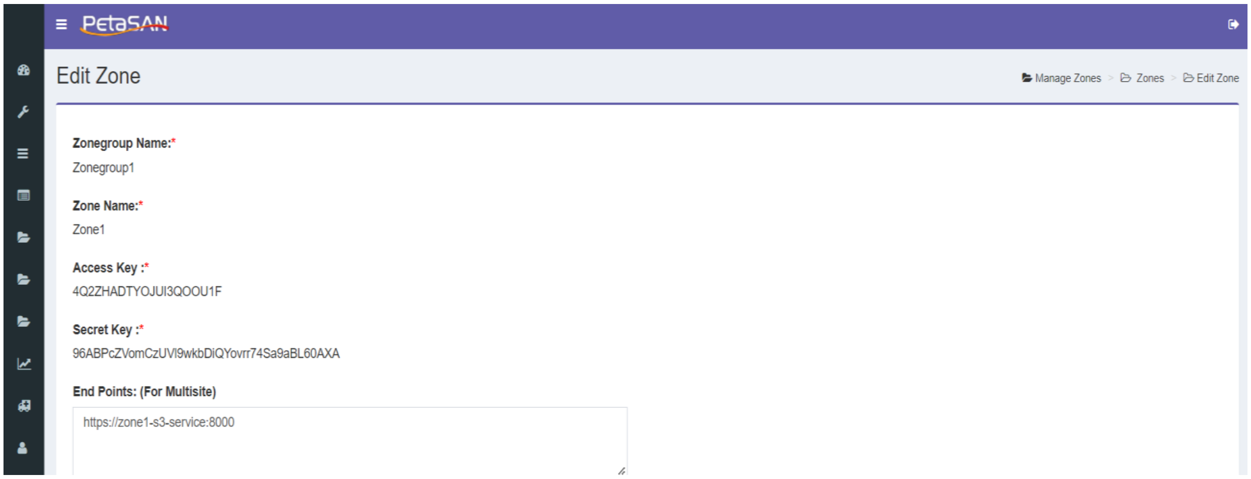

Zonegroup Name:*
Zonegroup1

Zone Name:*
Zone1

Access Key :*
4Q2ZHADTYOJUI3QOOU1F

Secret Key :*
96ABPcZVomCzUVl9wkbDiQYovrr74Sa9aBL60AXA

End Points: (For Multisite)
https://zone1-s3-service:8000

- Notice that the system has created user named “Synchronization-user”.

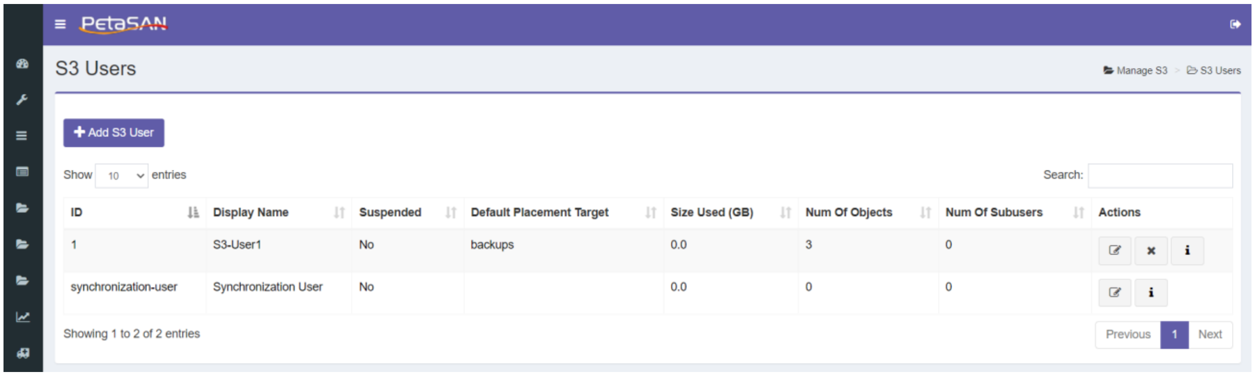

S3 Users

Manage S3 > S3 Users

Show 10 entries Search:

ID	Display Name	Suspended	Default Placement Target	Size Used (GB)	Num Of Objects	Num Of Subusers	Actions
1	S3-User1	No	backups	0.0	3	0	
synchronization-user	Synchronization User	No		0.0	0	0	

Showing 1 to 2 of 2 entries

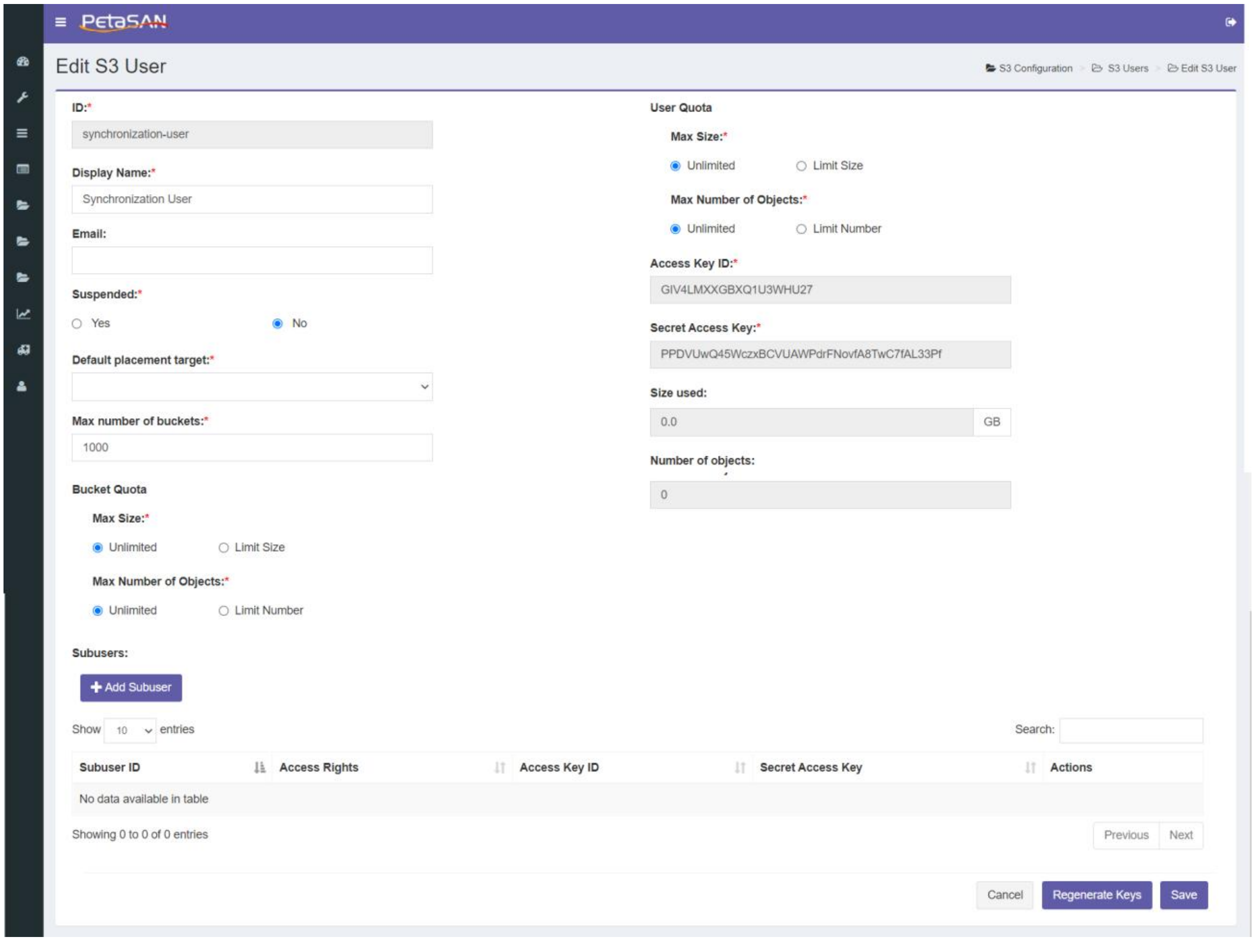

Edit S3 User

S3 Configuration > S3 Users > Edit S3 User

ID:* synchronization-user

Display Name:* Synchronization User

Email:

Suspended:* ○ Yes ● No

Default placement target:*

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:* GIV4LMXXGBXQ1U3WHU27

Secret Access Key:* PPDVUwQ45WczxBCVUAWPdrFNovfA8TwC7fAL33Pf

Size used: 0.0 GB

Number of objects: 0

Subusers:

Show 10 entries Search:

Subuser ID	Access Rights	Access Key ID	Secret Access Key	Actions
No data available in table				

Showing 0 to 0 of 0 entries

- You will use the “Synchronization-user” information in the Pull screen coming next.

5.1.4. Import Peer Certificate

- Import certificate of the second cluster in case of using self signed certificates.

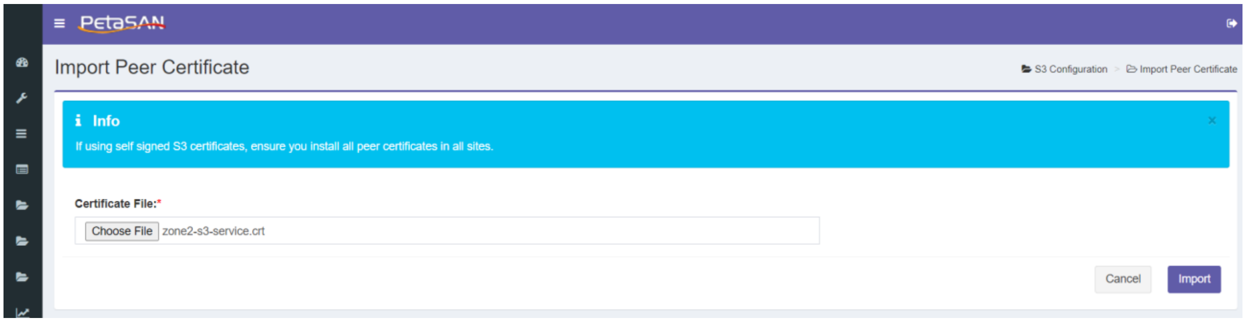

Import Peer Certificate

Info
If using self signed S3 certificates, ensure you install all peer certificates in all sites.

Certificate File:* zone2-s3-service.crt

- Your peer certificate has been imported successfully

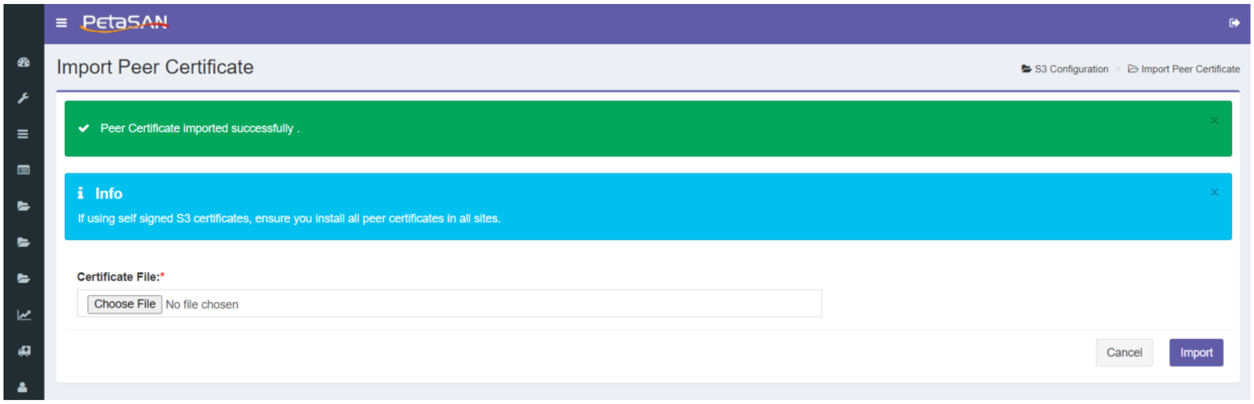

Import Peer Certificate

✓ Peer Certificate imported successfully .

Info
If using self signed S3 certificates, ensure you install all peer certificates in all sites.

Certificate File:* No file chosen

5.1.5. Define the service names in hosts files

In this example we need each zone to be able to access the other zone for data replication and configuration.

Update the hosts files on the first and second zones, PetaSAN always syncs the hosts file in consul server, so all nodes in cluster gets the same copy of the file. We need to do the following steps on both zones to correctly setup the hosts file:

Zone1

- stop auto sync service

  ```
  systemctl stop petasan-file-sync
  ```

- Update the hosts file in node1 of the first cluster
  Connect to node1 using WinSCP tool and go to path /etc/hosts or use command use winscp or nano

  ```
  nano /etc/hosts
  ```

  Edit the hosts file to add entry 10.0.3.90 zone2-s3-service which is the S3 service ip of zone2

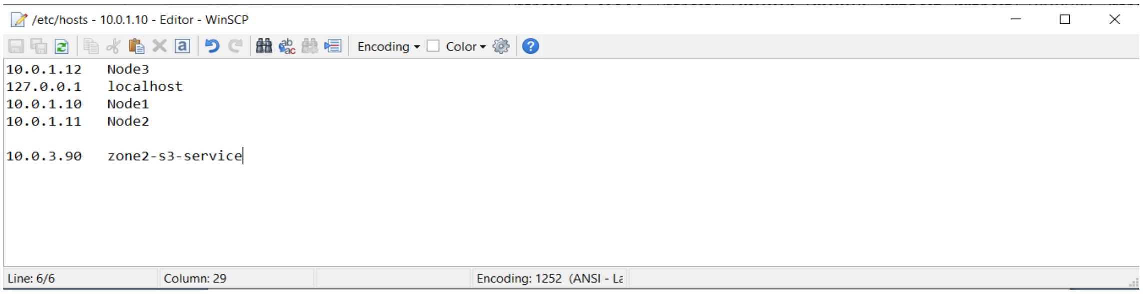
/etc/hosts - 10.0.1.10 - Editor - WinSCP
Encoding
Color

```
10.0.1.12    Node3
127.0.0.1    localhost
10.0.1.10    Node1
10.0.1.11    Node2

10.0.3.90    zone2-s3-service
```

Line: 6/6
Column: 29
Encoding: 1252 (ANSI - La

- sync the hosts file to consul
  /opt/petasan/scripts/util/sync_file.py /etc/hosts

- Restart the sync service on current node
  systemctl start petasan-file-sync

This will sync the updated hosts file to all nodes

Zone2

- stop auto sync service
  systemctl stop petasan-file-sync

- Update the hosts file in the second cluster node
  Connect to node90 using WinSCP tool and go to path /etc/hosts or use command
  use winscp or nano /etc/hosts

  Edit the hosts file to add entry 10.0.3.10 zone1-s3-service which is the S3 service ip of zone1

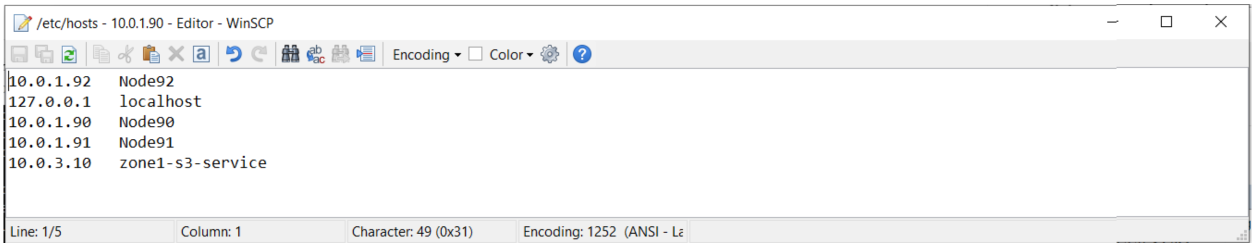
/etc/hosts - 10.0.1.90 - Editor - WinSCP
Encoding
Color

```
10.0.1.92    Node92
127.0.0.1    localhost
10.0.1.90    Node90
10.0.1.91    Node91
10.0.3.10    zone1-s3-service
```

Line: 1/5
Column: 1
Character: 49 (0x31)
Encoding: 1252 (ANSI - La

- sync the hosts file to consul
  /opt/petasan/scripts/util/sync_file.py /etc/hosts

- Restart the sync service on current node
  systemctl start petasan-file-sync

This will sync the updated hosts file to all nodes

- Add a temporary ip in the same zone1 network so we can access it to pull the zone information

```
ifconfig eth2 10.0.3.100 netmask 255.255.255.0
```

- Now you will be able to ping the first zone ip

```
ping 10.0.3.10
```

5.1.6. Import Peer Certificate

- For self signed certificates, import certificate of the first cluster.

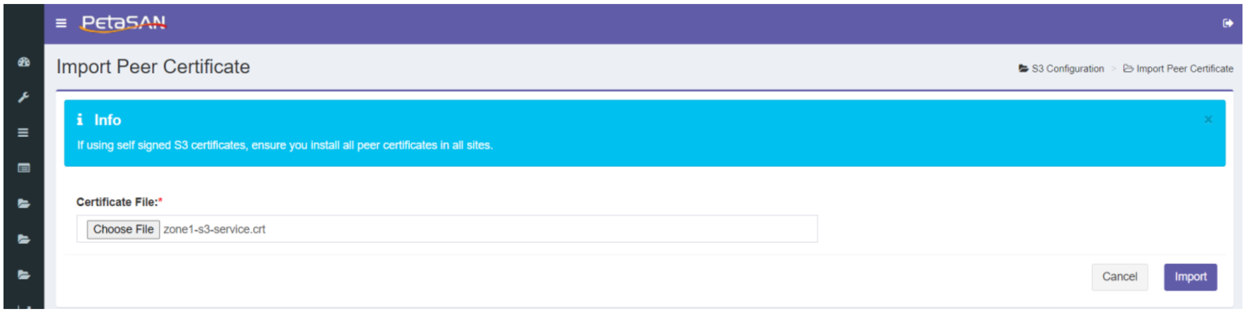

Import Peer Certificate

Info

If using self signed S3 certificates, ensure you install all peer certificates in all sites.

Certificate File:* zone1-s3-service.crt

- Your peer certificate has been imported successfully

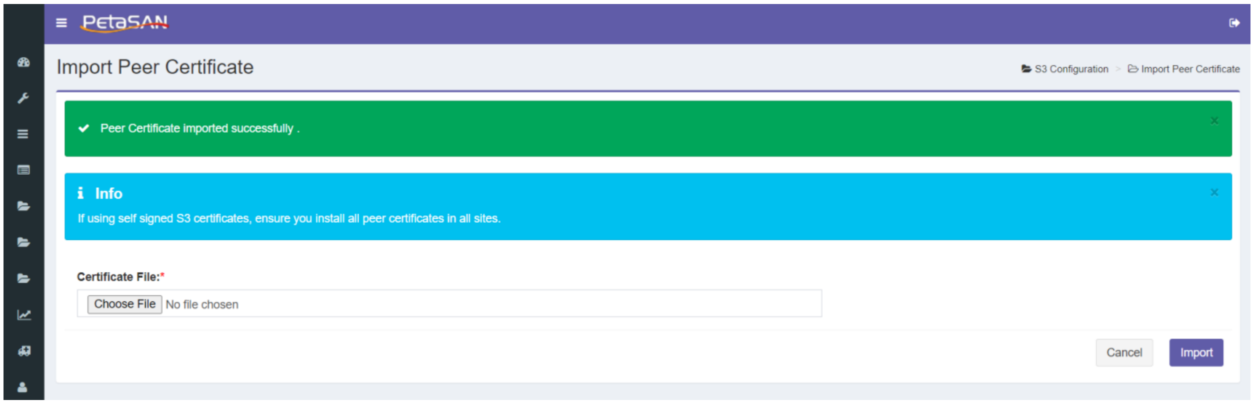

Import Peer Certificate

✔ Peer Certificate imported successfully .

Info

If using self signed S3 certificates, ensure you install all peer certificates in all sites.

Certificate File:* No file chosen

5.1.7. Pull First Cluster S3 Configuration

- From the menu select Configuration /S3 configuration/Zonegroups
- Pull the settings of the master zone by using its endpoint and Synchronization-user access and secret keys.

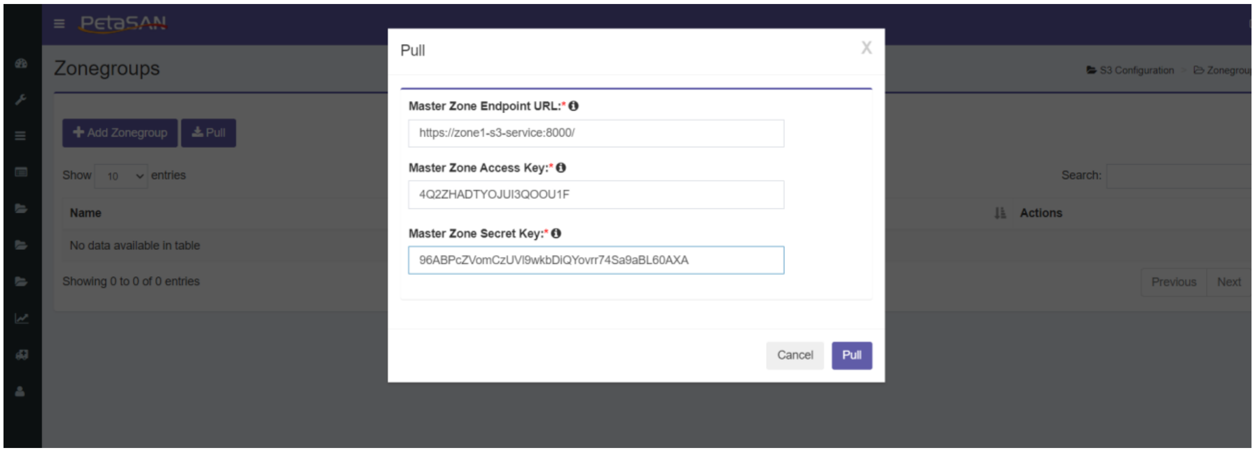

Pull

Master Zone Endpoint URL:* https://zone1-s3-service:8000/

Master Zone Access Key:* 4Q2ZHADTYOJUI3QOOU1F

Master Zone Secret Key:* 96ABPcZVomCzUVI9wkbDiQYovrr74Sa9aBL60AXA

Show 10 entries

Name	Actions
No data available in table	

Showing 0 to 0 of 0 entries

- The multisite zone configuration is pulled successfully and you can view the zonegroup information.

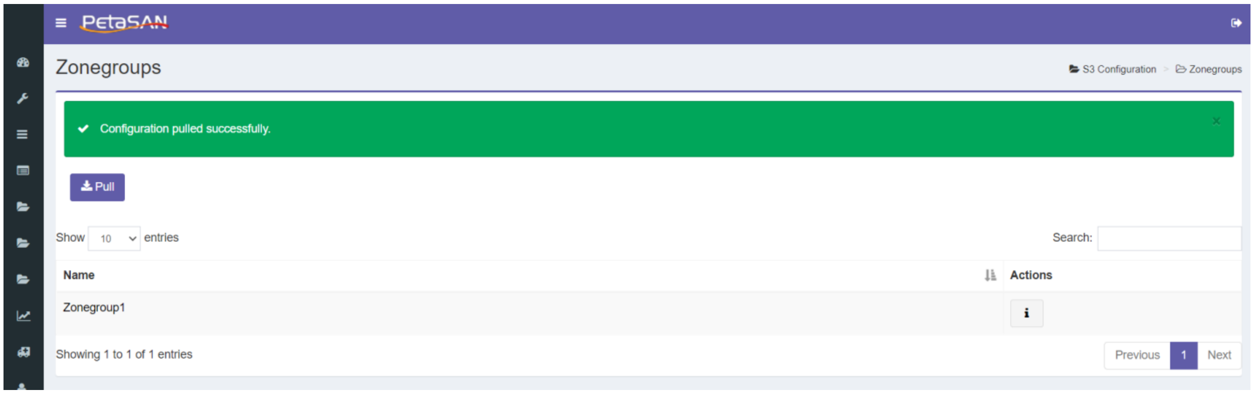

Configuration pulled successfully.

Show 10 entries

Name	Actions
Zonegroup1	i

Showing 1 to 1 of 1 entries

5.1.8. Add local zone

- You should now add a local zone to the second cluster

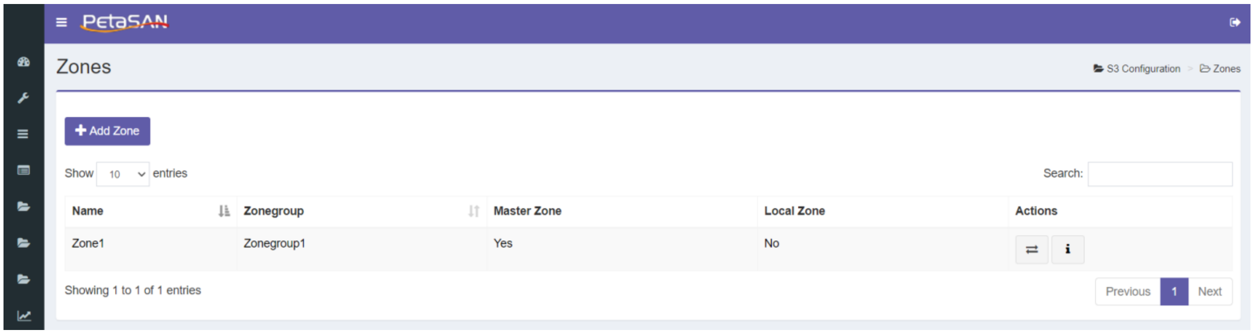

Show 10 entries

Name	Zonegroup	Master Zone	Local Zone	Actions
Zone1	Zonegroup1	Yes	No	⇄ i

Showing 1 to 1 of 1 entries

- Enter the zone name ,main pools and placement targets pools

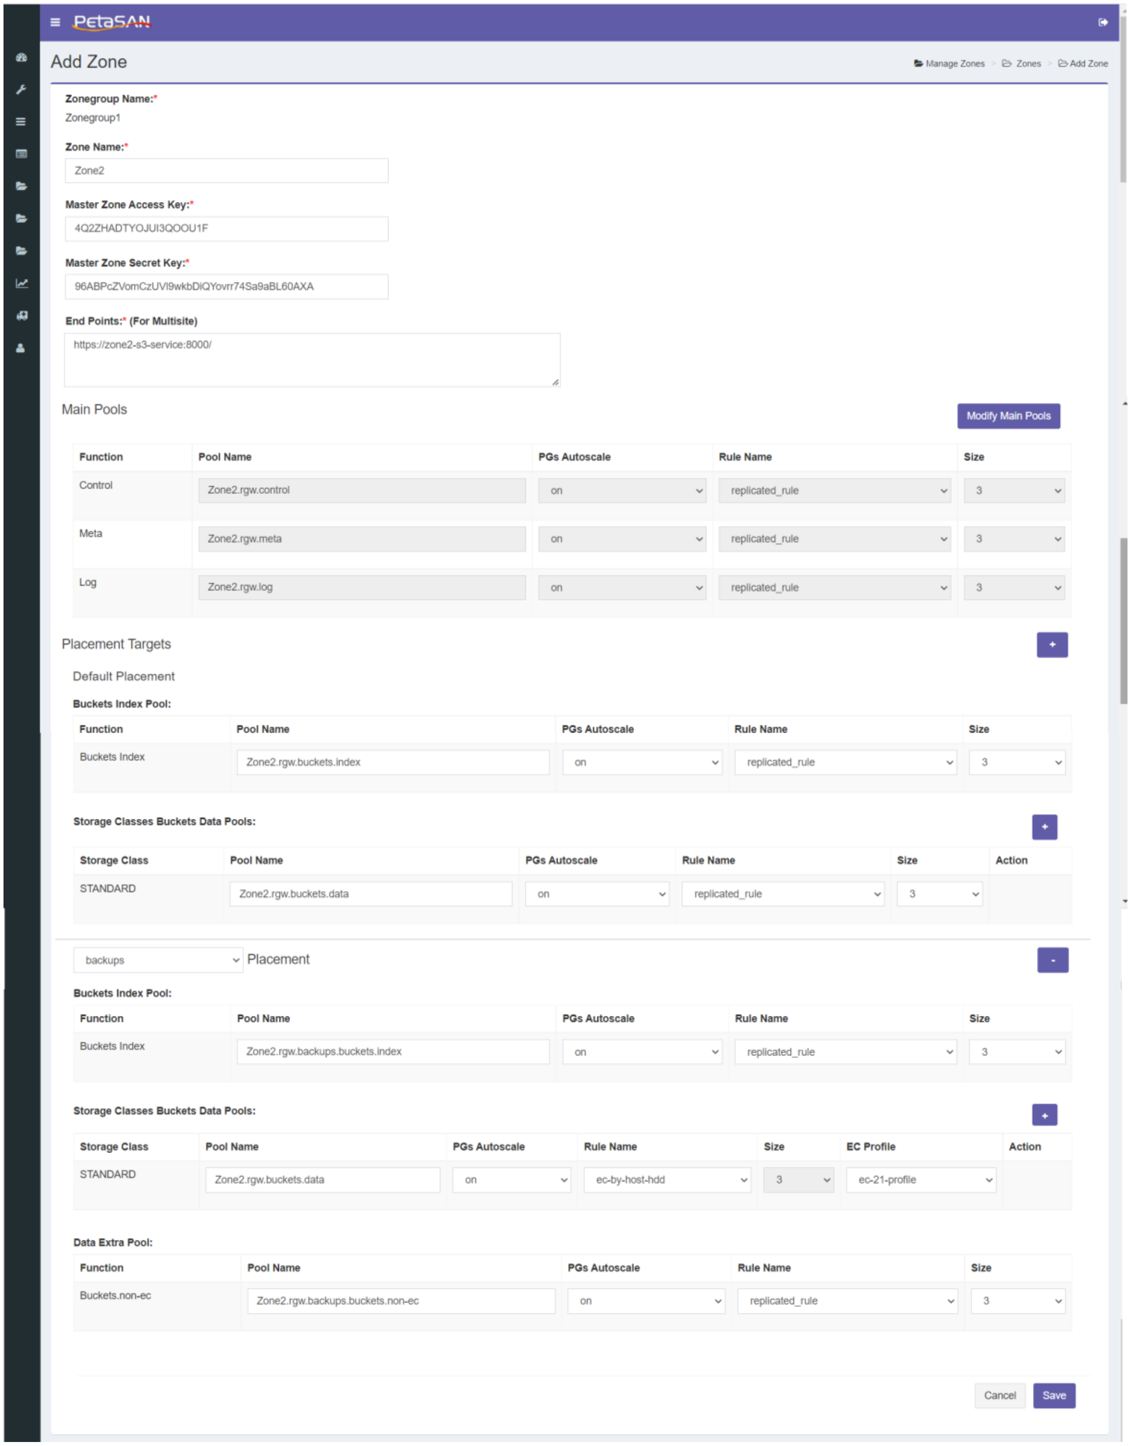

Add Zone

Zonegroup Name:*
Zonegroup1

Zone Name:*
Zone2

Master Zone Access Key:*
4Q2ZHADTYOJUI3QOOU1F

Master Zone Secret Key:*
96ABPcZVomCzUVl9wkbDiQYovrr74Sa9aBL60AXA

End Points:* (For Multisite)
https://zone2-s3-service:8000/

Main Pools

Function	Pool Name	PGs Autoscale	Rule Name	Size
Control	Zone2.rgw.control	on	replicated_rule	3
Meta	Zone2.rgw.meta	on	replicated_rule	3
Log	Zone2.rgw.log	on	replicated_rule	3

Placement Targets

Default Placement

Buckets Index Pool:

Function	Pool Name	PGs Autoscale	Rule Name	Size
Buckets Index	Zone2.rgw.buckets.index	on	replicated_rule	3

Storage Classes Buckets Data Pools:

Storage Class	Pool Name	PGs Autoscale	Rule Name	Size	Action
STANDARD	Zone2.rgw.buckets.data	on	replicated_rule	3	

backups Placement

Buckets Index Pool:

Function	Pool Name	PGs Autoscale	Rule Name	Size
Buckets Index	Zone2.rgw.backups.buckets.index	on	replicated_rule	3

Storage Classes Buckets Data Pools:

Storage Class	Pool Name	PGs Autoscale	Rule Name	Size	EC Profile	Action
STANDARD	Zone2.rgw.buckets.data	on	ec-by-host-hdd	3	ec-21-profile	

Data Extra Pool:

Function	Pool Name	PGs Autoscale	Rule Name	Size
Buckets.non-ec	Zone2.rgw.backups.buckets.non-ec	on	replicated_rule	3

- Zone2 is created successfully as shown

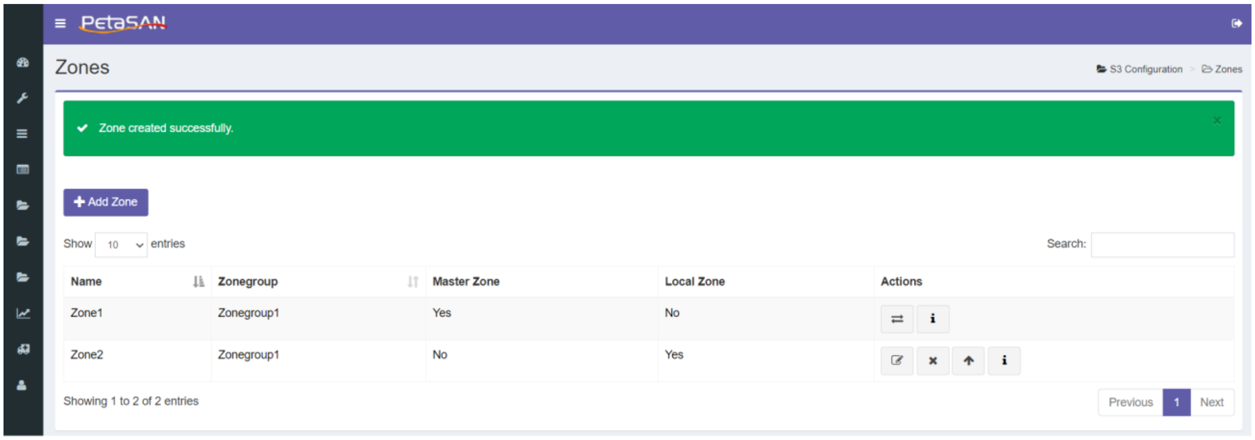

Zones

S3 Configuration > Zones

✔ Zone created successfully.

+ Add Zone

Show 10 entries Search:

Name	Zonegroup	Master Zone	Local Zone	Actions
Zone1	Zonegroup1	Yes	No	⇄ i
Zone2	Zonegroup1	No	Yes	✎ ✖ ↑ i

Showing 1 to 2 of 2 entries

Previous 1 Next

Repeat the client connectivity steps, installing the second zone certificate.

Now users can use any of the clusters to upload their data and data will be replicated automatically to the other cluster. Data can be written to both clusters in active/active manner.

6. Add S3 User

You can’t add, update or delete users from the secondary cluster, users must be maintained in the master cluster and they will be synched automatically to the second cluster.

In this example we added S3-User2 in the master cluster and it has been synched to the secondary cluster

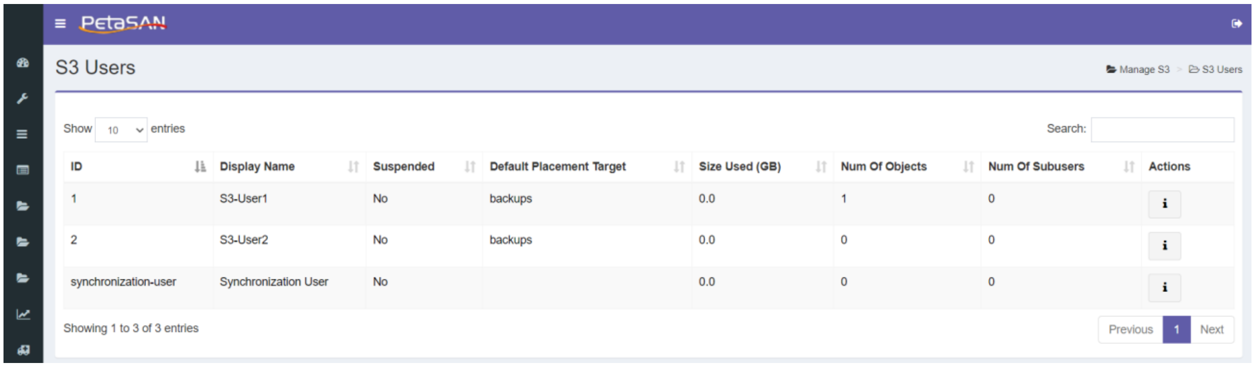

S3 Users

Manage S3 > S3 Users

Show 10 entries Search:

ID	Display Name	Suspended	Default Placement Target	Size Used (GB)	Num Of Objects	Num Of Subusers	Actions
1	S3-User1	No	backups	0.0	1	0	i
2	S3-User2	No	backups	0.0	0	0	i
synchronization-user	Synchronization User	No		0.0	0	0	i

Showing 1 to 3 of 3 entries

Previous 1 Next

7. Client Connectivity

7.1. S3 Browser

7.1.1. Define certificate in hosts file

Same as the first cluster section 4.1.1 but add the IP of the second zone service

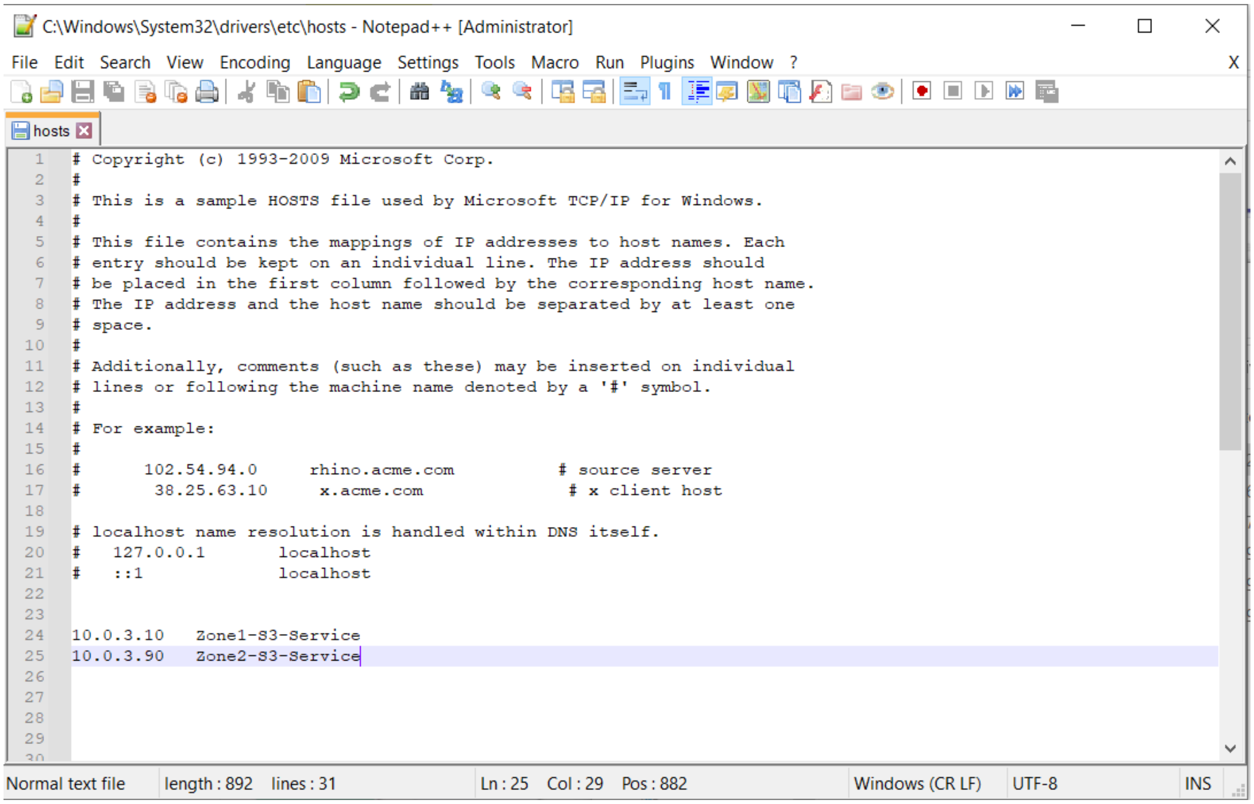
C:\Windows\System32\drivers\etc\hosts - Notepad++ [Administrator]
File Edit Search View Encoding Language Settings Tools Macro Run Plugins Window ?
hosts

```
# Copyright (c) 1993-2009 Microsoft Corp.
#
# This is a sample HOSTS file used by Microsoft TCP/IP for Windows.
#
# This file contains the mappings of IP addresses to host names. Each
# entry should be kept on an individual line. The IP address should
# be placed in the first column followed by the corresponding host name.
# The IP address and the host name should be separated by at least one
# space.
#
# Additionally, comments (such as these) may be inserted on individual
# lines or following the machine name denoted by a '#' symbol.
#
# For example:
#
#      102.54.94.0     rhino.acme.com          # source server
#       38.25.63.10     x.acme.com              # x client host

# localhost name resolution is handled within DNS itself.
#   127.0.0.1       localhost
#   ::1             localhost


10.0.3.10   Zone1-S3-Service
10.0.3.90   Zone2-S3-Service
```

7.1.2. Import s3-service certificate

- Import the second cluster certificate following the same steps in section 4.1.2 ,in this example we will import certificate zone2-s3-service.crt.

7.1.3. Create S3 Browser user account

- We will need to create an S3 browser account using the same S3User1 access key and secret keys but connecting to zone2-s3-service.

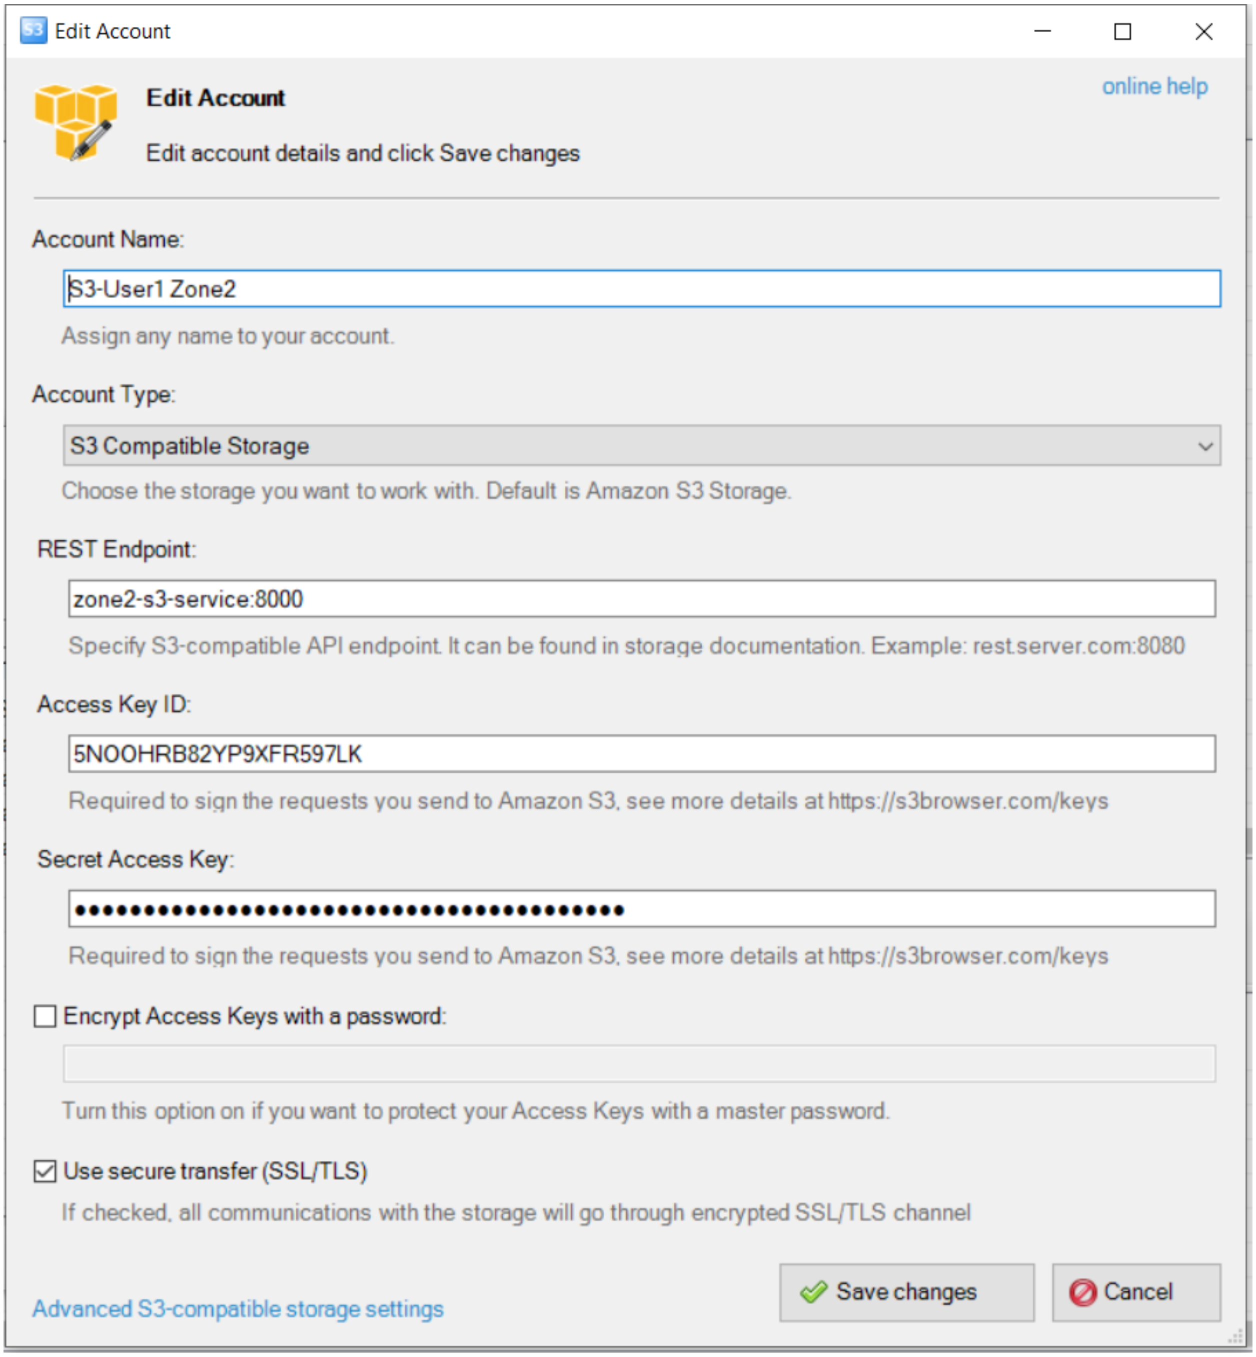

Edit Account

Edit Account

Edit account details and click Save changes

online help

Account Name:

S3-User1 Zone2

Assign any name to your account.

Account Type:

S3 Compatible Storage

Choose the storage you want to work with. Default is Amazon S3 Storage.

REST Endpoint:

zone2-s3-service:8000

Specify S3-compatible API endpoint. It can be found in storage documentation. Example: rest.server.com:8080

Access Key ID:

5NOOHRB82YP9XFR597LK

Required to sign the requests you send to Amazon S3, see more details at https://s3browser.com/keys

Secret Access Key:

••

Required to sign the requests you send to Amazon S3, see more details at https://s3browser.com/keys

☐ Encrypt Access Keys with a password:

Turn this option on if you want to protect your Access Keys with a master password.

☑ Use secure transfer (SSL/TLS)

If checked, all communications with the storage will go through encrypted SSL/TLS channel

Advanced S3-compatible storage settings

- You should see the same buckets as you have when you were connecting to zone1

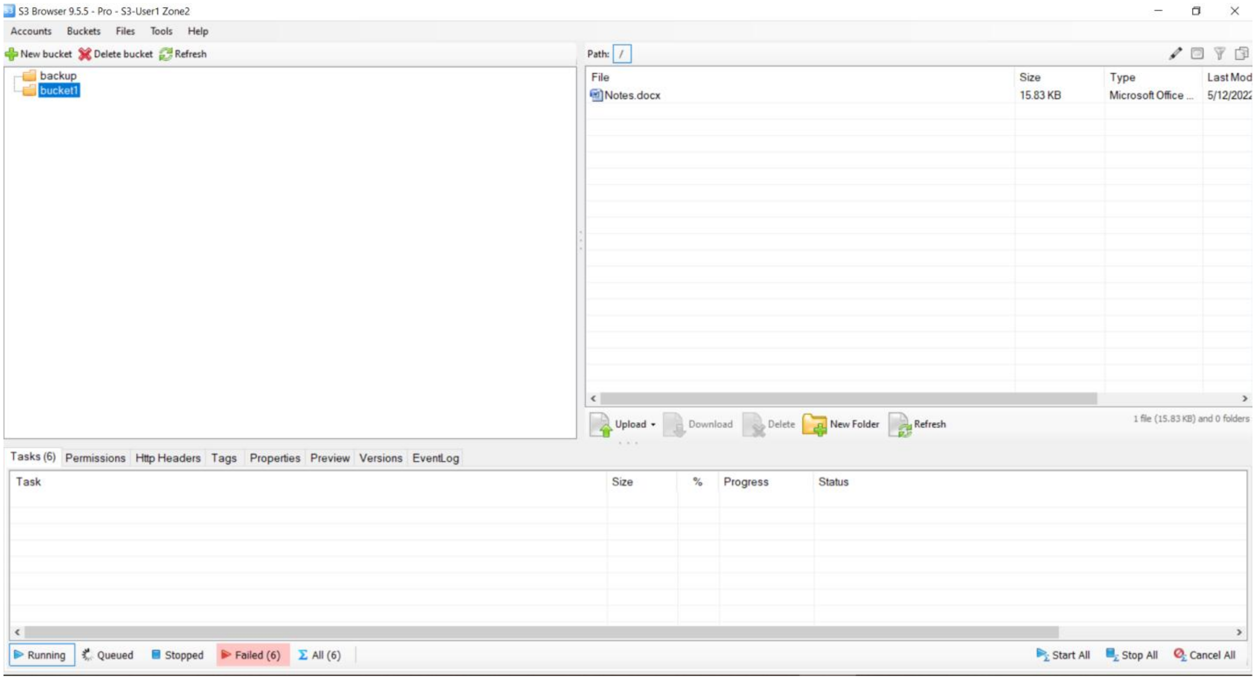

File	Size	Type	Last Mod
Notes.docx	15.83 KB	Microsoft Office ...	5/12/2022

7.2. Cyberduck

Same as done with the first cluster.

7.3. Amazon CLI Tool

Same as done with the first cluster.

8. Promote Zone

- In case the master zone (Currently in this example zone1) is down ,you can promote a non-master zone (Currently in this example zone2) to be a master zone by using the promote button in the zones view list.
- Make sure that all metadata (like zonegroup , zone or user updates) are synced before the promotion.

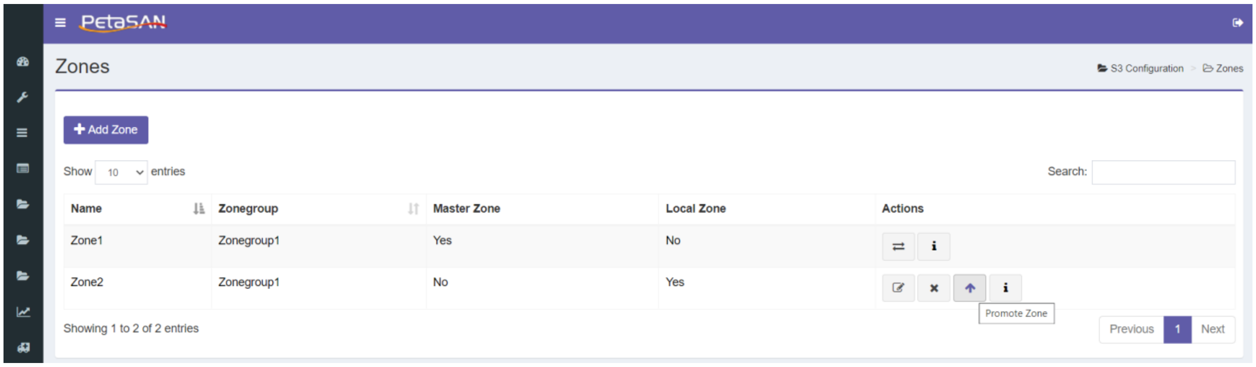

Zones

Name	Zonegroup	Master Zone	Local Zone	Actions
Zone1	Zonegroup1	Yes	No	
Zone2	Zonegroup1	No	Yes	

Showing 1 to 2 of 2 entries

If you promoted the zone2 to be the master zone then you will be able to do all the functionally of the master zone like adding S3 users.

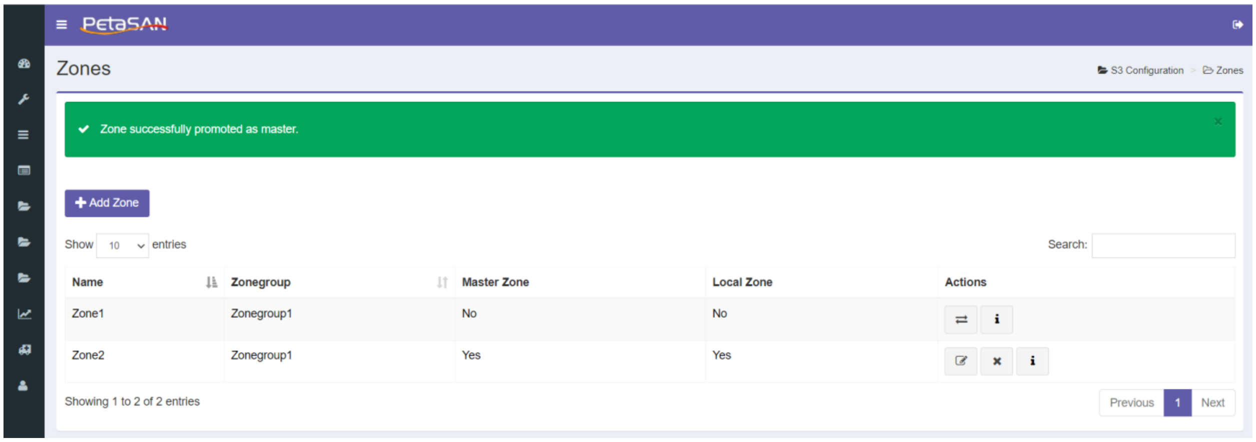

Zones

Zone successfully promoted as master.

Name	Zonegroup	Master Zone	Local Zone	Actions
Zone1	Zonegroup1	No	No	
Zone2	Zonegroup1	Yes	Yes	

Showing 1 to 2 of 2 entries