



Using PetaSAN S3 as a Veeam Backup Repository

Version 2.0

Revision History

Date	Version	Description
26/5/2022	1.0	Initial version
02/11/2023	2.0	Adding steps using Veeam 12



Contents

Revision History 1

1. Purpose 3

2. Pre-requisites 3

3. PetaSAN Configuration 3

 S3 Settings..... 3

 Add S3 user 5

4. Veeam Configuration 7

 Here we show how to create each repository type:..... 8

 Create Object Storage Repository 8

 Create Direct Attached Storage Repository..... 8

 Scale-out repository..... 8

5. Create a backup job 42



1. Purpose

The purpose of this document is to describe how PetaSAN S3 Object storage can be used as a Veeam Backup repository.

Currently Veeam supports using Direct-to-object storage to S3 Object storage (Veeam version 12)

While previously Veeam supported using S3 Object storage as a capacity tier (Veeam version 11)

2. Pre-requisites

This document assumes the user has the following components already installed:

- Veeam Backup and Replication version 12 with one Veeam server installed.
- PetaSAN Release 3.0.0 or higher is installed with S3 service running on one or more nodes, recommended to setup zone to use EC data pool for best storage efficiency. . Please refer to the S3 setup guide.
- In zonegroup a new placement target named “backups” has been created and is defined in the zone. Please refer to the S3 setup guide.

3. PetaSAN Configuration

S3 Settings

- Need to use https, in this example we are using a certificate with common name “s3-service” and using port 8000.
- For generated self signed certificates, the common name needs to match the S3 service URL that we will use to access the service, it needs to be configured in the hosts file as done in this example as shown below or needs to be defined in a round robin DNS. We also need to download certificate and then install it on the Veeam server as will be shown later.

Configuration > S3 Settings

S3 Settings

RadosGW http port:*

Load Balancer

Port:*

HTTPS ⓘ

Private Key:*

```
-----BEGIN PRIVATE KEY-----
MIIEQgIBADANBgkqhkiG9w0BAQEFAASCSSwggkoAgEAAoICAQDadvCvBvFrF25h2
utbgnC1//zmonojk3R5+Wj+Xe41S0y1wvFsnrHd9rqHeQsJEGY9M0gimqm4sFg0
ea8qETT81k0IRTQjZ7p//xc15x90rd5pouj8uSn3z6wmmT415cac3gw568uRbk
FqT1xugYvvhCkD+SmmP2RwMl6IFEKxwaS/YUo+eHydQv3ASDhNjhb1dub1DaSZ1
v/kQUARX7a8TzPumjG1VsTtEBxRvv0BI//ShbAZomo2PqmkQ7FzmqlLY4ckKmb
```

Generate

Certificate:*

```
-----BEGIN CERTIFICATE-----
MIIE0TCCARmgAwIBAgIUf6L0yASC4HzoA1C4bdAki/9yMxwDQYJKoZIhvcNAQEL
BQAwFETMBEGA1UEAwkuzMzZVydmljZTAeFw0yMjA0MjA0MjA0MjA0MjA0MjA0
MDEuNDAA4MjNaMBUxEzARBglNVBAwMjLXN1cnZpY2UwggEjMA0GCSqGSIb3DQEB
AQUAA4ICDwAwggTKAoICAQDadvCvBvFrF25h2utbgnC1//zmonojk3R5+Wj+Xe41S
0y1wvFsnrHd9rqHeQsJEGY9M0gimqm4sFg0ea8qETT81k0IRTQjZ7p//xc15x
```

Certificate CN:

Interface:*

Subnet Mask:*

 VLAN Tagging

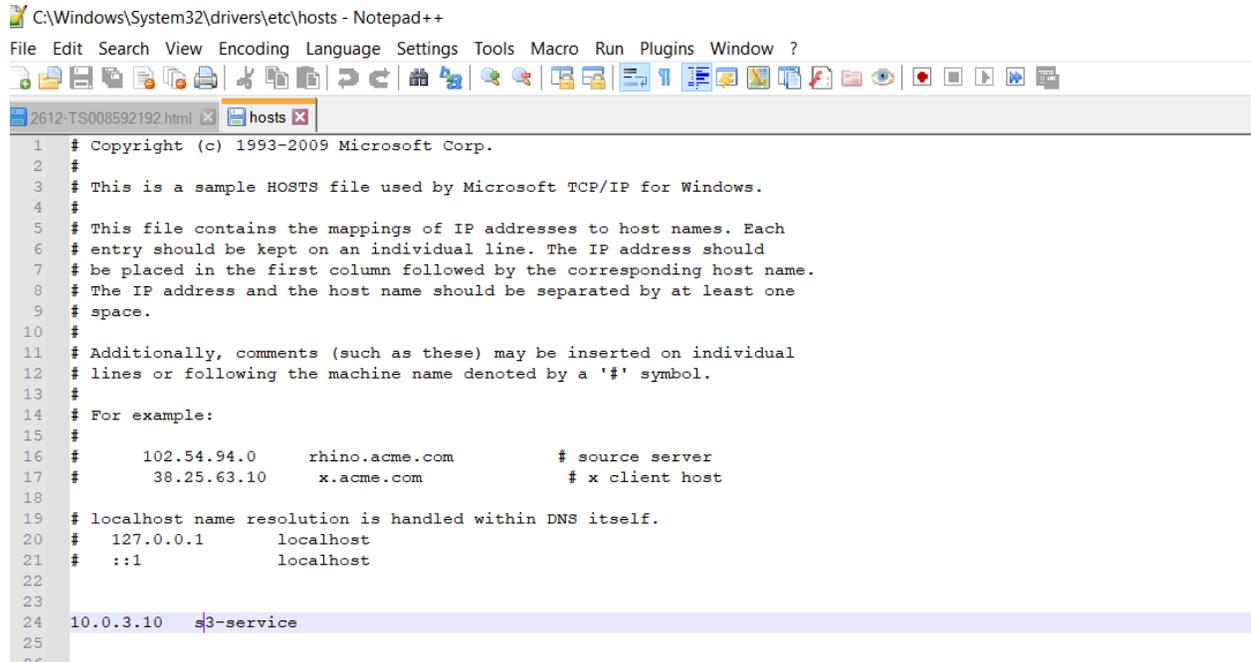
Public IP Range: From: * ⓘ

 To: *

Gateway:

Default Custom

Cancel Download Certificate Save



The screenshot shows a Notepad++ window titled "C:\Windows\System32\drivers\etc\hosts - Notepad++". The menu bar includes File, Edit, Search, View, Encoding, Language, Settings, Tools, Macro, Run, Plugins, Window, and ?. The toolbar contains various icons for file operations. The active tab is "hosts". The content of the file is as follows:

```
1 # Copyright (c) 1993-2009 Microsoft Corp.
2 #
3 # This is a sample HOSTS file used by Microsoft TCP/IP for Windows.
4 #
5 # This file contains the mappings of IP addresses to host names. Each
6 # entry should be kept on an individual line. The IP address should
7 # be placed in the first column followed by the corresponding host name.
8 # The IP address and the host name should be separated by at least one
9 # space.
10 #
11 # Additionally, comments (such as these) may be inserted on individual
12 # lines or following the machine name denoted by a '#' symbol.
13 #
14 # For example:
15 #
16 #       102.54.94.0     rhino.acme.com       # source server
17 #       38.25.63.10   x.acme.com           # x client host
18
19 # localhost name resolution is handled within DNS itself.
20 #   127.0.0.1       localhost
21 #   ::1             localhost
22
23
24 10.0.3.10   s3-service
25
26
```

Add S3 user

- In this example we will create new S3 user named Veeam that stores its data in backups placement target which was created earlier.

≡ PetaSAN ➔

Add S3 User S3 Configuration > S3 Users > Add S3 User

ID:

Display Name:

Email:

Suspended:
 Yes No

Default placement target: ⓘ

Max number of buckets:

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 adding the S3 user, view the user's generated access and secret keys which will be used later while configuring Veeam S3 object storage repository.

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

ID:
10

Display Name:
Veeam

Email:

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

Subusers:
[+ Add Subuser](#)

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

User Quota

Max Size:
 Unlimited Limit Size

Max Number of Objects:
 Unlimited Limit Number

Access Key ID:
G7EZNDZ0V64860NBDQCL

Secret Access Key:
Zuv1TgskURPADMDHmBOXIZFhh0rAUJ0eCXs8GWgQ

Size used:
0.0 GB

Number of objects:
0

Search:

Previous Next

Cancel [Regenerate Keys](#) [Save](#)

4. Veeam Configuration

You can setup Veeam using one of the following configurations:

- **PetaSAN S3 as a direct backup repository**
In this setup, Veeam backups directly to the S3 repository. This feature is supported in Veeam starting from version 12.
- **PetaSAN S3 as a capacity tier in a scale-out repository**
In this setup, a Scale Out Backup Repository (SOBR) is configured. The SOBR consists of 2 tiers:
-A non S3 backup repository acting as a performance tier
-An S3 repository acting as a capacity tier
this feature is supported in Veeam starting from version 11.

Here we show how to create each repository type:

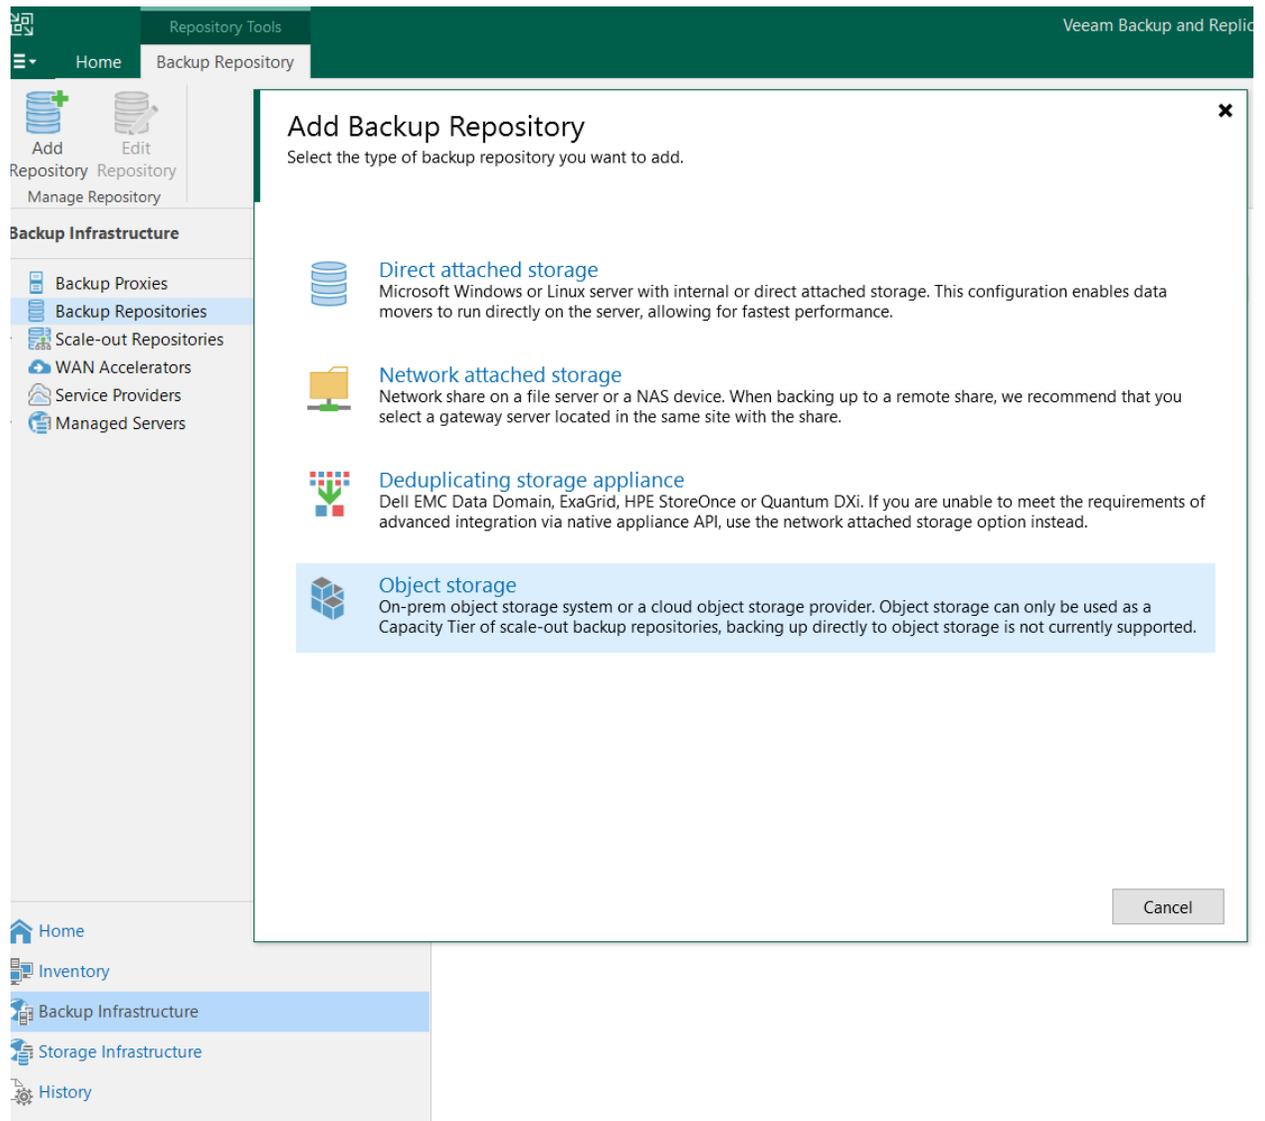
[Create Object Storage Repository](#)

[Create Direct Attached Storage Repository](#)

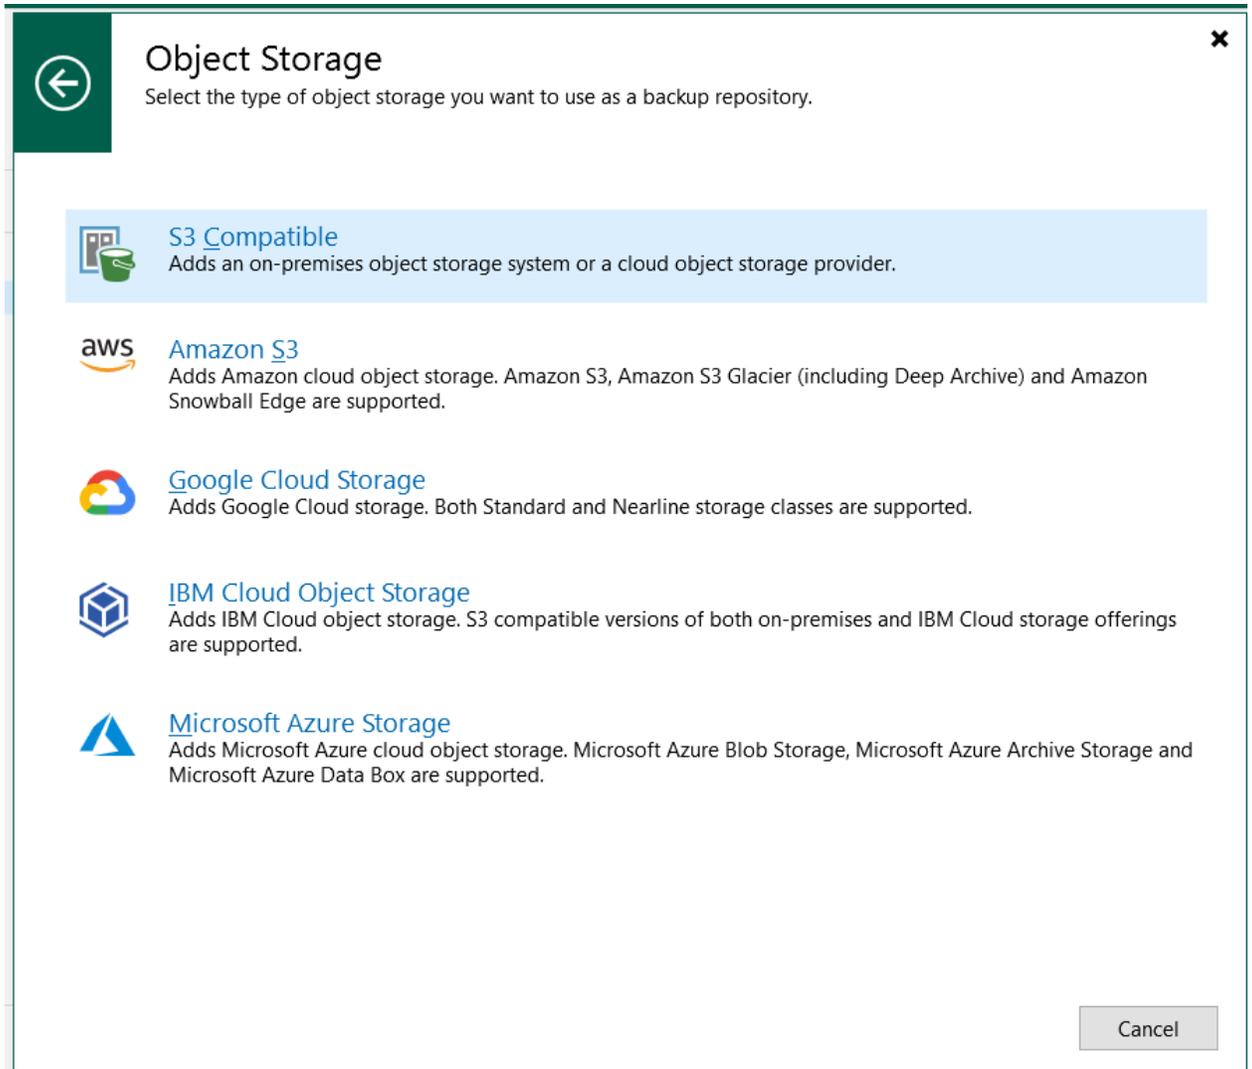
[Scale-out repository](#)

Create Object Storage Repository

From the Backup Infrastructure section select to add new backup repository and then select the object storage type



- Select the S3 Compatible Option



- Enter the S3 Object Repository Name

New Object Storage Repository ✕

 **Name**
Type in a name and description for this object storage repository.

Name	Name: <input type="text" value="PetaSAN Object storage repository"/>
Account	Description: <input type="text" value="Created by DESKTOP-95RGNP3\lenovo at 4/8/2022 1:13 AM."/>
Bucket	
Summary	

Limit concurrent tasks to:

Use this setting to limit the maximum number of tasks that can be processed concurrently in cases when your object storage is overloaded or cannot keep up with the number of API requests issued by multiple object storage offload tasks.

- Enter the S3 service name and port number, then the access key and secret key of the PetaSAN S3 user we want to use. In this example we will use the S3 user we have created named Veam

New Object Storage Repository

Account
Specify account to use for connecting to S3 compatible storage system.

Name

Account

Bucket

Summary

Service point: s3-service:8000

Region: zonegroup

Credentials: [Add...](#)

[Manage cloud accounts](#)

Credentials

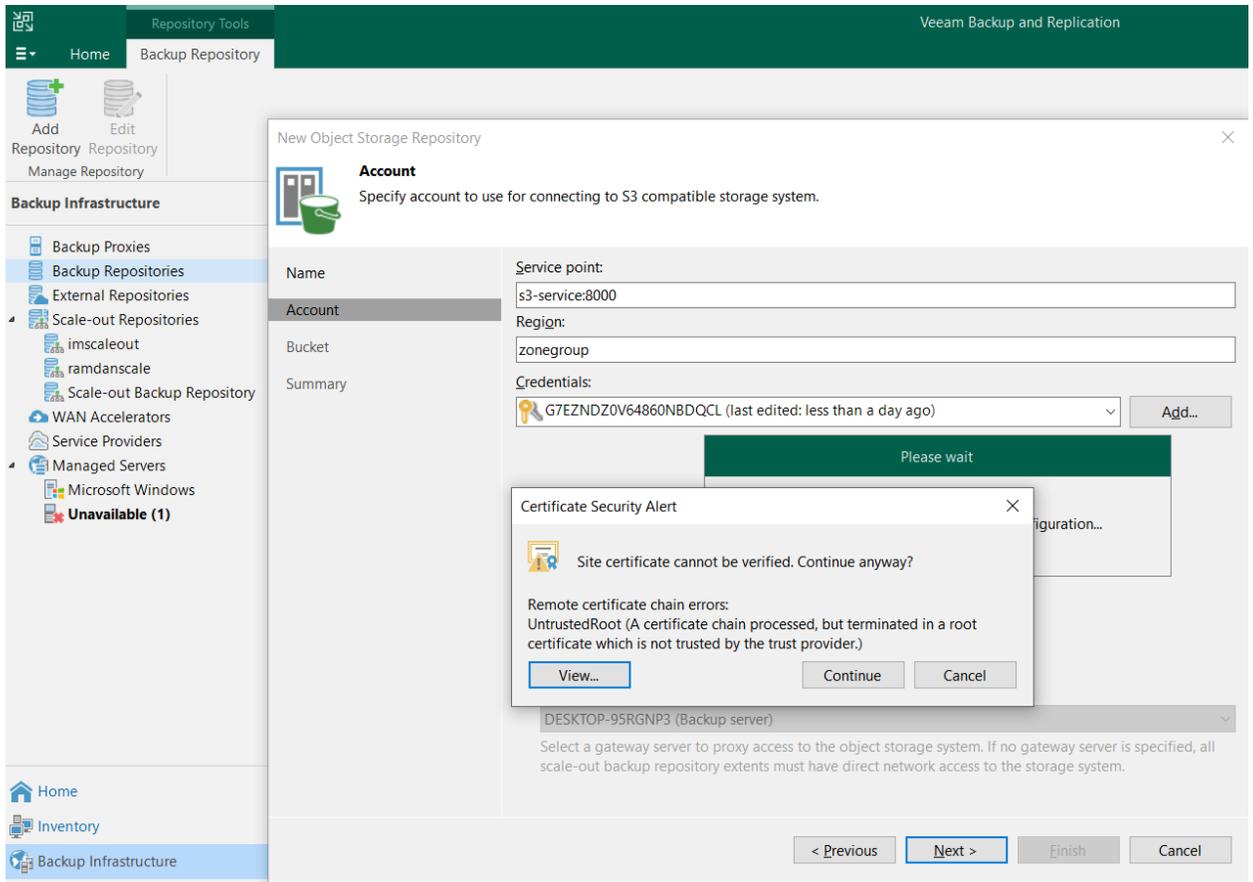
Access key: G7EZNDZ0V64860NBDQCL

Secret key:

Description:

OK Cancel

< Previous Next > Finish Cancel

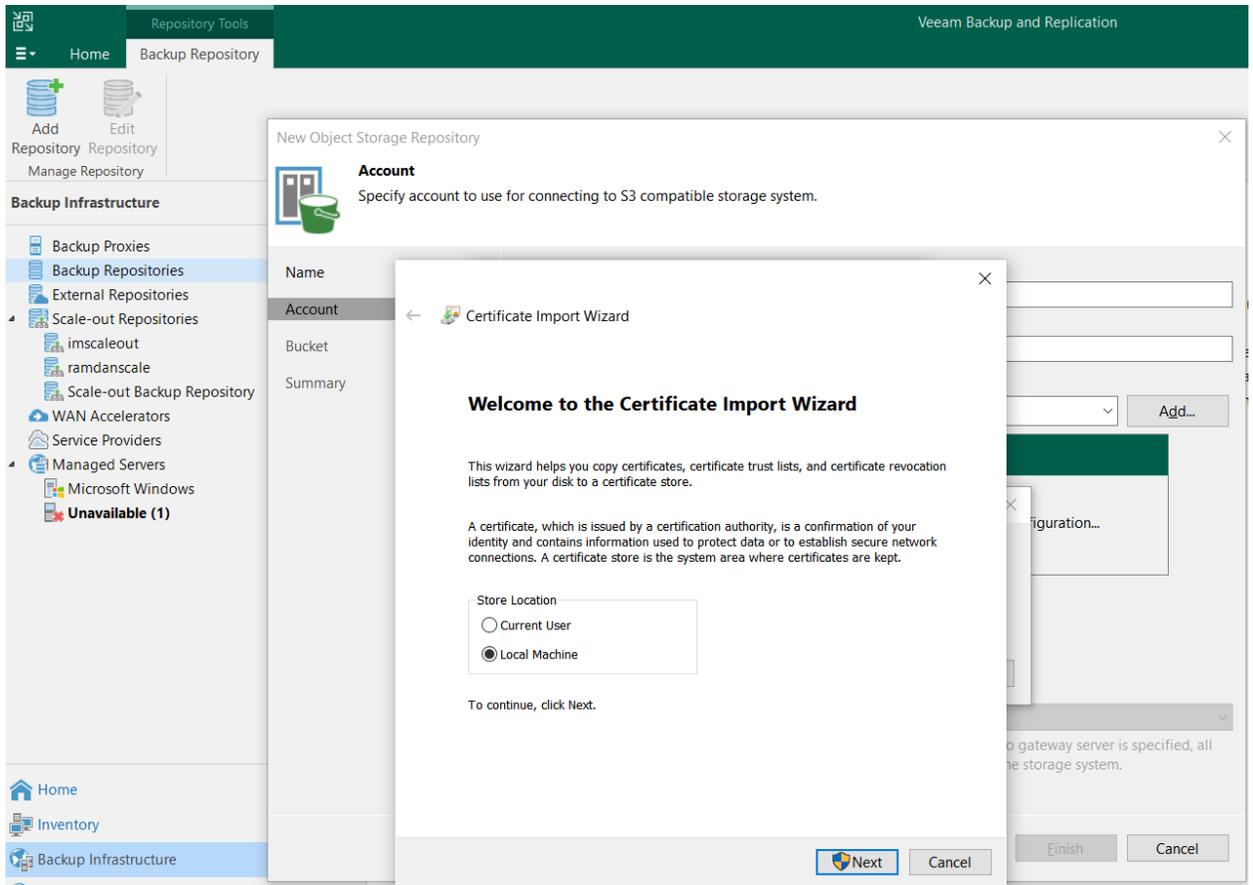


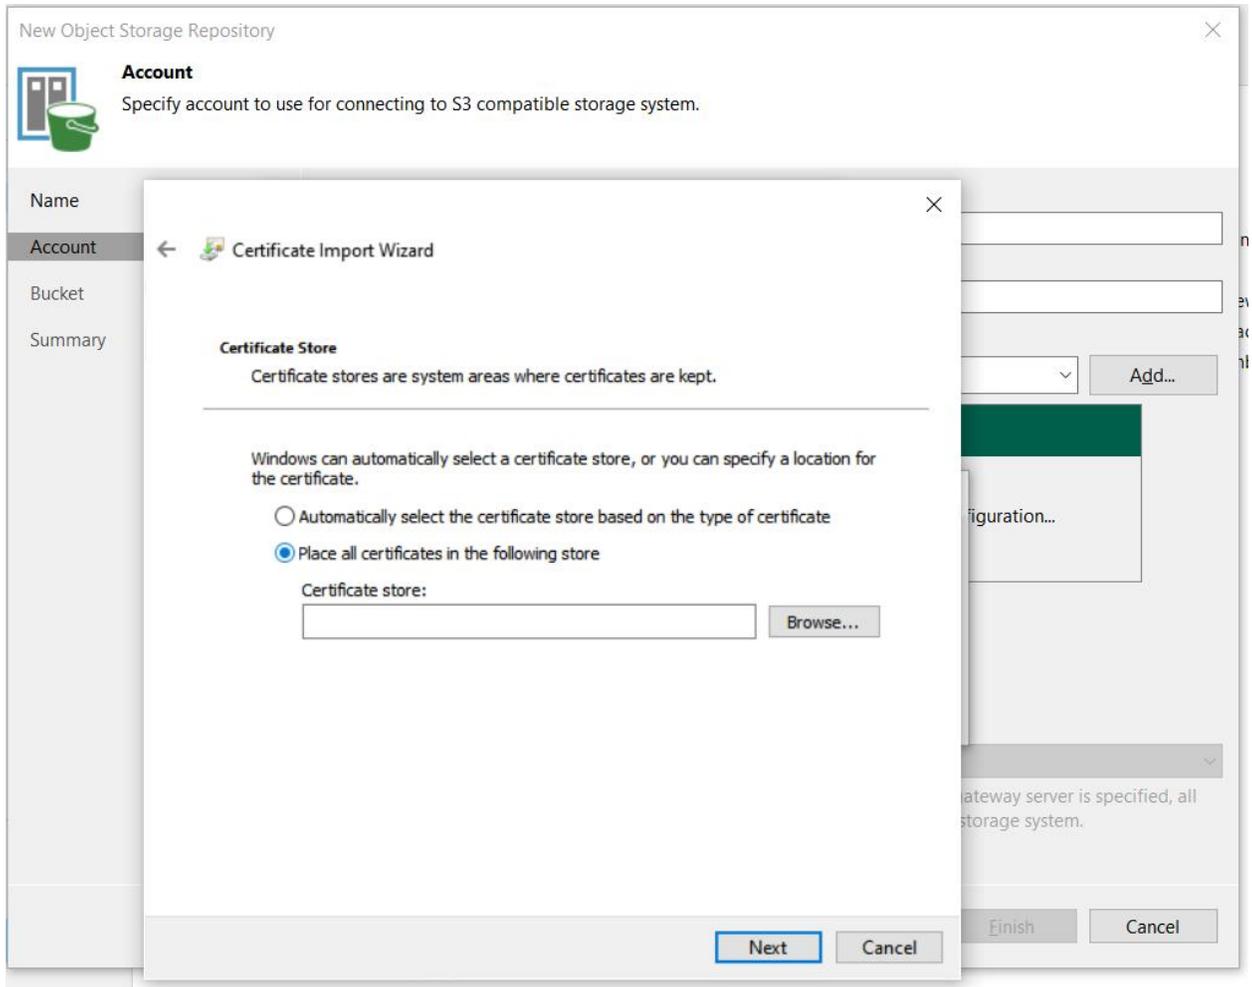
- Certificate Security Alert dialog will appear, click on View Certificate to be able to install the certificate we had downloaded from the PetaSAN S3 settings.

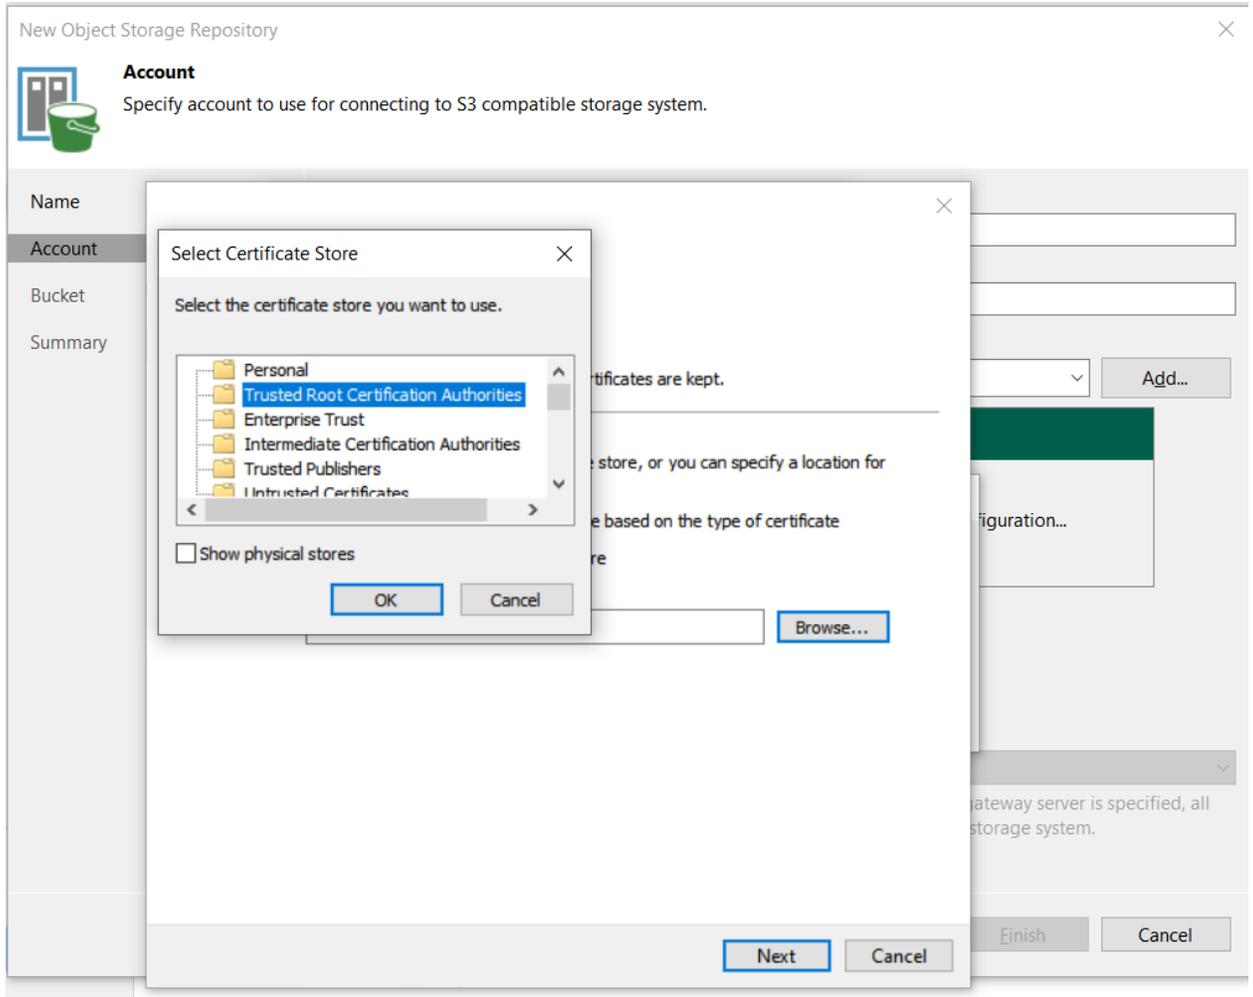
The screenshot displays the Veeam Backup and Replication interface. The main window is titled 'New Object Storage Repository' and is in the 'Account' step. A 'Certificate' dialog box is open, showing the following information:

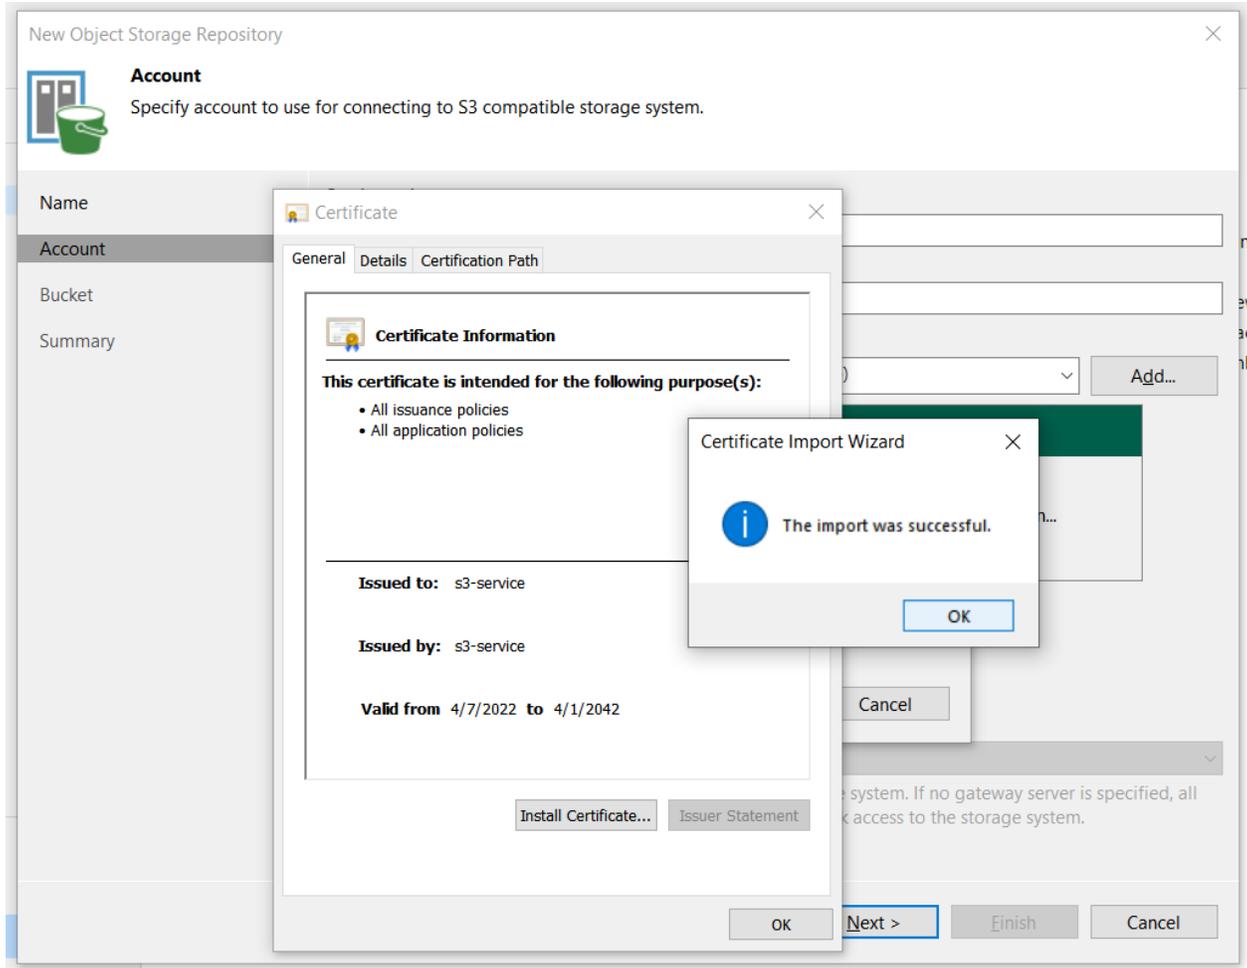
- Certificate Information**
- This certificate is intended for the following purpose(s):**
 - All issuance policies
 - All application policies
- Issued to:** s3-service
- Issued by:** s3-service
- Valid from:** 4/7/2022 to 4/1/2042

The dialog box has buttons for 'Install Certificate...', 'Issuer Statement', 'OK', and 'Cancel'. The background window shows a sidebar with 'Backup Infrastructure' expanded to 'Backup Repositories', and a main area with fields for 'Name', 'Account', 'Bucket', and 'Summary'. The 'Next >' button is highlighted.









- After installing the certificate click back then next again ,you should notice the certificate security alert has been disappeared

New Object Storage Repository

Account
Specify account to use for connecting to S3 compatible storage system.

Name	Service point: s3-service:8000
Account	Region: zonegroup
Bucket	Credentials: G7EZNDZ0V64860NBDQCL (last edited: less than a day ago) Add...
Summary	Manage cloud accounts

Use the following gateway server:
DESKTOP-95RGNP3 (Backup server)

Select a gateway server to proxy access to the object storage system. If no gateway server is specified, all scale-out backup repository extents must have direct network access to the storage system.

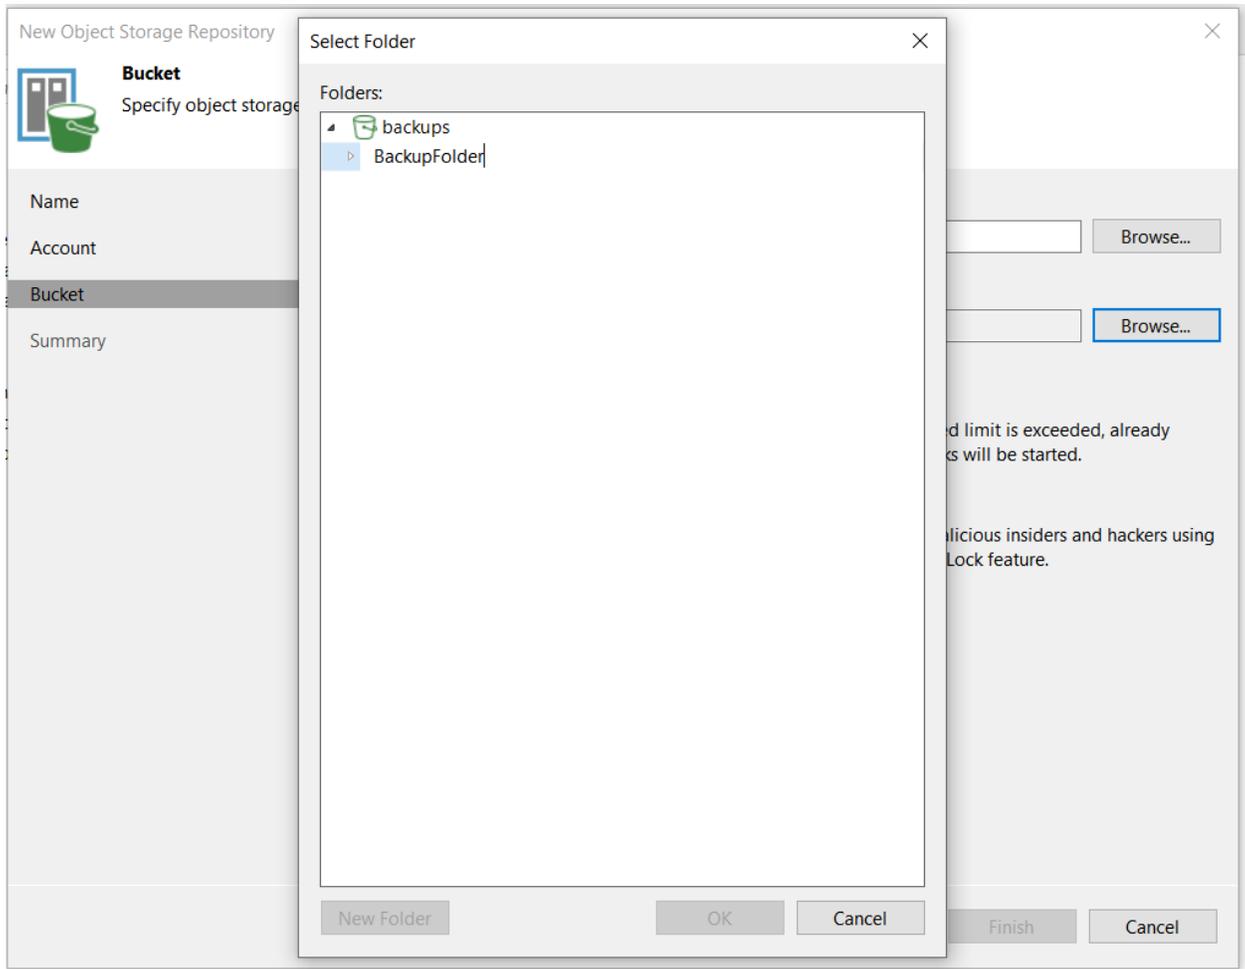
< Previous Next > Finish Cancel

- Select next and then click on browse to list the existing buckets, you can create buckets using S3 client tools like S3 Browser (refer to the S3 setup guide).

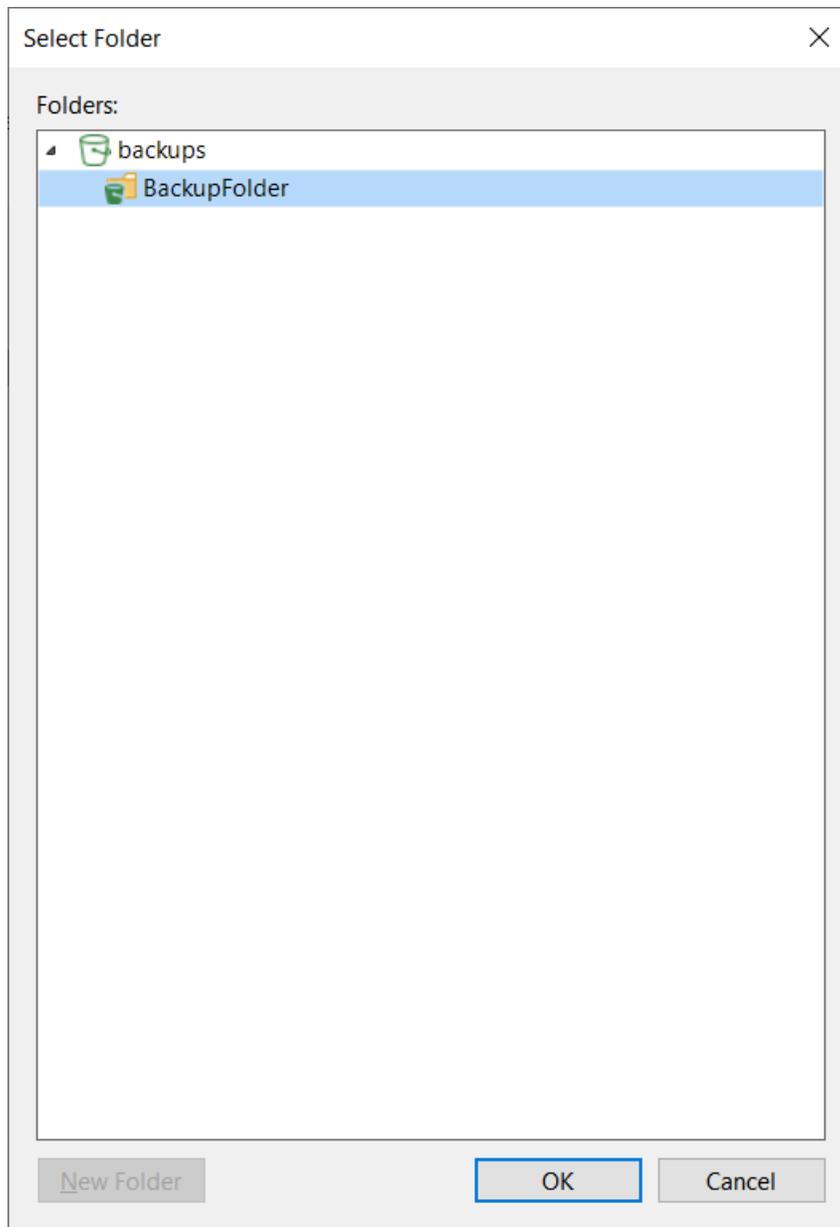
New Object Storage Repository ×

 **Bucket**
Specify object storage system bucket to use.

Name	Bucket:	<input type="text"/>	<input type="button" value="Browse..."/>
Account			
Bucket	Folder:	<input type="text"/>	<input type="button" value="Browse..."/>
Summary	<input type="checkbox"/> Limit object storage consumption to: <input type="text" value="10"/> <input type="text" value="TB"/> This is a soft limit to help control your object storage spend. If the specified limit is exceeded, already running backup offload tasks will be allowed to complete, but no new tasks will be started.		
	<input type="checkbox"/> Make recent backups immutable for: <input type="text" value="30"/> days Protects recent backups from modification or deletion by ransomware, malicious insiders and hackers using native object storage capabilities. Object storage must support S3 Object Lock feature.		



- Select a folder under the selected bucket



- Select the repository settings, whether to limit the object storage size or not and if you want to store recent backups as immutable and for how many days.

New Object Storage Repository ✕

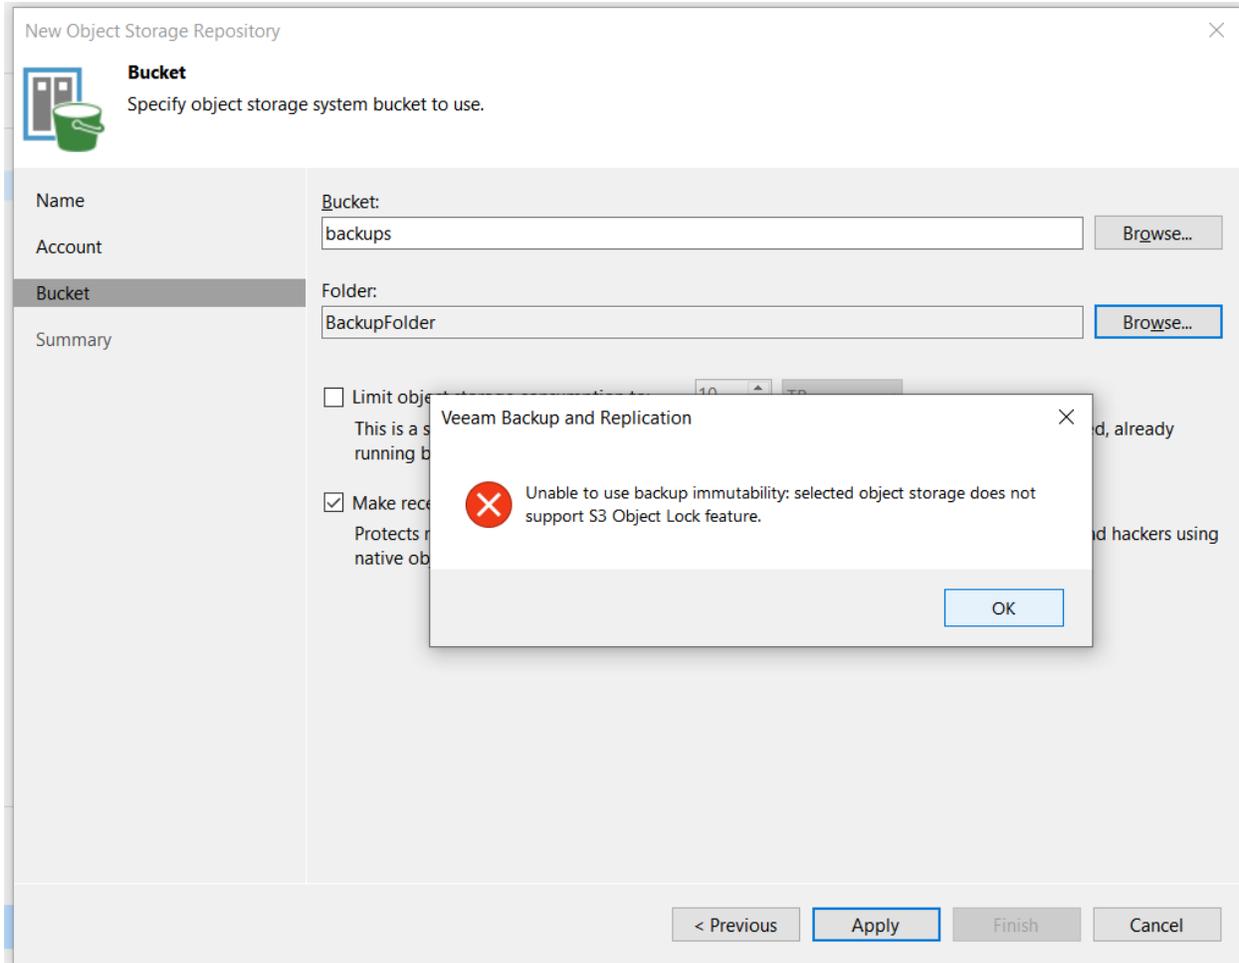
 **Bucket**
Specify object storage system bucket to use.

Name	Bucket:	<input type="text" value="backups"/>	<input type="button" value="Browse..."/>
Account			
Bucket	Folder:	<input type="text" value="BackupFolder"/>	<input type="button" value="Browse..."/>
Summary			

Limit object storage consumption to:
This is a soft limit to help control your object storage spend. If the specified limit is exceeded, already running backup offload tasks will be allowed to complete, but no new tasks will be started.

Make recent backups immutable for: days
Protects recent backups from modification or deletion by ransomware, malicious insiders and hackers using native object storage capabilities. Object storage must support S3 Object Lock feature.

- To enable the use of immutable option, bucket has to be lock enabled during bucket creation, otherwise the following error will appear:



Note:

Immutability is a method that prevents data deletion or modification from storage.

Knowing that data is critical to all businesses—having an immutable copy of your backup data ensures that there is an untouched version of that source data which is always recoverable.

To create a bucket with lock option enabled, use Amazon Cli tool (refer to the S3 setup guide for more detail).

You can install it from

<https://docs.aws.amazon.com/cli/latest/userguide/install-cliv2-windows.html>

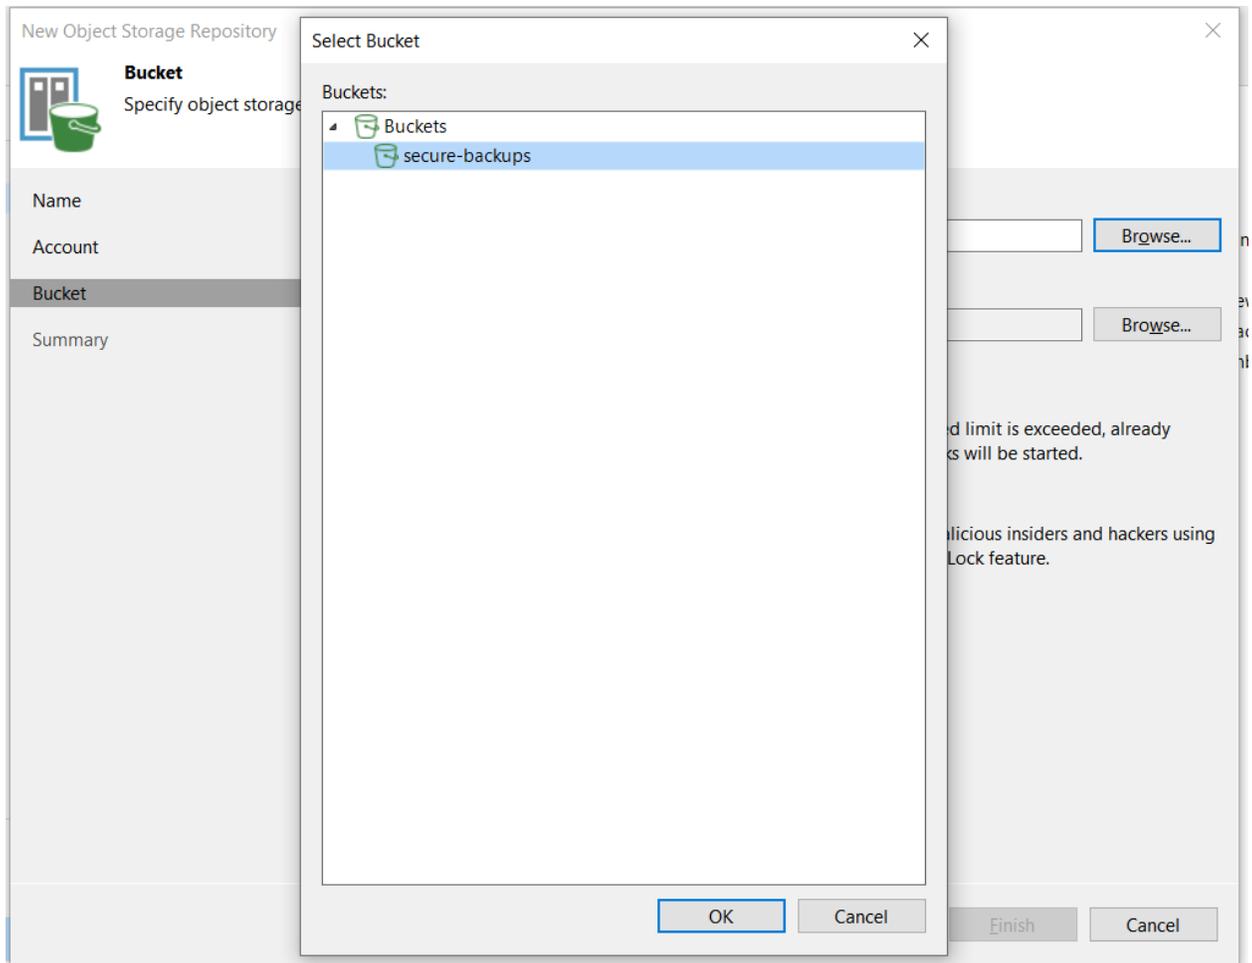
- Setup the service certificate
`aws configure set default.ca_bundle "c:\s3-service.crt"`
- After installation we will configure the aws using command
`aws configure`

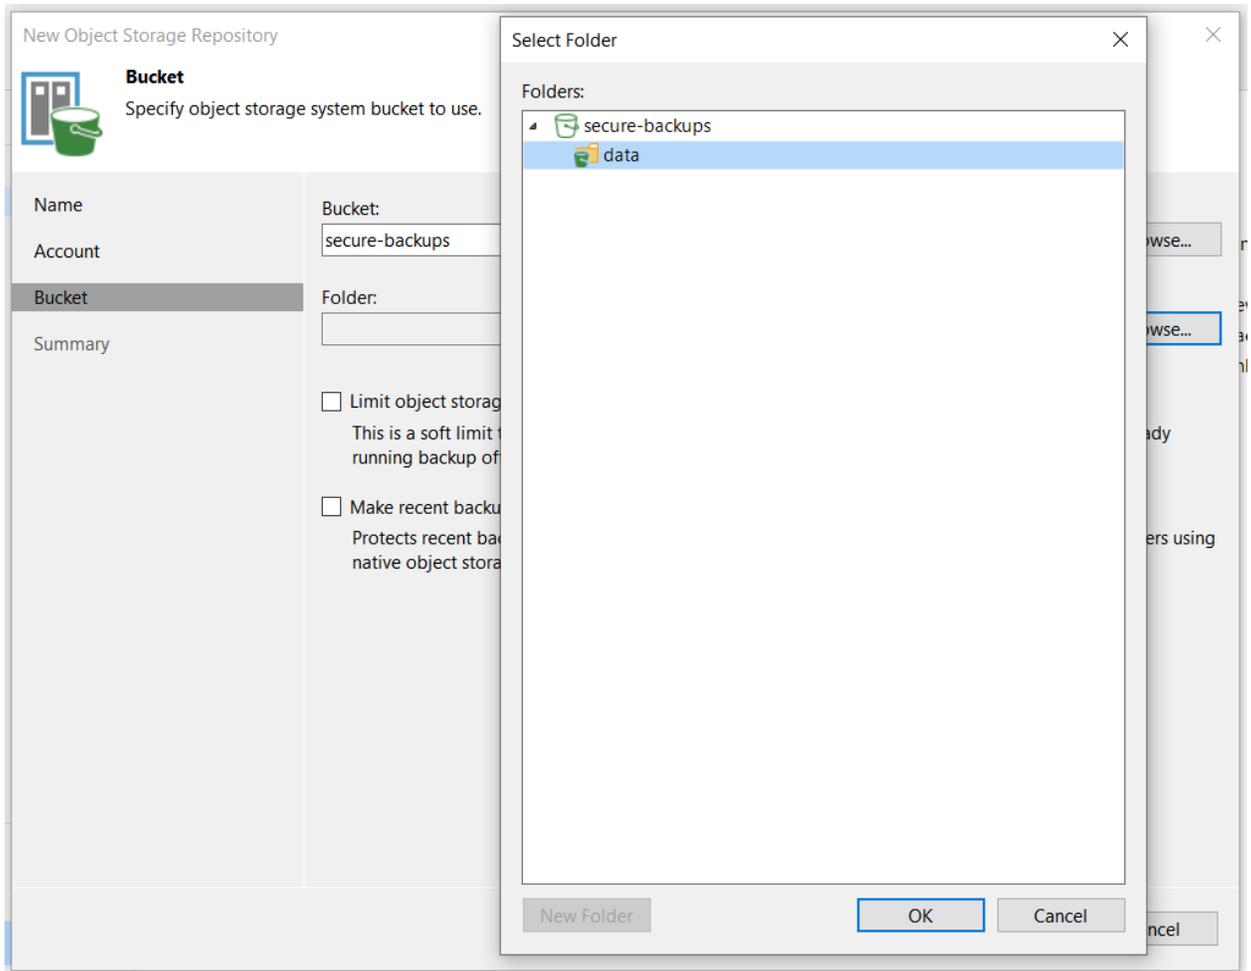
In this example we use the PetaSAN S3 user named Veeam and we enter its access and secret keys as well as the zonegroup name.

```
Command Prompt - aws configure
Microsoft Windows [Version 10.0.19044.1586]
(c) Microsoft Corporation. All rights reserved.

C:\Users\lenovo>aws configure
AWS Access Key ID [*****DQCL]: G7EZNDZ0V64860NBDQCL
AWS Secret Access Key [*****GwgQ]: Zuv1TgskURPADMDHmBOXIZFhh0rAUj0eCXs8GwgQ
Default region name [default]: zonegroup
Default output format [None]:
```

- Use the following command to create a bucket named “secure-backups” with lock is enabled
`aws s3api create bucket --bucket secure-backups --endpoint-url https://s3-service:8000 --object-lock-enabled-for-bucket`
- We continue this example using the newly created bucket with the immutability feature





- Select Make recent backups immutable for 30 days

New Object Storage Repository ✕

 **Bucket**
Specify object storage system bucket to use.

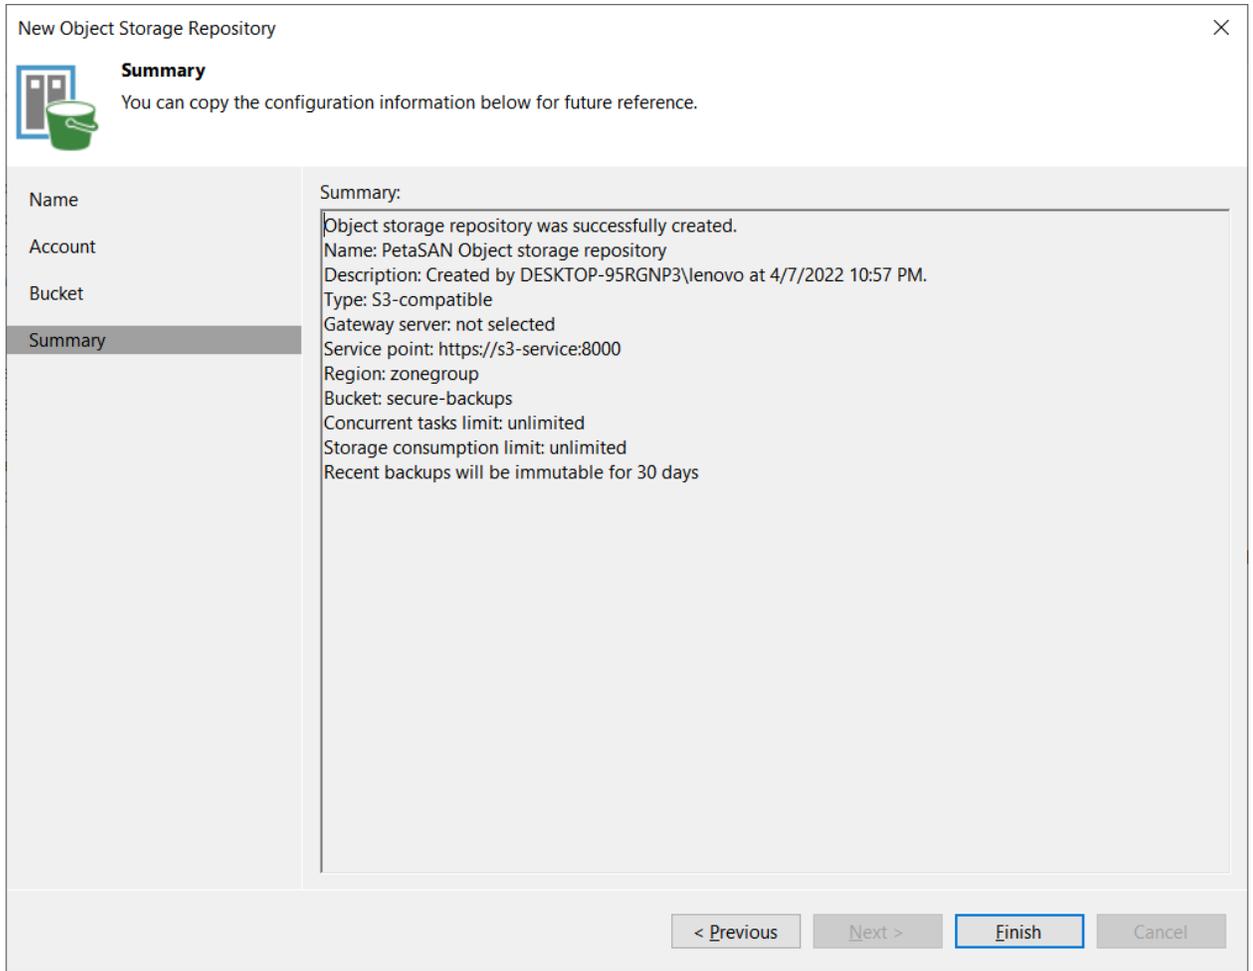
Name	Bucket:	<input type="text" value="secure-backups"/>	<input type="button" value="Browse..."/>
Account			
Bucket	Folder:	<input type="text" value="data"/>	<input type="button" value="Browse..."/>
Summary			

Limit object storage consumption to:

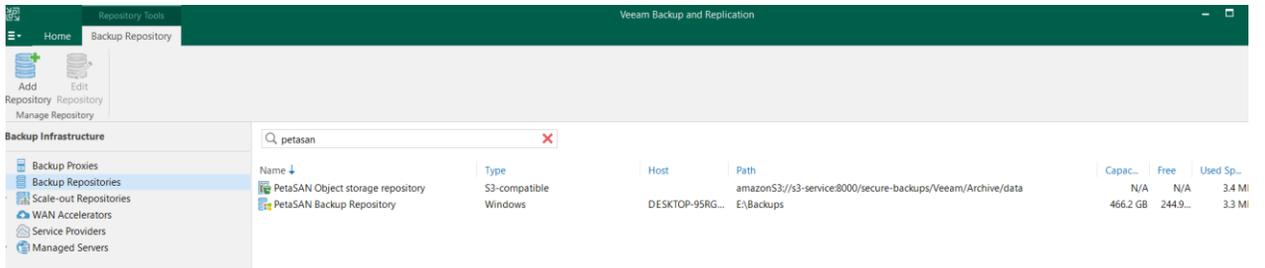
This is a soft limit to help control your object storage spend. If the specified limit is exceeded, already running backup offload tasks will be allowed to complete, but no new tasks will be started.

Make recent backups immutable for: days

Protects recent backups from modification or deletion by ransomware, malicious insiders and hackers using native object storage capabilities. Object storage must support S3 Object Lock feature.

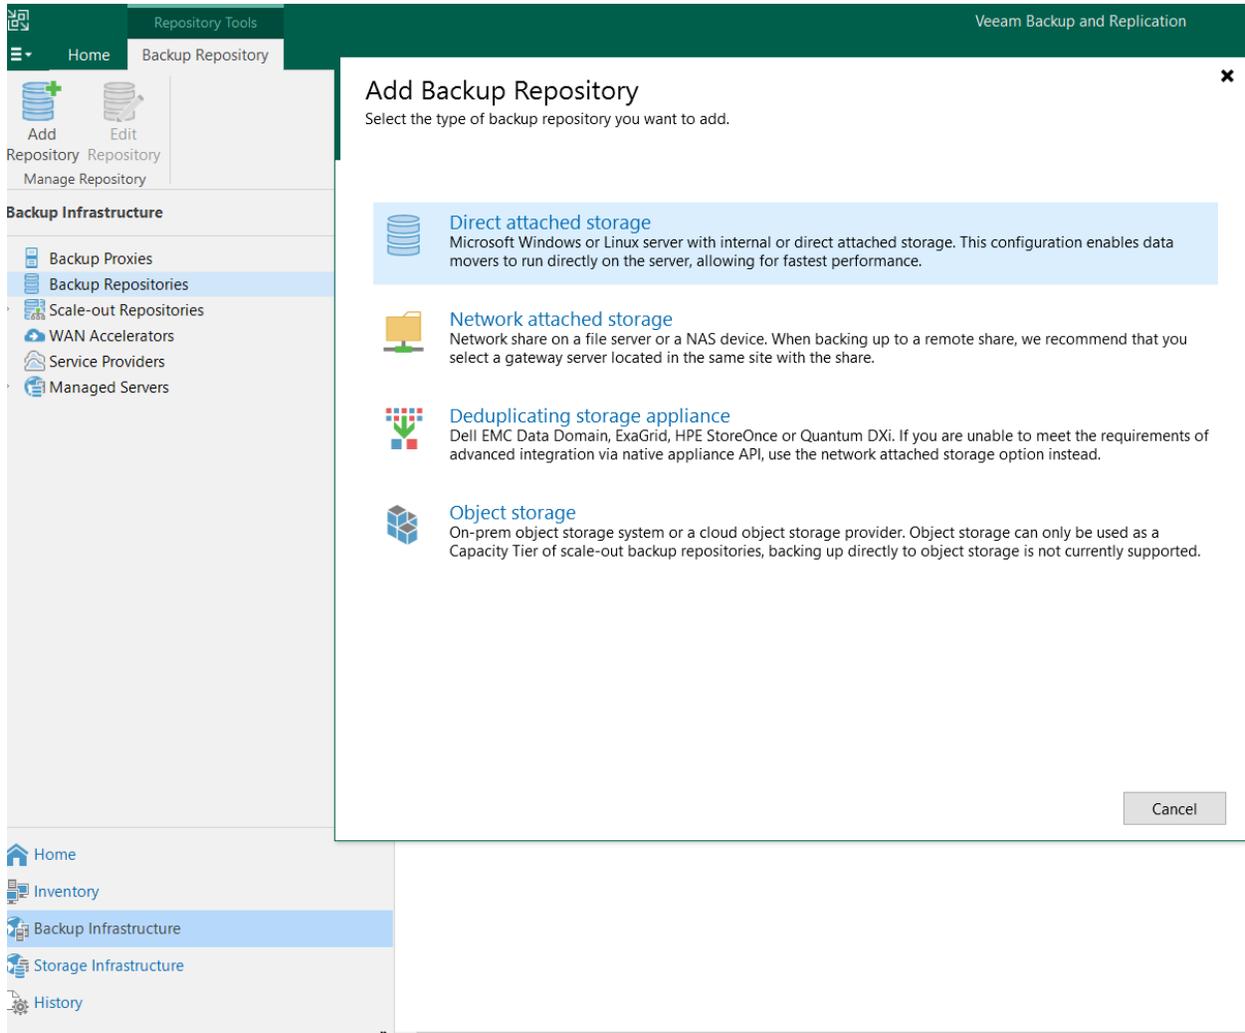


After saving the object storage repository you can view it in the repositories list

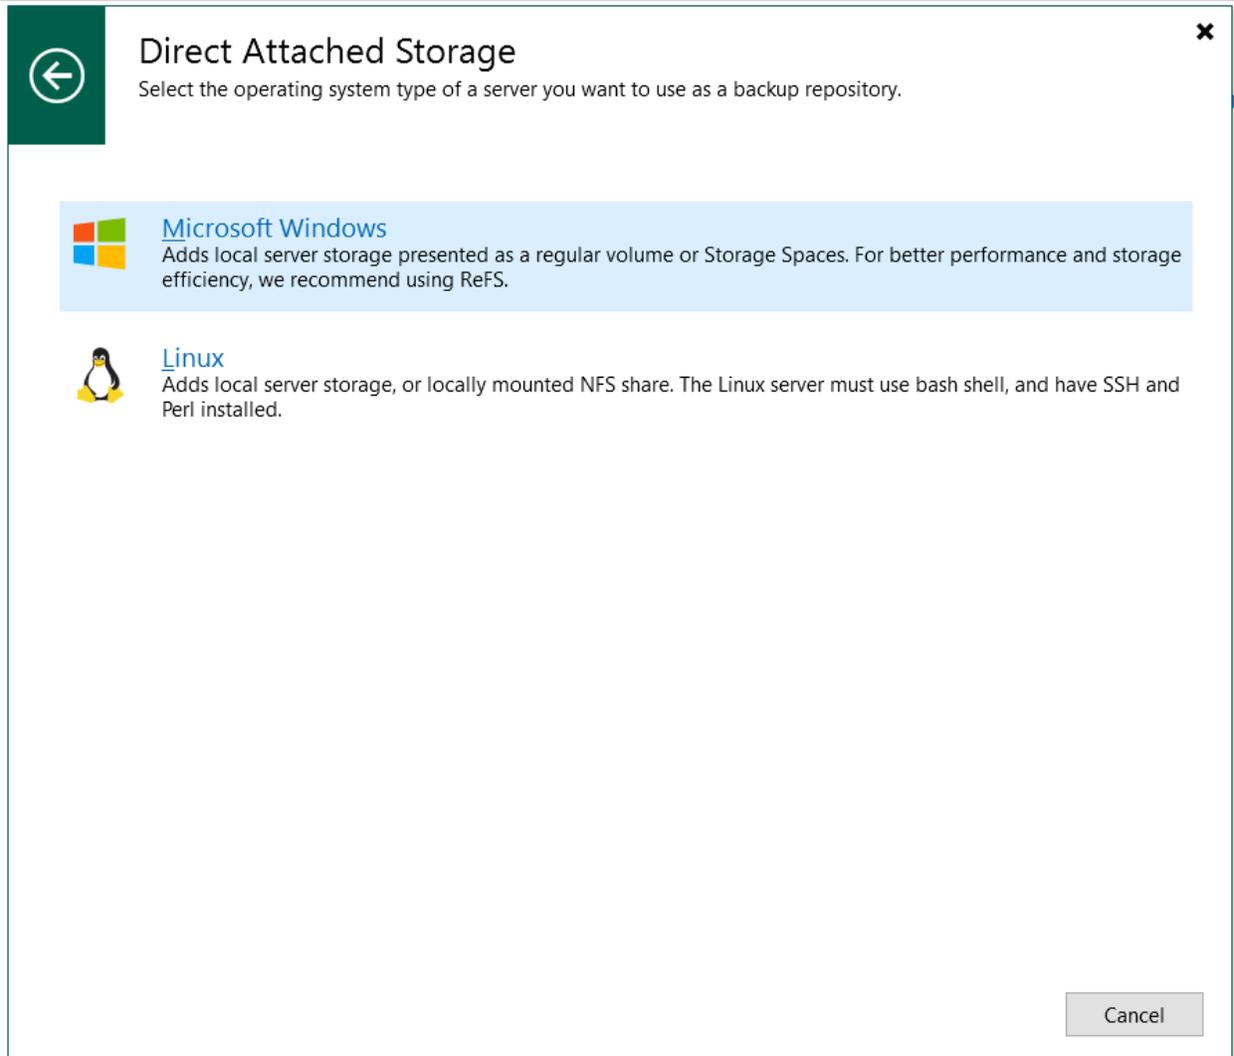


Create Direct Attached Storage Repository

From the Backup Infrastructure section select to add new backup repository and then select the A Direct attached storage type which needs to be created to be used as the performance tier in case of Veeam version 11.



- Then select Microsoft Windows



- Enter the backup repository name

New Backup Repository ✕

 **Name**
Type in a name and description for this backup repository.

Name	Name: <input type="text" value="PetaSAN Backup Repository"/>
Server	Description: <input type="text" value="Created by DESKTOP-95RGNP3\lenovo at 4/7/2022 10:50 PM."/>
Repository	
Mount Server	
Review	
Apply	
Summary	

New Backup Repository ✕

 **Repository**
Type in path to the folder where backup files should be stored, and set repository load control options.

Name	Location
Server	Path to folder: <input type="text" value="E:\Backups"/> Browse...
Repository	 Capacity: 466.2 GB Populate
Mount Server	 Free space: 245.6 GB
Review	Load control
Apply	Running too many concurrent tasks against the repository may reduce overall performance, and cause I/O timeouts. Control storage device saturation with the following settings:
Summary	<input checked="" type="checkbox"/> Limit maximum concurrent tasks to: <input type="text" value="4"/>
	<input type="checkbox"/> Limit read and write data rate to: <input type="text" value="1"/> MB/s

Click Advanced to customize repository settings. ⚙️ Advanced

< Previous Next > Finish Cancel

- Select the mount server

New Backup Repository ✕

 **Mount Server**
Specify a server to mount backups to when performing advanced restores (file, application item and instant VM recoveries). Instant recoveries require a write cache folder to store changed disk blocks in.

Name	Mount server:
Server	DESKTOP-95RGNP3 (Backup server) Add New...
Repository	Instant recovery write cache folder: D:\ProgramData\Veeam\Backup\IRCache\ Browse...
Mount Server	Ensure that the selected volume has sufficient free disk space to store changed disk blocks of instantly recovered VMs. We recommend placing write cache on an SSD drive. <input checked="" type="checkbox"/> Enable vPower NFS service on the mount server (recommended) Ports... Unlocks instant recovery of any backup (physical, virtual or cloud) to a VMware vSphere VM. vPower NFS service is not used for instant recovery to a Microsoft Hyper-V VM.
Review	
Apply	
Summary	

< Previous Next > Finish Cancel

- Review your setup then apply it to create the repository

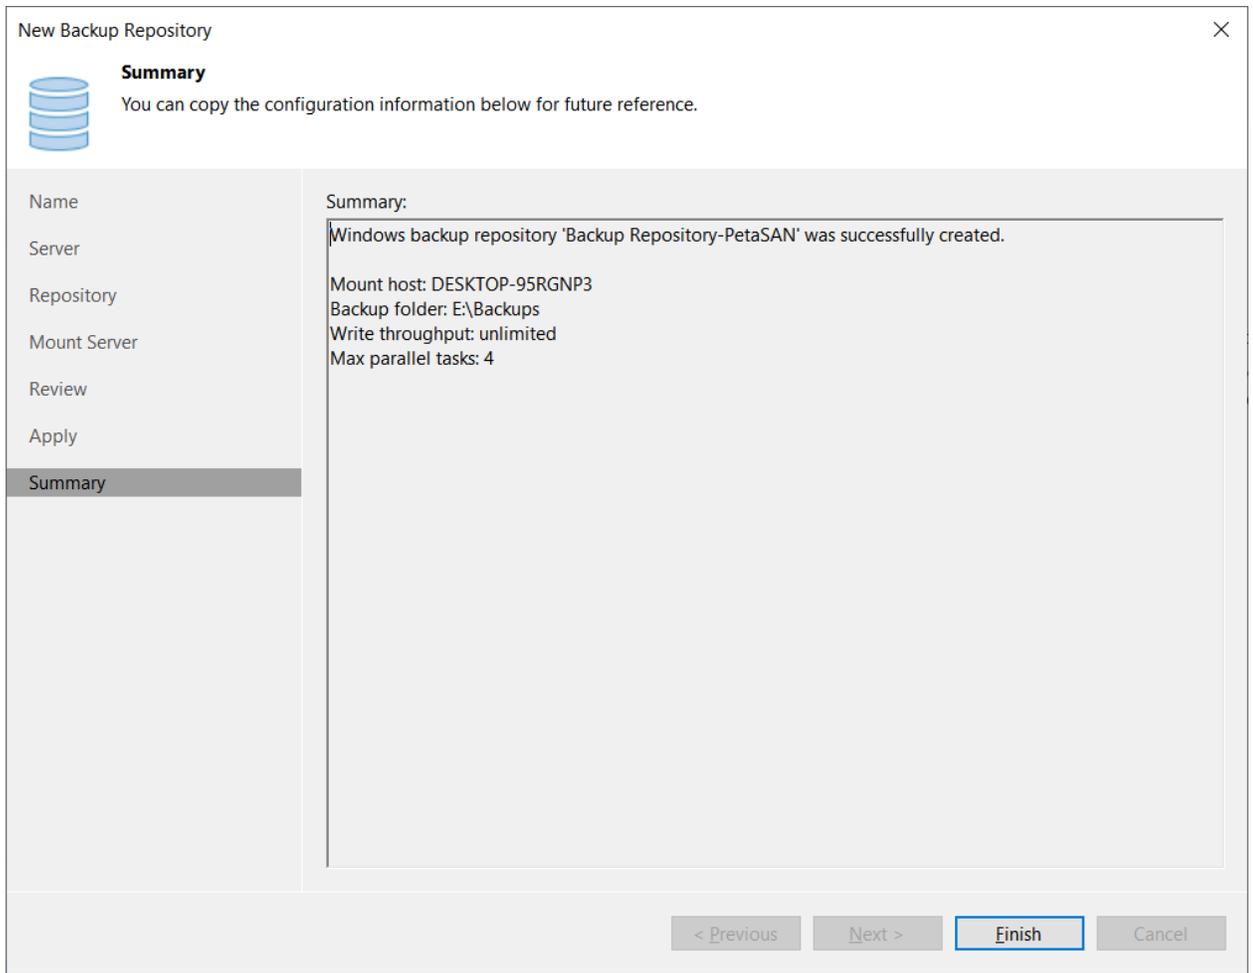
New Backup Repository ✕

 **Review**
Please review the settings, and click Apply to continue.

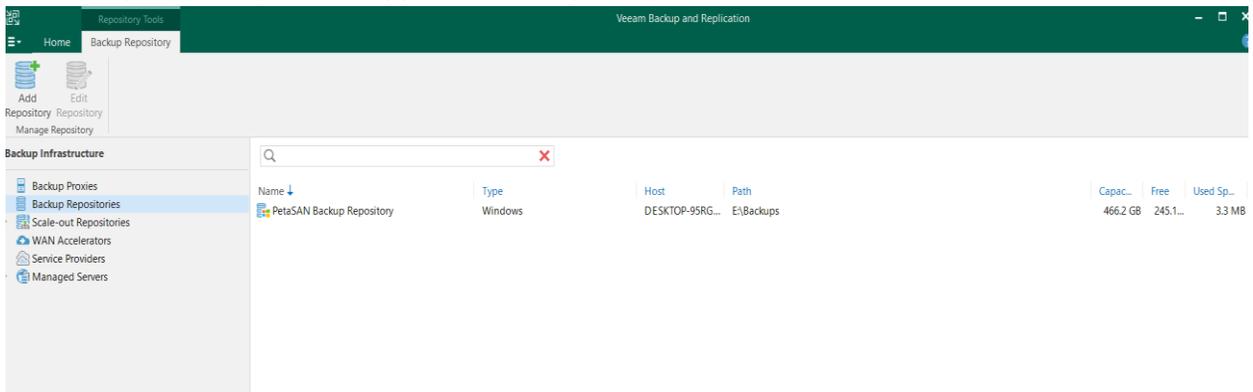
	The following components will be processed on server DESKTOP-95RGNP3:	
Name	Component name	Status
Server	Transport	already exists
Repository	vPower NFS	already exists
Mount Server	Mount Server	already exists
Review		
Apply		
Summary		

Search the repository for existing backups and import them automatically

Import guest file system index data to the catalog



The backup repository created appears in the repositories list

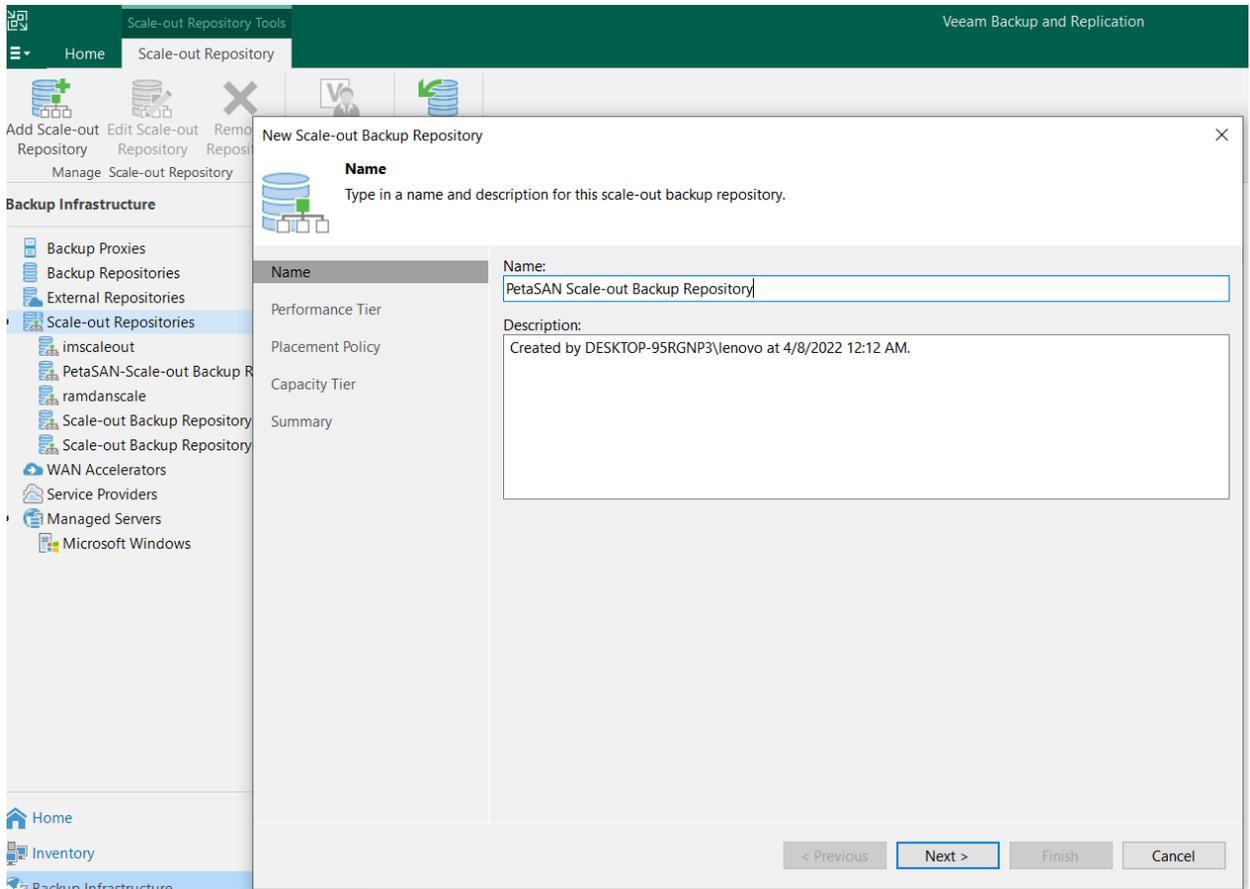


Create Scale-Out Repository

From the Backup Infrastructure section select to add new scale-out repository and then select the object storage type

Create a new scale-out repository using the previously created repositories as performance and capacity tiers ,in case of Veeam 11

- Enter the scale out repository name



- Select the performance tier repository

New Scale-out Backup Repository

Performance Tier
Select backup repositories to use as the landing zone and for the short-term retention.

Name
Performance Tier
Placement Policy
Capacity Tier
Summary

Extents:
Name

Click Advanced to sp

OK Cancel

< Previous Next > Finish Cancel

Add... Remove

Advanced

Extents

Select backup repositories to include in this scale-out backup repository.

Backup repositories:

Name	
<input checked="" type="checkbox"/> PetaSAN Backup Repository	<input type="button" value="Select All"/>
<input type="checkbox"/> Default Backup Repository	<input type="button" value="Clear All"/>
<input type="checkbox"/> Destination	
<input type="checkbox"/> repo1	

OK Cancel

New Scale-out Backup Repository ×

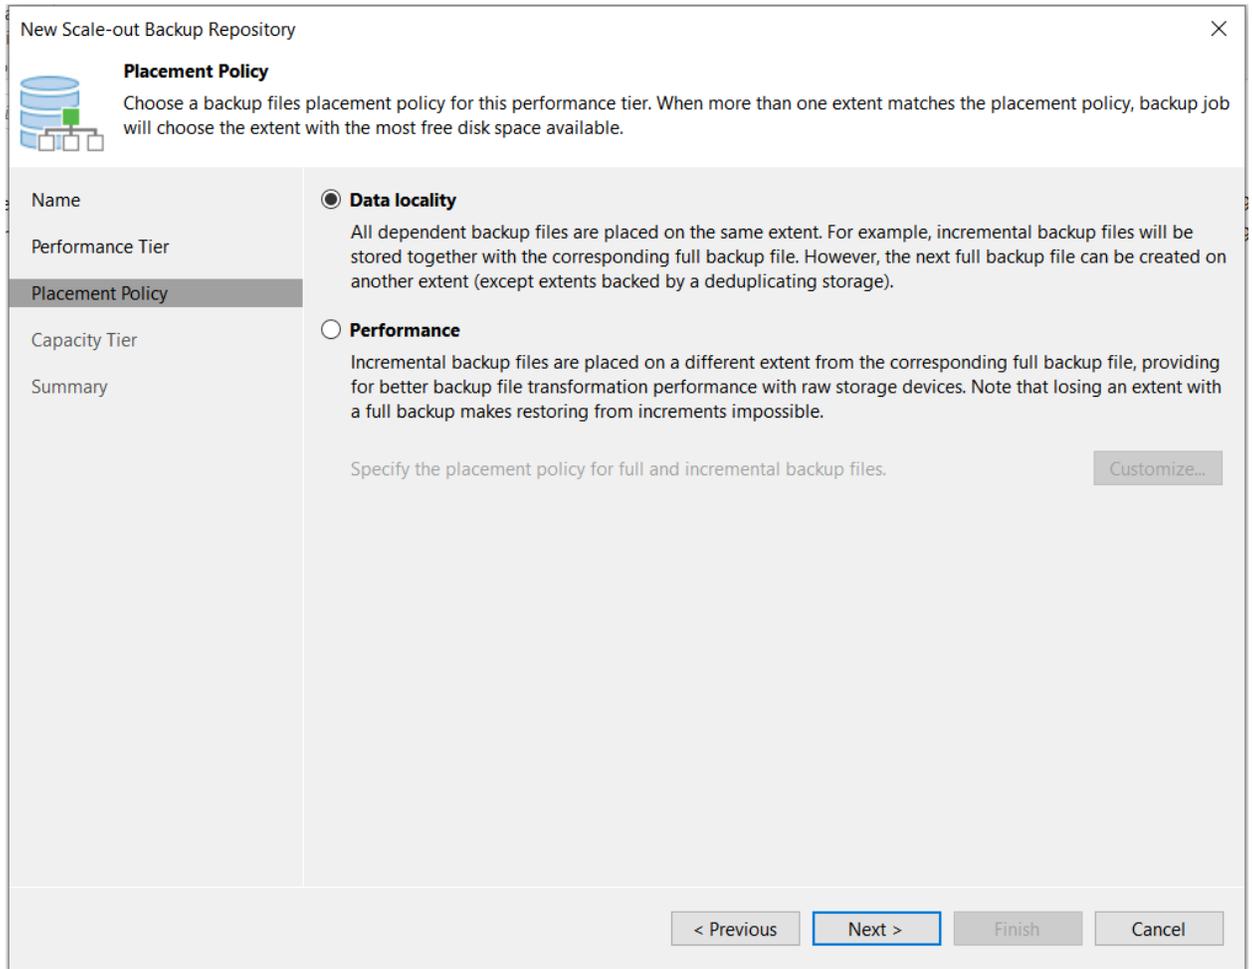
 **Performance Tier**
Select backup repositories to use as the landing zone and for the short-term retention.

Name	Extents:	
Performance Tier	Name	Add...
Placement Policy	 PetaSAN Backup Repository	Remove
Capacity Tier		
Summary		

Click Advanced to specify additional scale-out backup repository options. Advanced

< Previous Next > Finish Cancel

- Select placement policy



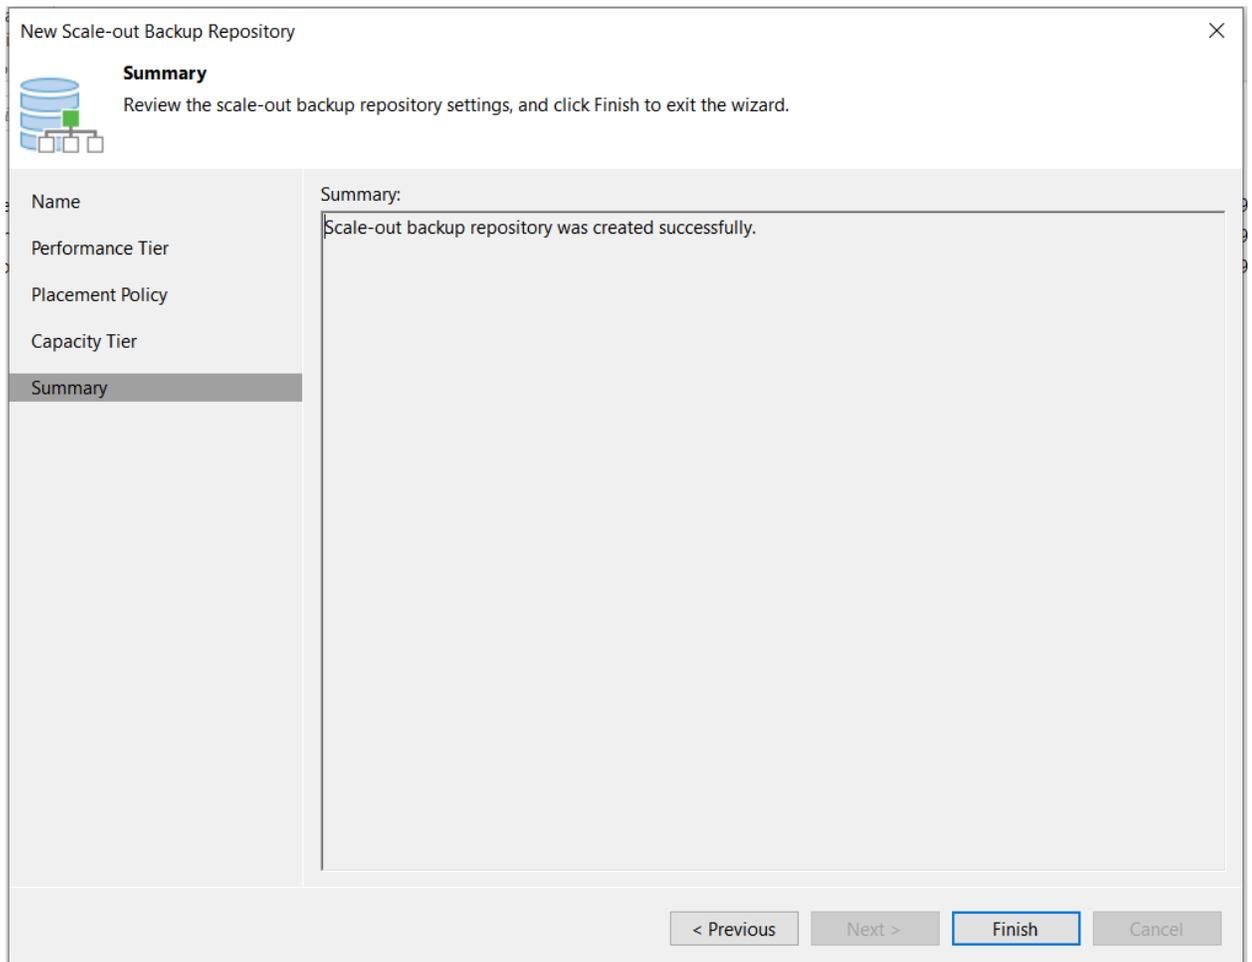
- Select the capacity tier repository and select to copy backups as soon as they are created

New Scale-out Backup Repository ✕

 **Capacity Tier**
Specify object storage to copy backups to for redundancy and DR purposes. Older backups can be moved to object storage completely to reduce long-term retention costs while preserving the ability to restore directly from offloaded backups.

Name	<input checked="" type="checkbox"/> Extend scale-out backup repository capacity with object storage: PetaSAN Object storage repository Add...
Performance Tier	
Placement Policy	Define time windows when uploading to capacity tier is allowed Window...
Capacity Tier	<input checked="" type="checkbox"/> Copy backups to object storage as soon as they are created Create additional copy of your backups for added redundancy by having all backups copied to the capacity tier as soon as they are created on the performance tier.
Summary	<input checked="" type="checkbox"/> Move backups to object storage as they age out of the operational restore window Reduce your long-term retention costs by moving older backups to object storage completely while preserving the ability to restore directly from offloaded backups. Move backup files older than <input type="text" value="14"/> days (your operational restore window) Override...
	<input type="checkbox"/> Encrypt data uploaded to object storage Password: <input type="text"/> Add... Manage passwords

< Previous Apply Finish Cancel



5. Create a backup job

Go to the Home Section and select to create new backup job

In case of direct backup repository

We will create the job to connect directly to the created object storage repository.

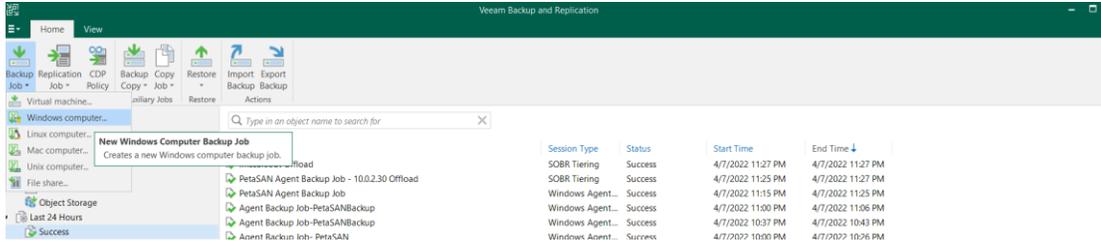
In case of Scale- out Backup Repository (SOBR)

You will create the job to connect to the scale-out repository that was created using (both the Direct Storage and object storage repositories as performance and capacity tiers respectively) so it will be executed using two sub jobs:

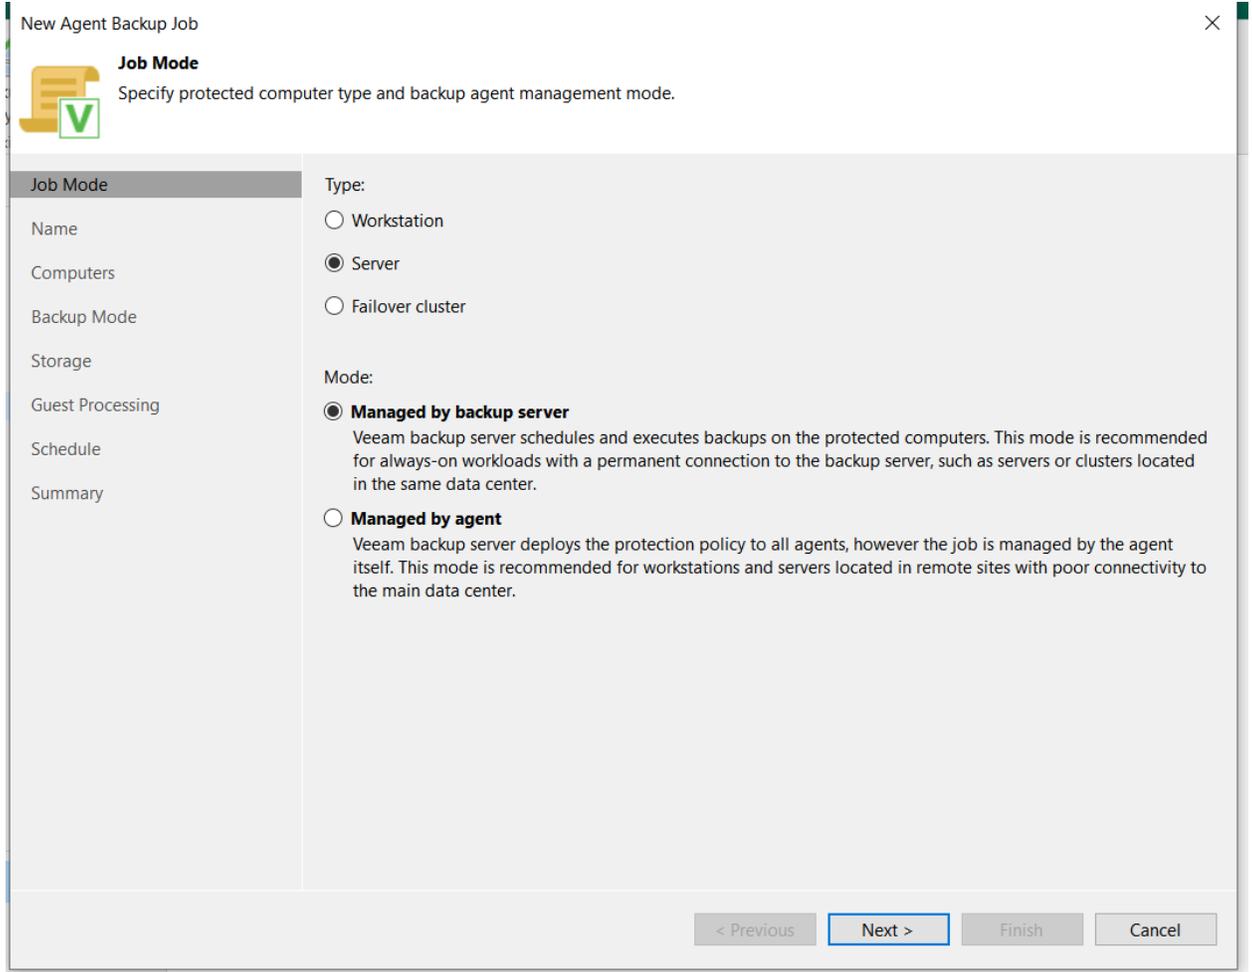
-Backup job

-Offload job

- Add new Backup job, in this example we will create Windows computer backup job to backup data from a Windows server with IP 10.0.2.30



- Select mode as server



- Enter the job name

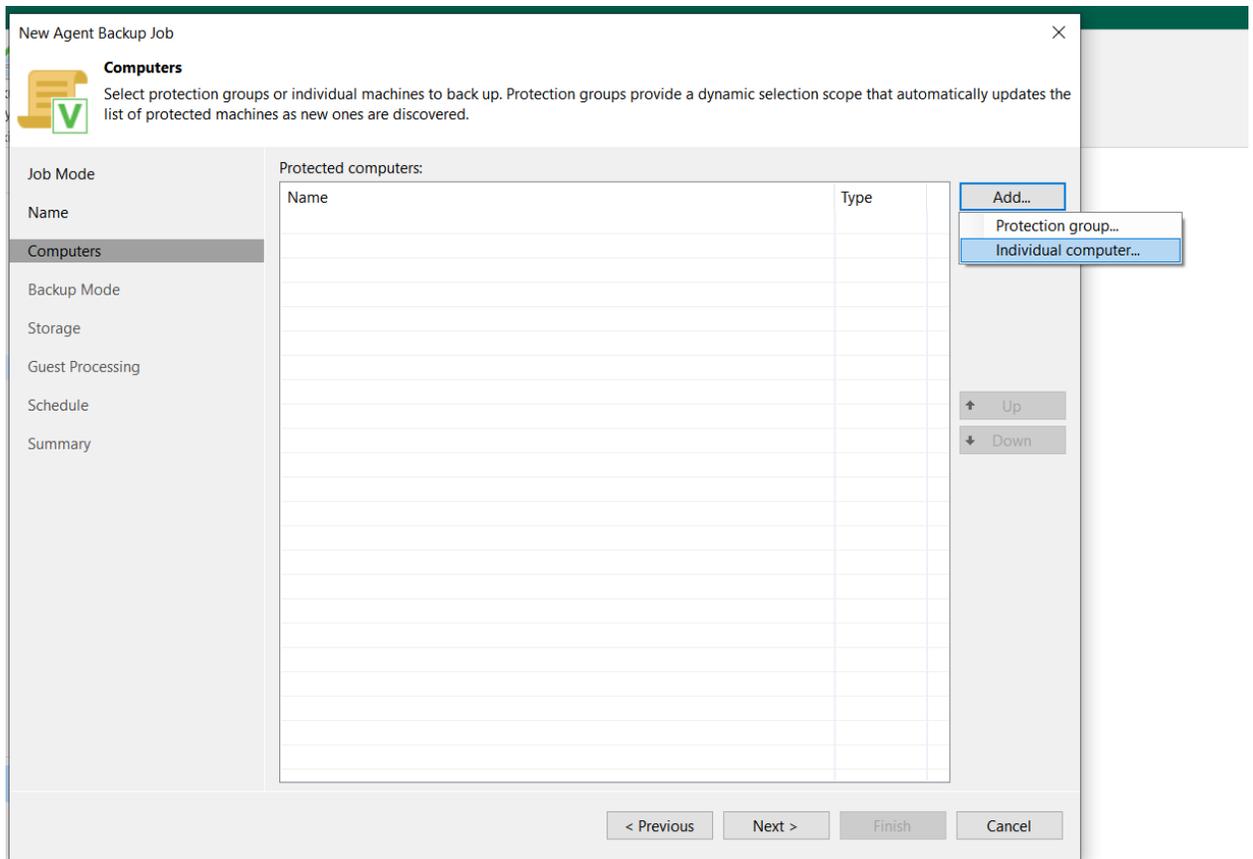
New Agent Backup Job ✕

 **Name**
Type in a name and description for this agent backup job.

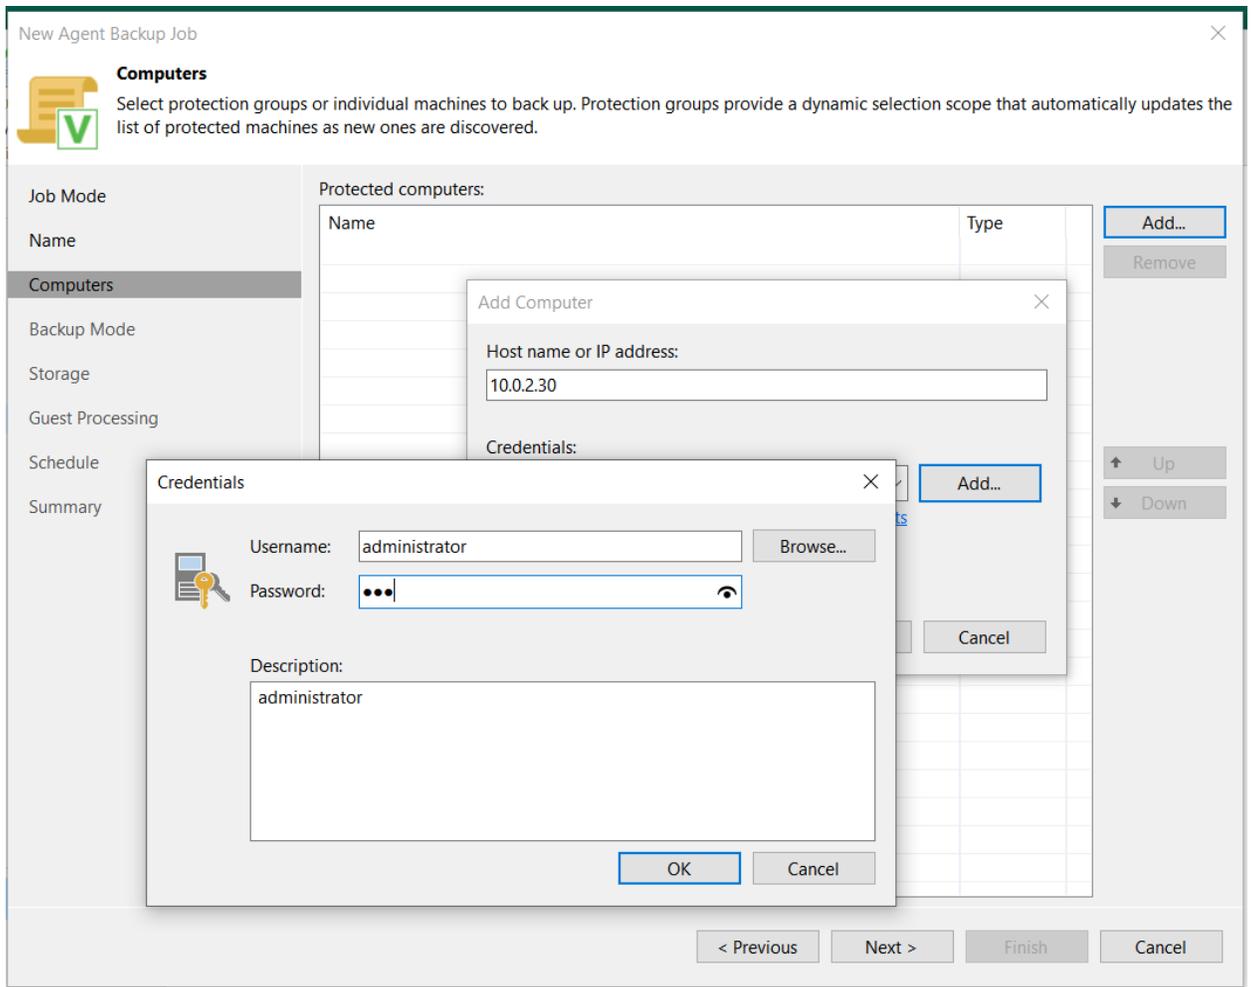
Job Mode	Name: <input type="text" value="PetaSAN Agent Backup Job"/>
Name	Description: <input type="text" value="Created by DESKTOP-95RGNP3\lenovo at 4/7/2022 11:11 PM."/>
Computers	
Backup Mode	
Storage	
Guest Processing	
Schedule	
Summary	

High priority
Backup infrastructure resources are offered to high priority jobs first. Use this option for jobs sensitive to the start time, or jobs with strict RPO requirements.

- Choose what will be backed up, in this example we will backup data on host 10.0.2.30



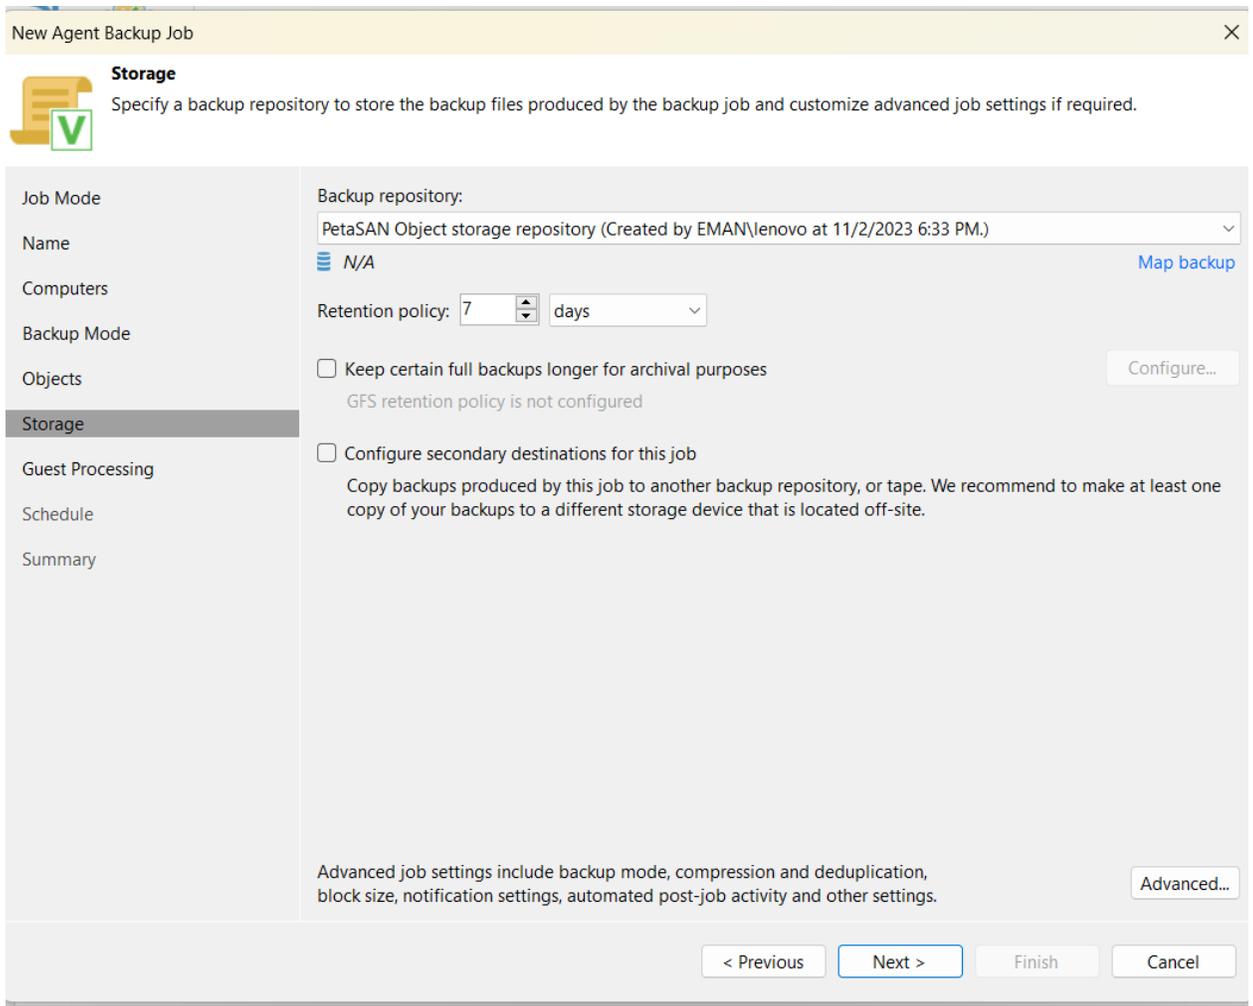
- Enter the authentication information of computer 10.0.2.30



New Agent Backup Job ✕

 **Backup Mode**
Choose what data you want to back up from selected computers.

Job Mode	<input type="radio"/> Entire computer Back up entire computer image for fast recovery on any level. Deleted, temporary and page files are automatically excluded from the image to reduce the backup size. <input type="checkbox"/> Include external USB drives
Name	
Computers	
Backup Mode	<input type="radio"/> Volume level backup Back up images of specified volumes, for example only data volumes. Deleted, temporary and page files are automatically excluded from the image to reduce the backup size.
Objects	<input checked="" type="radio"/> File level backup (slower) Back up selected files and directories only. This mode still produces an image-based backup, but only with protected file system objects included in the image.
Storage	
Guest Processing	
Schedule	
Summary	



But in case of Scale-Out Backup Repository (SOBR), you will select the Scale-Out repository

New Agent Backup Job ✕

 **Storage**
Specify a backup repository to store the backup files produced by the backup job and customize advanced job settings if required.

Job Mode	Backup repository:
Name	PetaSAN-Scale-out Backup Repository (Created by DESKTOP-95RGNP3\lenovo at 4/7/2022 11:05 PM.)
Computers	245 GB free of 466 GB Map backup
Backup Mode	Retention policy: 7 days
Objects	<input type="checkbox"/> Keep certain full backups longer for archival purposes Configure... GFS retention policy is not configured
Storage	<input type="checkbox"/> Configure secondary destinations for this job Copy backups produced by this job to another backup repository, or tape. We recommend to make at least one copy of your backups to a different storage device that is located off-site.
Guest Processing	
Schedule	
Summary	

Advanced job settings include backup mode, compression and deduplication, block size, notification settings, automated post-job activity and other settings. ⚙️ Advanced

< Previous Next > Finish Cancel

- Choose application processing options

New Agent Backup Job ✕

 **Guest Processing**
Choose application processing options.

Job Mode	<input checked="" type="checkbox"/> Enable application-aware processing Detects and prepares applications for consistent backup, performs transaction logs processing, and configures the OS to perform required application restore steps upon first boot.
Name	
Computers	Customize application handling options for individual machines and applications Applications...
Backup Mode	
Objects	<input type="checkbox"/> Enable guest file system indexing Creates catalog of guest files to enable browsing, searching and 1-click restores of individual files. Indexing is optional, and is not required to perform instant file level recoveries.
Storage	Customize advanced guest file system indexing options for individual machines Indexing...
Guest Processing	
Schedule	
Summary	

< Previous Next > Finish Cancel

- Select the schedule

New Agent Backup Job ✕

Schedule
Specify the scheduling options. If you do not set the schedule, the job will need to be controlled manually.

Job Mode Run the job automatically

Name Daily at this time: 10:00 PM Everyday

Computers Monthly at this time: 10:00 PM Fourth Saturday

Backup Mode Periodically every: 1 Hours

Objects After this job: Agent Backup Job 1 (Created by DESKTOP-95RGNP3\lenovo at 4/6/2022 1:12 AM.)

Storage **Automatic retry**

Retry failed items processing: 3 times

Wait before each retry attempt for: 10 minutes

Guest Processing **Backup window**

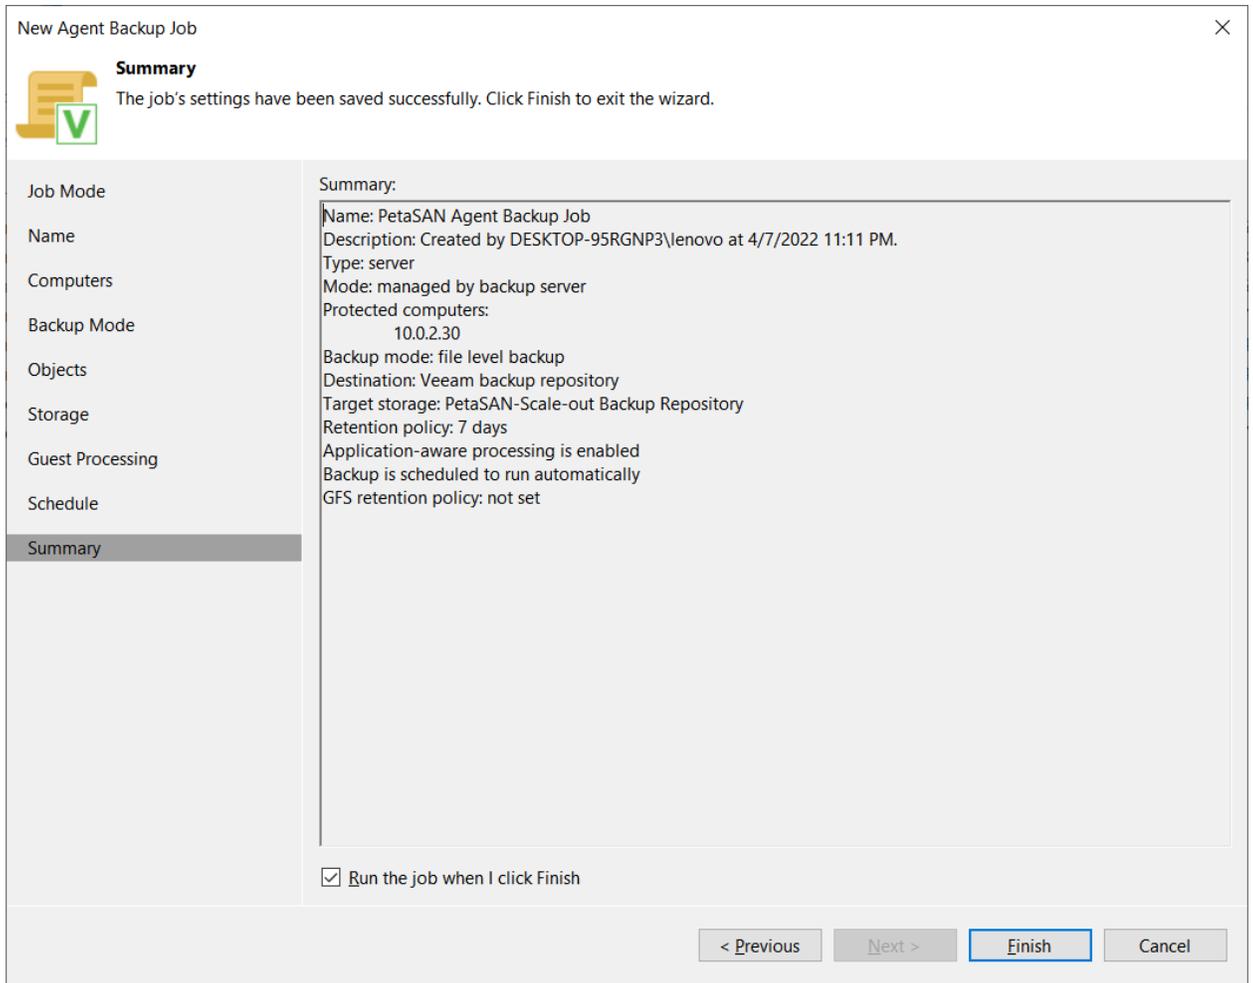
Terminate job if it exceeds allowed backup window

If the job does not complete within allocated backup window, it will be terminated to prevent snapshot commit during production hours.

Schedule

Summary

- Select to run the job after I click finish



- The job will start immediately
- The agent backup job will run directly in case of direct backup.

Agent Backup Job 4

Job progress: 100% 1 of 1 hosts

SUMMARY		DATA		STATUS	
Duration:	06:53	Processed:	19.9 GB (100%)	Success:	1
Processing rate:	1 MB/s	Read:	10.2 MB	Warnings:	0
Bottleneck:	N/A	Transferred:	10.2 MB (1x)	Errors:	0

THROUGHPUT (ALL TIME)

Name	Status	Action
DESKTOP-DME0...	Success	<ul style="list-style-type: none"> Creating VSS snapshot Calculating digests Getting list of local users (C): enumerating directories (C) (18.3 GB) 9.3 MB read at 6 MB/s Finalizing Truncating transaction logs Updating metadata Full backup created Network traffic verification detected no corrupted blocks Processing finished at 11/2/2023 7:48:46 PM

Speed: 0 KB/s

Hide Details OK

The agent backup job will run first then the offload job in case Scale-Out Backup (SOBR)

Veeam Backup and Replication

Home View Session Tools

Stop Statistics Report

Actions Details

Home

Jobs

- Backups
- Backups
- Disk
- Object Storage
- Last 24 Hours
- Success
- Warning
- Failed

Home Inventory Backup Infrastructure Storage Infrastructure History

Job Name ↑

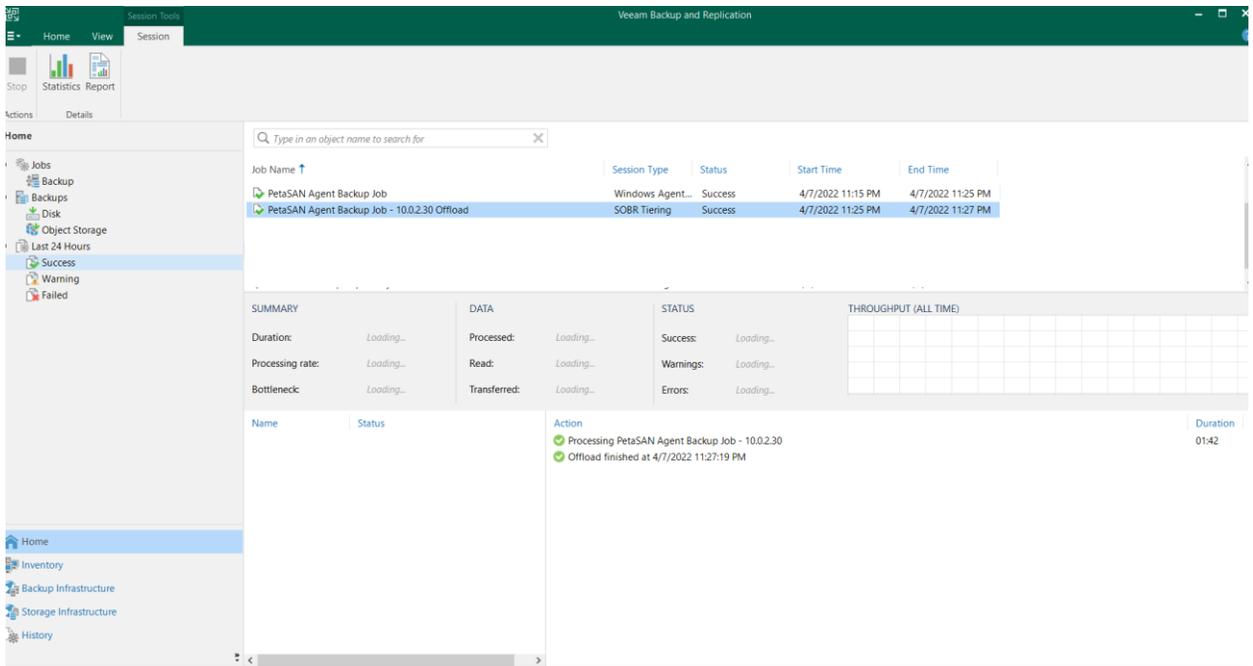
Job Name	Session Type	Status	Start Time	End Time
PetaSAN Agent Backup Job	Windows Agent...	Success	4/7/2022 11:15 PM	4/7/2022 11:25 PM
PetaSAN Agent Backup Job - 10.0.2.30 Offload	SOBR Tiering	Success	4/7/2022 11:25 PM	4/7/2022 11:27 PM

SUMMARY		DATA		STATUS	
Duration:	09:50	Processed:	18.4 GB (100%)	Success:	1
Processing rate:	927 KB/s	Read:	9 MB	Warnings:	0
Bottleneck:	N/A	Transferred:	9 MB (1x)	Errors:	0

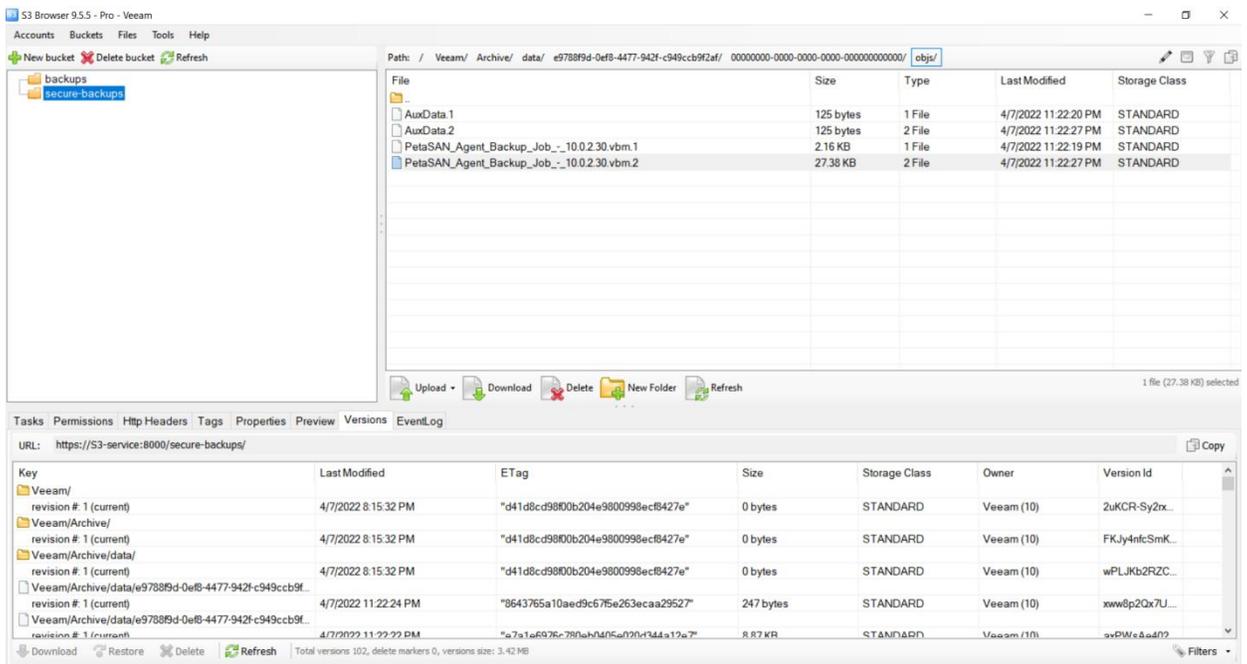
THROUGHPUT (ALL TIME)

Speed: 565 KB/s

Name	Status	Action	Duration
DESKTOP-DME0...	Success	<ul style="list-style-type: none"> Job started at 4/7/2022 11:15:24 PM Building hosts list Processing DESKTOP-DME0TUC Job finished at 4/7/2022 11:25:14 PM 	00:56 03:49



- You will notice that data has been backed up in buckets that we selected



Note:

You may notice that you can delete objects from the bucket even though it is immutable, don't worry they are there but just hidden. If you go into the bucket and select show versions button, the objects will show up.

S3 Browser 9.5.5 - Pro - Veeam

Accounts Buckets Files Tools Help

New bucket Delete bucket Refresh

backups
secure-backups

Path: / Veeam/ Archive/ data/ e9788f9d-0ef8-4477-942f-c949ccb9f2af/ 00000000-0000-0000-0000-000000000000/ obj/

File	Size	Type	Last Modified	Storage Class
AuxData.1	125 bytes	1 File	4/7/2022 11:22:20 PM	STANDARD
AuxData.2	125 bytes	2 File	4/7/2022 11:22:27 PM	STANDARD

Upload Download Delete New Folder Refresh

2 files (250 bytes) and 0 folders

Tasks Permissions Http Headers Tags Properties Preview Versions Event Log

URL: https://S3-service:8000/secure-backups/

Key	Last Modified	ETag	Size	Storage Class	Owner	Version Id
revision # 1 (current)	4/7/2022 11:22:27 PM	"73cfd5a3274a123772bf4d0202c984ebd"	125 bytes	STANDARD	Veeam (10)	ts0p.TQ9Vao...
Veeam/Archive/data/e9788f9d-0ef8-4477-942f-c949ccb9f2af/revision # 2 (deleted)	4/8/2022 12:48:25 AM				Veeam (10)	USOP3bvX66...
revision # 1	4/7/2022 11:22:19 PM	"fa6fe4d3e739953b1680cac7600d7c73"	2.16 KB	STANDARD	Veeam (10)	DsOJOBU4o...
Veeam/Archive/data/e9788f9d-0ef8-4477-942f-c949ccb9f2af/revision # 2 (deleted)	4/8/2022 12:48:25 AM				Veeam (10)	fhH8nAyrTp...
revision # 1	4/7/2022 11:22:27 PM	"e6e0deeeaffd183659f140eae8d3df56"	27.38 KB	STANDARD	Veeam (10)	kXM1ISEYFc...
Veeam/Archive/data/e9788f9d-0ef8-4477-942f-c949ccb9f2af/revision # 1 (current)	4/7/2022 11:22:25 PM	"3d7b884257eaf01e70ce3a7d0b276d3"	588 bytes	STANDARD	Veeam (10)	XUUKsynLSK...
Veeam/Archive/data/e9788f9d-0ef8-4477-942f-c949ccb9f2af/revision # 2 (deleted)						

Download Restore Delete Refresh Total versions 102, delete markers 2, versions size: 3.42 MB

Filters