



VAST DATA UNIVERSAL STORAGE WITH VEEAM® BACKUP & REPLICATION

Configuration Guide

VERSION 1.0



TABLE OF CONTENTS

Revision History	2
Introduction	3
Veeam Configuration	4
Adding a vCenter Server	5
Adding an NFS Share	8
Veeam Transport Modes Discussion	10
Automatic Selection	10
Direct NFS Access Mode	12
VAST Cluster as Backup Repositories	14
Creating an NFS Backup Repository	14
VAST NFS Repository Settings	19
Creating an S3 Backup Repository	20
Creating an S3 User on VAST	20
Creating an S3 Bucket on VAST	22
Create Veeam Backup S3 Repository	23
Creating a Scale-Out Backup Repository	26
Configuring Veeam Backup Jobs	30
Veeam Backup Job to a VAST NFS Repository	30
Veeam Backup Jobs to an S3 Repository	34
Backup Job for NFS Share onto VAST S3	34
Backup job for SOBR with VAST S3	36
Restoring From a Veeam Backup Job	37
Restoring VMs From a VAST NFS Repository	37
Restoring Files to NFS File Share	39
Restoring Data From an S3 Repository	41

REVISION HISTORY

Revision Number	Description	Date
1.0	Initial Release	May 2022



INTRODUCTION

VAST Data's Universal Storage, for the first time, redefines the economics of flash storage, making flash affordable for all applications, from the highest performance databases to the largest data archives. The Universal Storage concept blends game-changing storage innovations to lower the acquisition cost of flash with an exabyte-scale file and object storage architecture breaking decades of storage tradeoffs. Veeam®

With the advantage of new, enabling technologies that weren't available before 2018, this new Universal Storage concept can achieve a previously impossible architecture design point. The system combines low-cost hyperscale flash Drives and Storage Class Memory with stateless, containerized storage services all connected over new low-latency NVMe over Fabrics networks to create VAST's Disaggregated Shared Everything (DASE) scale-out architecture. Next-generation global algorithms are applied to this DASE architecture to deliver new levels of storage efficiency, resilience, and scale.

Veeam Backup & Replication 11 is a sophisticated data protection and disaster recovery solution. With Veeam Backup & Replication, you can create image-level backups of virtual, physical, cloud machines and restore from them instantly. Intelligent technology used in the product optimizes data transfer and resource consumption providing superior recovery time and point objectives as well as helping to minimize storage costs. Veeam Backup & Replication provides a centralized console for administering backup/restore/replication operations in all supported platforms (virtual, physical, cloud).

VAST Clusters with Veeam Backup & Replication jointly deliver an integrated, simple, and powerful data protection solution for today's complex virtualized environments. Veeam's Backup & Replication software takes advantage of the industry-leading inline deduplication and compression technology built into the VAST Cluster. In addition, the VAST Cluster introduces Similarity-Based Data Reduction which rethinks data reduction algorithms to deliver unprecedented storage efficiency. When combined, the result is a fast and compact backup, providing a high availability and site recovery solution. This joint solution delivers dramatic reductions in the cost and complexity of protecting virtual server infrastructure while meeting increasingly aggressive SLAs with the power of flash.

This document covers basic configuration steps to integrate Veeam Backup and Replication with a VAST Data Cluster. It covers specific Veeam architecture basics, minimal click VAST Cluster configuration for both NFS and S3 and highlights recommended practices for protecting virtual machines in VMware vSphere environments using VAST clusters and Veeam's Backup & Replication solution.

This guide will assist individuals who are responsible for the design and deployment of data protection and disaster recovery solutions of virtual machines deployed on a VAST Cluster.



VEEAM CONFIGURATION

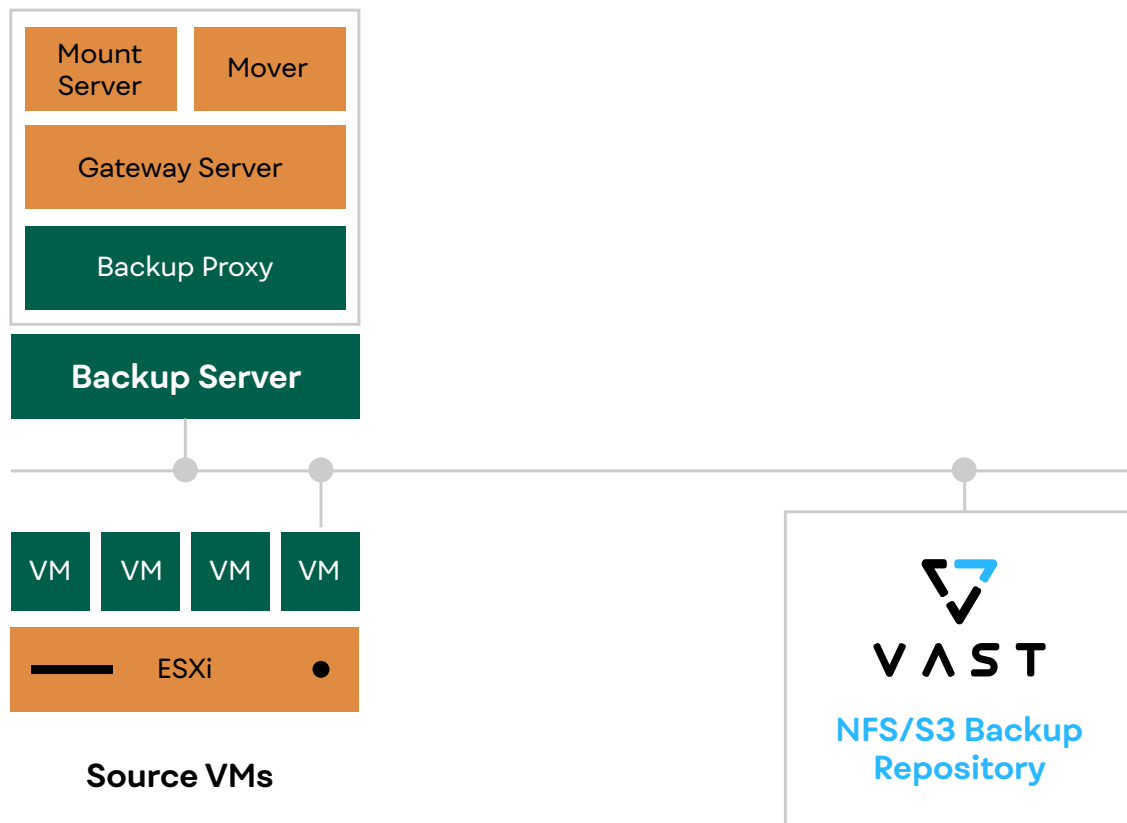


Figure 1 - Simple NFS/S3 Deployment

This document was written using a basic deployment of Veeam Backup and Replication where all the components were installed and running on the backup server (Figure 1). For a detailed discussion of how to expand and scale a Veeam deployment refer to the [Veeam Backup and Replication documentation](#).

One of the first tasks for implementing a backup with Veeam Backup and Replication, is to define what needs to be protected. Two things will be added as data sources to be backed up, a Managed Server and an NFS file share. The managed server could be a standalone ESXi host, a supported hypervisor cluster, or a typical windows or Linux host with associated datastores. For this example, the two will be a vSphere vCenter server and an NFS share from a second VAST cluster.



ADDING A VCENTER SERVER

With the Veeam Backup and Replication Console open and the **Inventory** tab on the left pane selected right click on **Virtual Infrastructure** and select **Add Server** (Figure 2). There are several ways to access this menu, but this is meant to be the most intuitive.

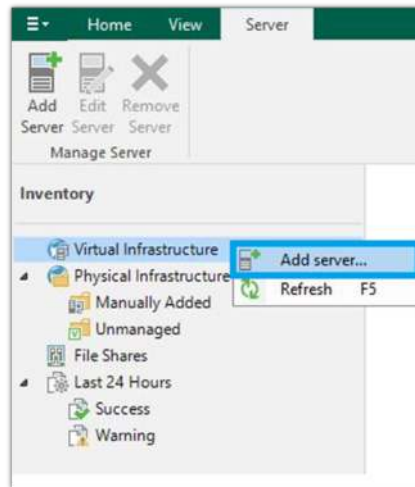


Figure 2 – Adding vSphere vCenter Server to Veeam

This will pop up an **Add Server** (Figure 3) window where VMware vSphere should be selected.

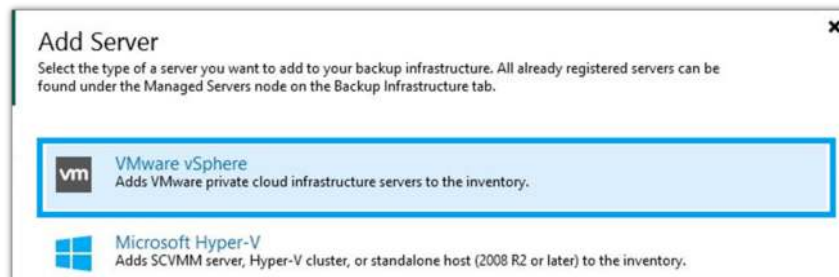


Figure 3 – Picking Hypervisor Server

On the next window select **vSphere** to add the vCenter Server to add to Veeam (Figure 4).

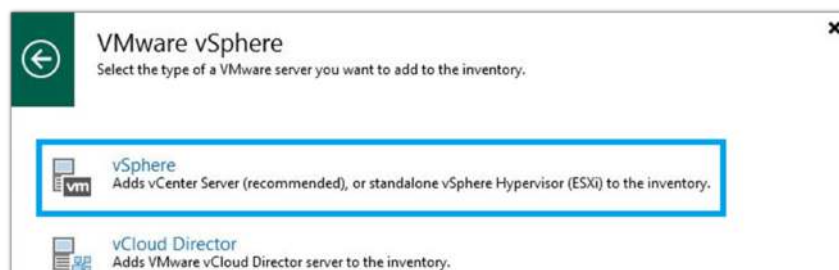
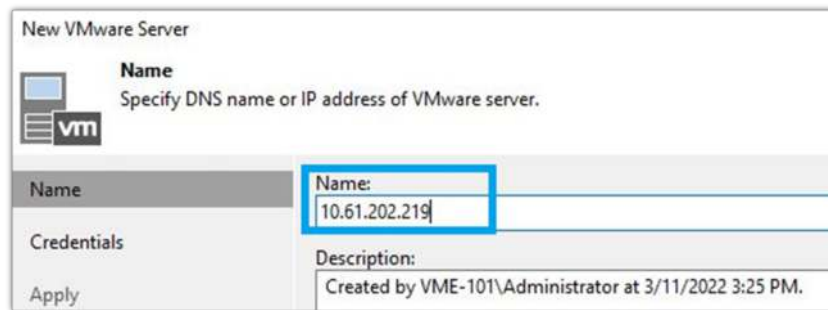


Figure 4 – Select vSphere vCenter Server



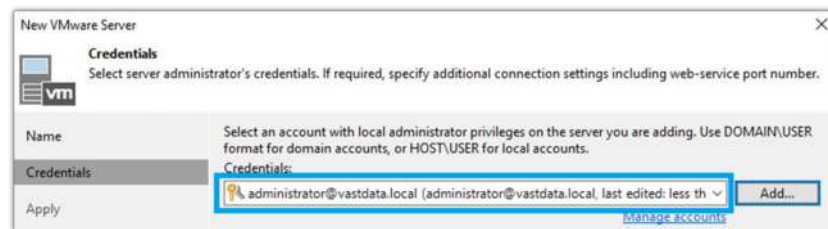
On the next window select vSphere to add the vCenter Server to add to Veeam (Figure 4).



The 'New VMware Server' dialog box is shown. The 'Name' field is highlighted with a blue box and contains the IP address '10.61.202.219'. The 'Description' field contains the text 'Created by VME-101\Administrator at 3/11/2022 3:25 PM.'.

Figure 5 - Adding VMware Server

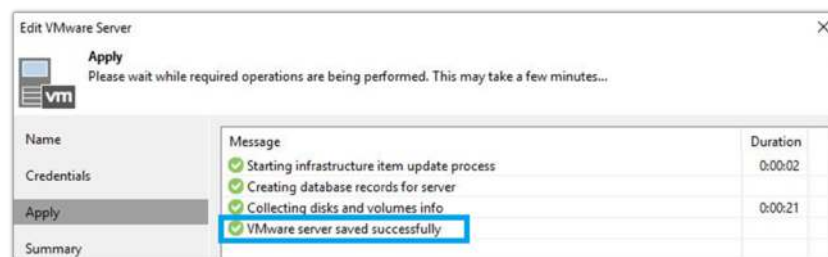
After clicking **Next**, the **Credentials** window will appear (Figure 6). The user credentials for the VMware Server may already be in the pull-down menu but if not then click Add to and enter the user and password for that vCenter Server.



The 'New VMware Server' dialog box is shown with the 'Credentials' tab selected. The 'Name' field is highlighted with a blue box and contains the text 'administrator@vastdata.local (administrator@vastdata.local, last edited: less th...'. The 'Add...' button is also highlighted with a blue box.

Figure 6 - Add vCenter Server Credentials

When finished click **Apply** and Veeam will begin the process of importing all the information about the vCenter Server including all its VMs (Figure 7).



The 'Edit VMware Server' dialog box is shown. The 'Apply' button is highlighted with a blue box. The 'Message' field contains the text 'Starting infrastructure item update process'. The 'Duration' field contains the text '0:00:02'. The 'Summary' field contains the text 'VMware server saved successfully'.

Figure 7 - Import of vCenter Server



When the scan finishes simply click Next and review the Summary screen (Figure 8). If any changes need to be made, click **Back** otherwise click **Finish**. The vCenter Server is now added to Veeam as a managed server.

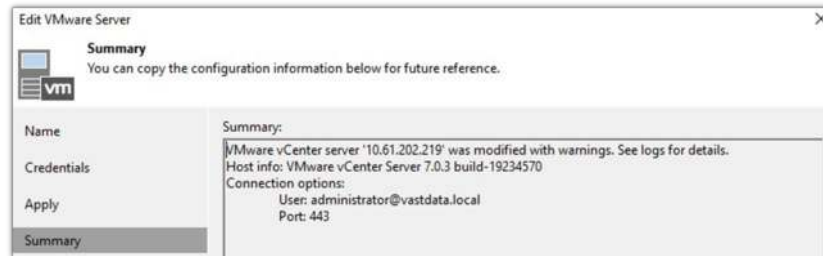


Figure 8 - Adding vCenter Server Summary

The vCenter server should now show under the Virtual Infrastructure list (Figure 9).

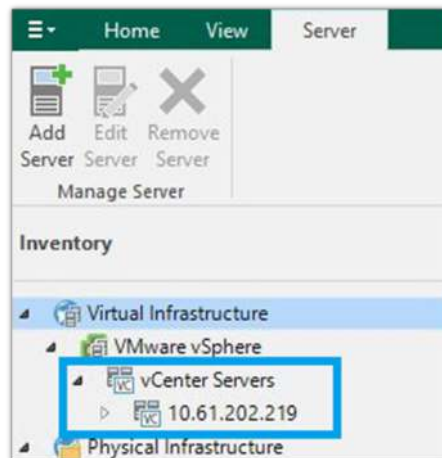


Figure 9 - vCenter Server Successfully Added to Veeam



ADDING AN NFS SHARE

An example used later in this document needs to backup data from an NFS share. This shows how to add a VAST NFS share as the source of data to be backed up. From the **Inventory** section right click on **File Shares** and select Add file share (Figure 10).

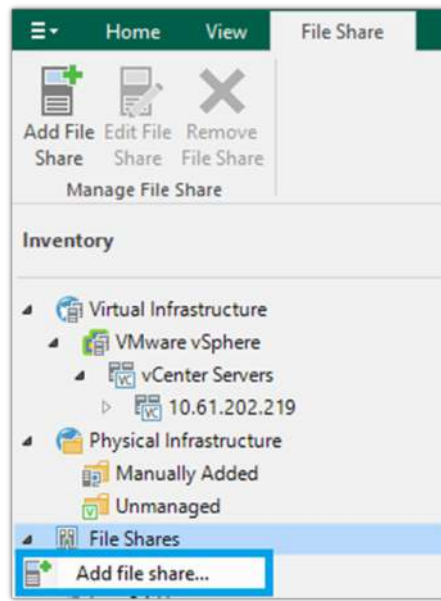


Figure 10 – Select Add file share

The next window (Figure 11) presents several options for adding a file share, select NFS Share.

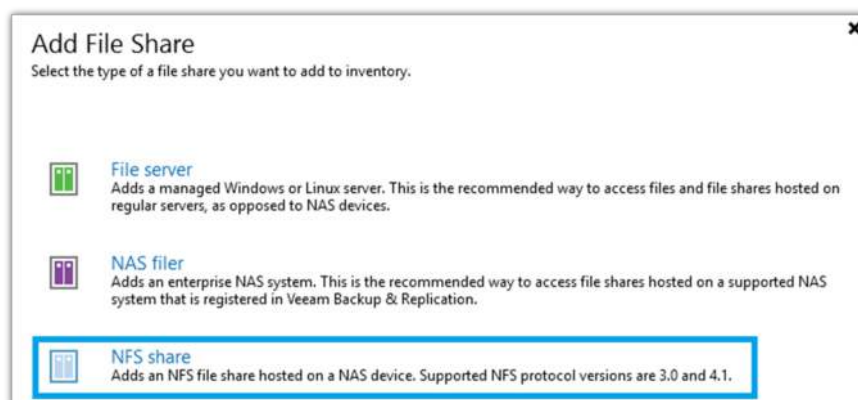


Figure 11 – Add NFS File Share



Add the file share in the standard server:/folder format as shown in Figure 12.




Figure 12 - Add file share specifics

The rest of the settings are left at their defaults for simplicity. Simply click through the rest of the settings and review on the last window. Once complete the NFS share with the source data should show up as in Figure 13.

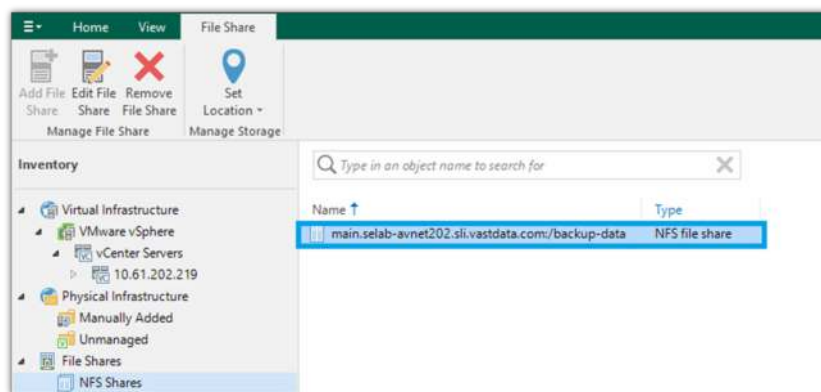


Figure 13 - NFS Share Added



VEEAM TRANSPORT MODES DISCUSSION

This section will cover a brief discussion on Veeam Backup & Replication transport modes, and how they're used to efficiently backup NFS datastores.

As part of the method that protects and backups data, Veeam copies data from a source to a backup repository through what is called a transport mode. It uses the Veeam Data Mover (Figure 1) to retrieve VM data from the source and write VM data to the target using one of three transport modes.

- Direct Storage Access
- Virtual Appliance (HotAdd)
- Network (LAN Transport)

Automatic Selection

By default, Veeam's Backup & Replication will use an automatic backup proxy transport selection. In this mode, the backup proxy and the connected NFS datastores are analyzed to determine the most efficient transport mode. If several transport modes are available for the same backup proxy, Veeam Backup & Replication will choose the mode in the order that is listed above.

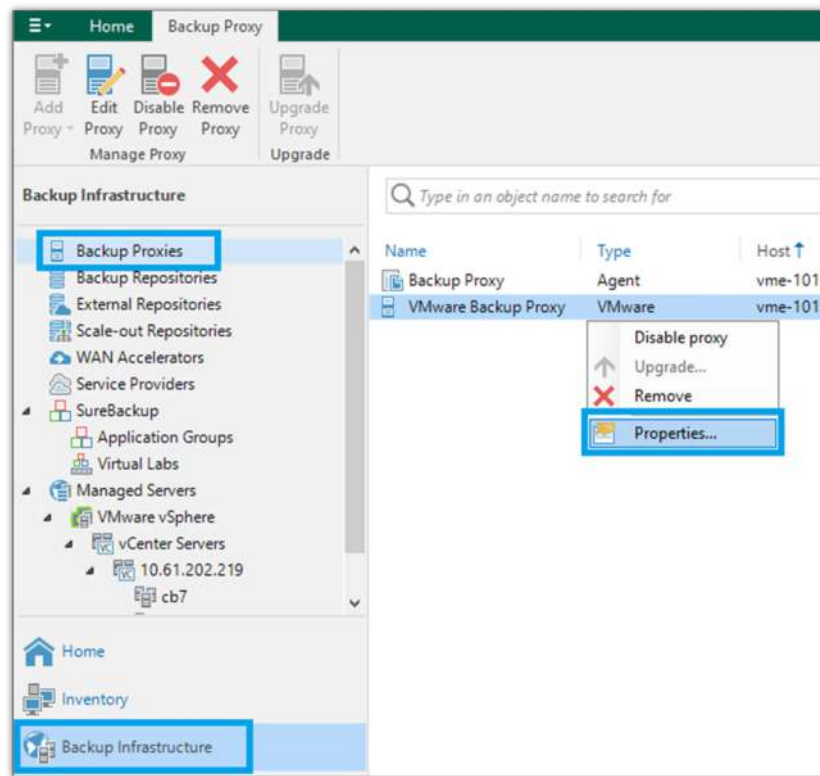


Figure 14 - Selecting the Backup Proxy



The default mode can be altered by editing the properties of the backup proxy. To edit the mode, click on **Backup Infrastructure** in the left pane and then click on **Backup Proxies**. On the right windowpane right click on the **VMware Backup Proxy** and select **Properties** (Figure 14).

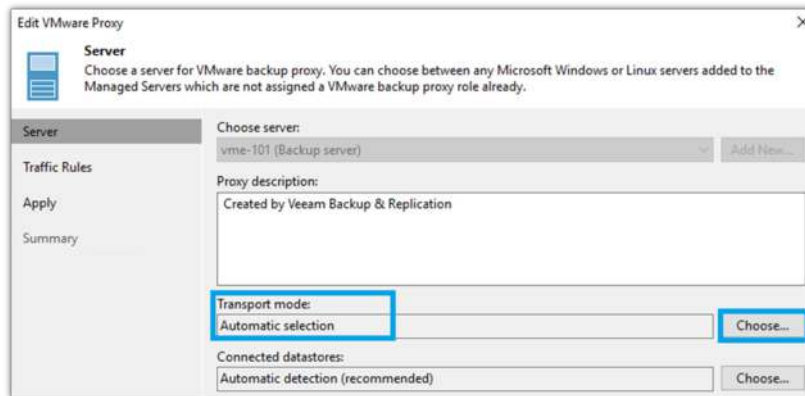


Figure 15 – Transport Mode

The mode is then modified by clicking on **Choose** (Figure 15) which will bring up the screen in Figure 16. The default settings shown are recommended unless a specific transport mode is needed. However, even if a specific mode is desired it, automatic mode will use that mode if it's the best mode available. It just simplifies the decision of which mode to select.

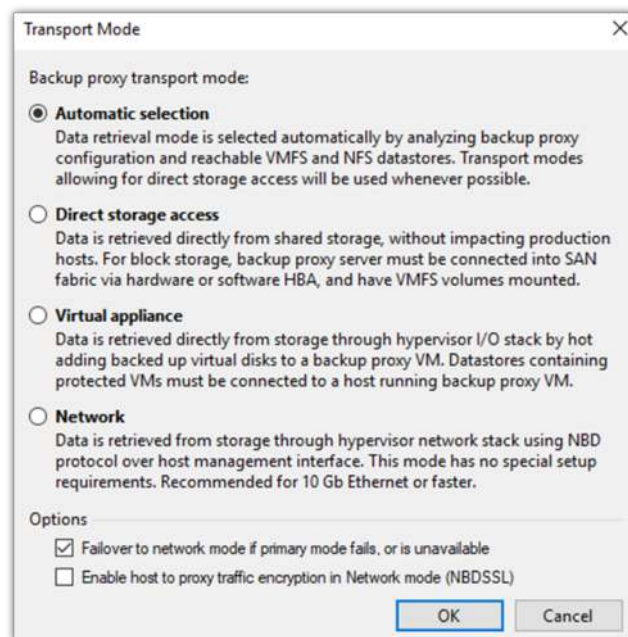


Figure 16 – Modifying the Transport Mode



Direct NFS Access Mode

Prior to the introduction of this transport mode, data on NFS datastores could only be backed up using Virtual Appliance or Network Transport Mode. These modes place a heavy load on the LAN and use VMware VDDK to communicate with the ESXi host which produces additional load. With Direct NFS access mode, Veeam Backup & Replication bypasses the ESXi host by deploying its native NFS client on the backup proxy(ies) and uses it for VM data transport. This way there is no additional load on the ESXi host.

Direct NFS Access mode is ideal for the scenario where both the source VM data and the backup repository reside on NFS shares. Veeam Backup & Replication will read and write data directly from an NFS source to an NFS backup repository and as mentioned bypass the ESXi host thereby offloading the workload from that host.

This is the most efficient of the three transport modes. To use this transport mode the backup proxy must have access to the source NFS datastore where the VMs exist.

VAST AS A SOURCE NFS DATASTORE

The focus of this document is around using the VAST Cluster NFS shares as a target for backup repositories on Veeam. However, if VAST is being used as the source datastore for the VMs there are a couple of prerequisites to ensure that dNFS transport mode will function properly. To read and write data in the Direct NFS transport mode, the backup proxy must meet the following requirements:

1. The backup proxy(ies) must have access to the NFS datastore(s) where the VM disks are located.
2. The backup proxy(ies) must have Read/Write permissions and root access to the NFS datastore.

You can also choose networks over which Veeam Backup & Replication must transport data when you perform data protection and disaster recovery tasks. This option, a preferred network, can be helpful if you have a non-production network and want to route data traffic over this network instead of the production one. A preferred network is not specific to direct NFS but can ensure its use in the transport decision hierarchy.

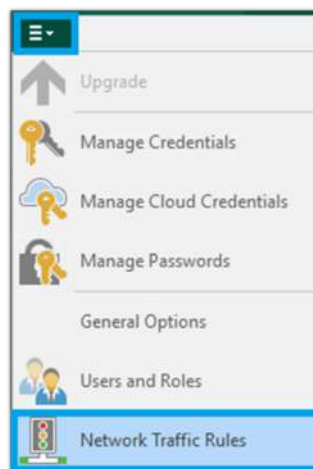


Figure 17 - Configure Network Traffic Rules



To configure a preferred network in the Veeam UI click on the icon at the top with the three lines and select **Network Traffic Rules** (Figure 17). In the Global Network Traffic Rules window click on the Networks button (Figure 18). Then add the networks that contain the NFS repositories and Veeam Backup proxies.

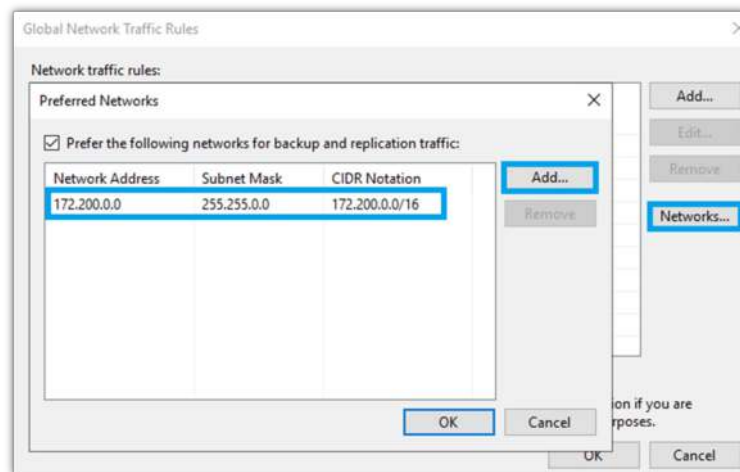


Figure 18 - Setting Preferred Network



VAST CLUSTER AS BACKUP REPOSITORIES

A backup repository is a location used by Veeam Backup & Replication jobs to store backup files, VM copies, and metadata for replicated VMs. The following sections provide the appropriate steps to configure the VAST cluster and deploy it as both an NFS and S3 backup repository on Veeam.

CREATING AN NFS BACKUP REPOSITORY

It's simple to create an NFS share and present it to Veeam. From the VAST UI Dashboard click on **Views** (Figure 19).



Figure 19 - Selecting Views to Create NFS Directory

This brings up the Element Store window with the Views tab selected. In the upper right corner select Create View (Figure 20).



Figure 20 - Select Create View

The Add View window (Figure 21) opens with several fields to fill in. Starting with the Path, enter a path that adheres to any organizational structure that may exist. Under **Protocols**, this example uses NFS but NFSv4 is also supported.

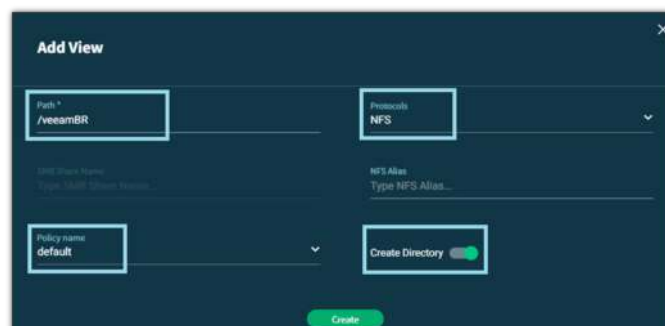


Figure 21 - Adding View Parameters



The **Policy Name** here is set to **default**, which is configurable based on a user's needs. In general, policies define the protocol(s) available, group membership, read/write and squash rules to name a few. Configuring **View Policies** is not discussed here but Figure 22 shows how this specific default policy is defined. The Read/Write access and no squash are wide open using the asterisk as a wild card.



Name	↑	Atime Freq...	Flavor	Group Mem...	Read Write	Read Only	All Squash	No Squash	Root Squash	Trash	Actions
default			Select fla...		Search...	Search...	Search...	Search...	Search...	Search...	
default			NFS	Client	*						

Figure 22 - Default View Policy

An alternate mount point can be defined with the **NFS Alias** setting but was left blank in this example. Since this is a new view the **Create Directory** toggle has been turned on.

NFS exports have the familiar server:/folder format for mounting purposes. The VAST cluster creates access to the server through the concept of Virtual IPs. With a minimum four nodes in a VAST cluster multiple IPs can be grouped together into a Virtual IP Pool to allow for better distribution of connections and for better performance and resiliency through multipathing. To that end the full path of the export (server:/folder) in this example is defined as:

```
main.selab-avnet202.sli.vastdata.com:/veeamBR
```

This will be used in the next step within the Veeam UI. For additional information on configuring and understanding virtual IPs see the VAST Cluster Administrator's Guide.




Figure 23 - Add Backup Repository

To add the export, in the Veeam UI, select the category Backup Infrastructure and then right click Backup Repositories. Click on Add backup repository (Figure 23). This will bring up the **Add Backup Repository** window where **Network attached storage** should be selected (Figure 24). If the VAST export or share is mounted on a separate Linux or Windows system (not discussed in this guide) then direct attached storage would be selected.

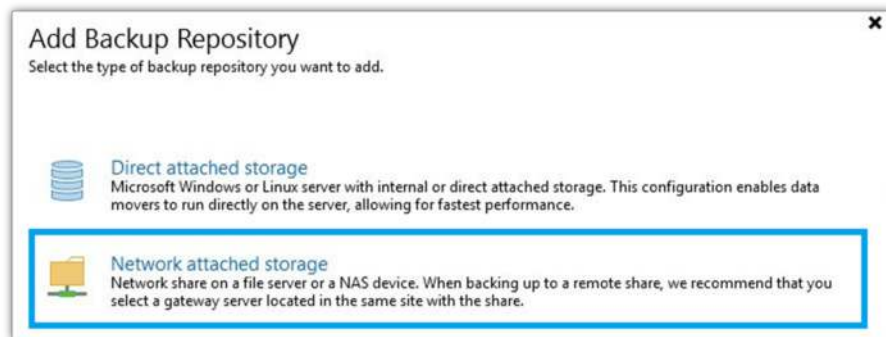


Figure 24 - Select Network Attached Storage

In the next window simply click NFS share (Figure 25).

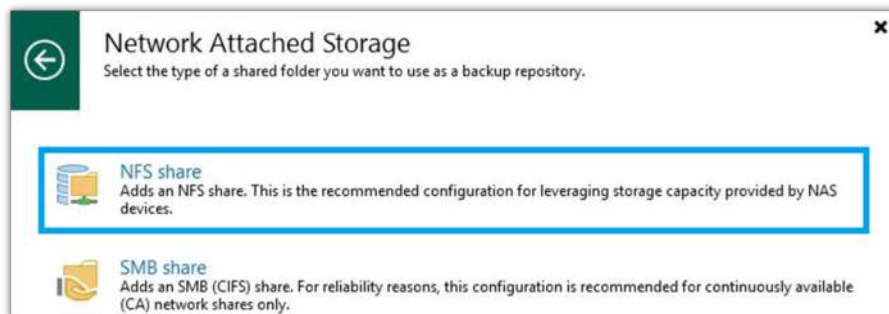


Figure 25 - Select NFS Share

The next screen (Figure 26) defines the repository name and optional description.

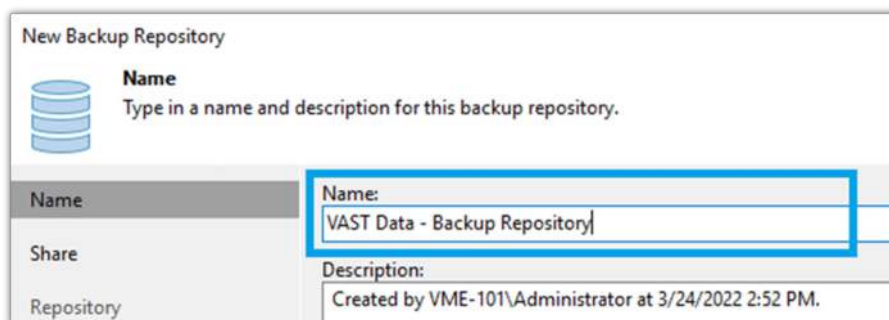
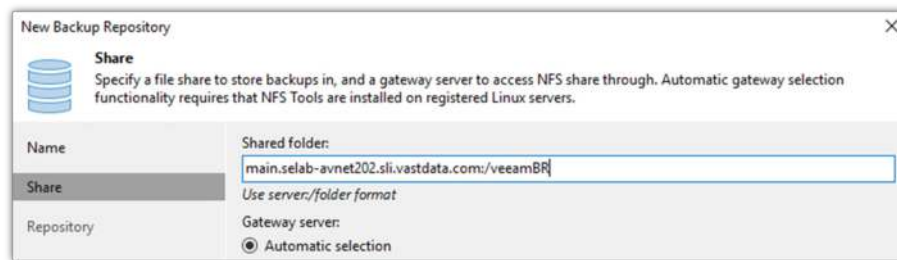


Figure 26 - Define the Repository Name

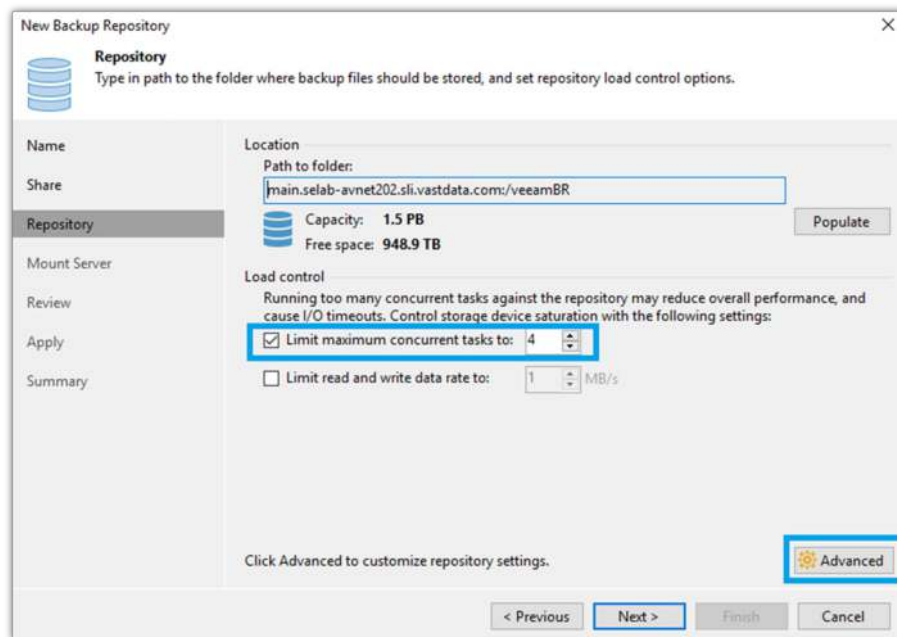
On the next screen (Figure 27) enter the full path (server:/folder) of the view that was discussed at the beginning of this section. This is a simplistic deployment with the gateway server on the same server as Veeam so that setting is left as automatic.



The 'New Backup Repository' window is shown with the 'Share' tab selected. The 'Name' field is empty. The 'Share' field contains the path 'main.selab-avnet202.sli.vastdata.com:/veeamBR'. The 'Repository' field is empty. The 'Gateway server' section has 'Automatic selection' selected.

Figure 27 - Enter Shared folder

Continuing with the Repository settings in Figure 28 the maximum concurrent tasks may have to be limited based on the resources available in the Veeam server. Refer to Veeam documentation for better guidance on this. Click on the **Advanced** button.



The 'New Backup Repository' window is shown with the 'Repository' tab selected. The 'Name' field is empty. The 'Share' field contains the path 'main.selab-avnet202.sli.vastdata.com:/veeamBR'. The 'Repository' field is empty. The 'Mount Server' field is empty. The 'Review' field is empty. The 'Apply' field is empty. The 'Summary' field is empty. The 'Location' section shows 'Path to folder:' as 'main.selab-avnet202.sli.vastdata.com:/veeamBR', 'Capacity:' as '1.5 PB', and 'Free space:' as '948.9 TB'. The 'Load control' section has 'Limit maximum concurrent tasks to:' checked and set to '4'. The 'Limit read and write data rate to:' is unchecked and set to '1 MB/s'. The 'Advanced' button is highlighted. The 'Next >' button is also highlighted.

Figure 28 - Recommended Settings for VAST Repositories

The Storage Compatibility Settings window will appear (Figure 29). Ensure the boxes are checked as shown. Both settings, as described, will improve backup performance.

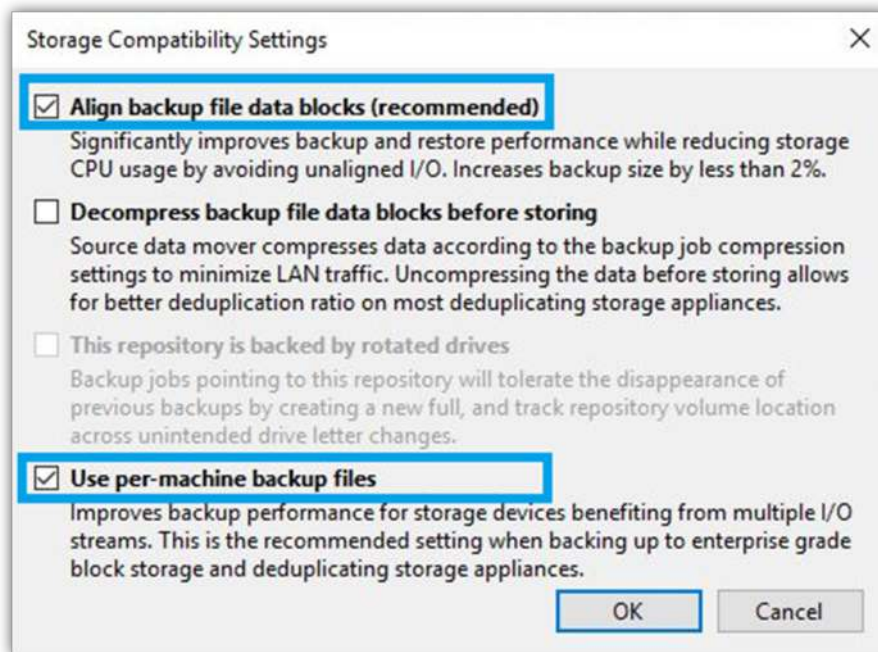


Figure 29 – Storage Compatibility Settings

The next window (Figure 30) is used to define the mount server where backup files are mounted to this server to allow for file recovery. There is a mount server associated with every backup repository. With VAST Cluster NFS backup repository it's recommended to also enable the vPower NFS service to allow for instant recovery of VMs.

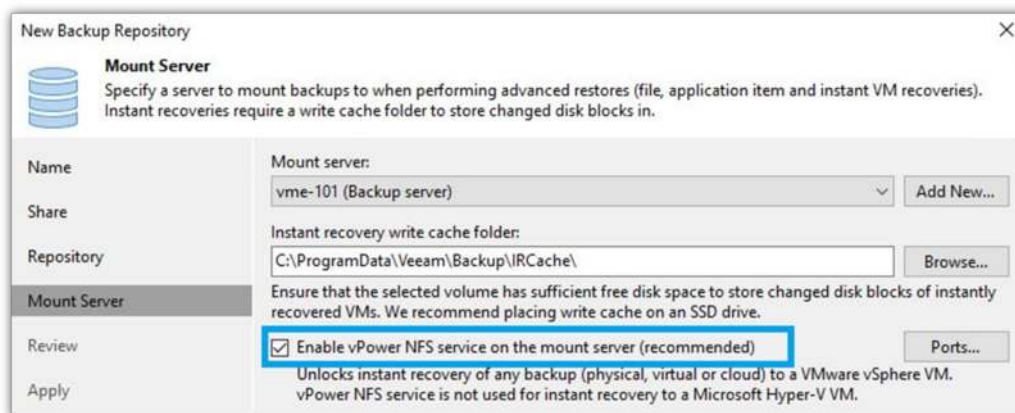


Figure 30 – Mount Server Settings

On the Review window simply verify all settings and then click **Next**. The next window (Figure 31) will show several processes being performed, installation of services and adding of the backup repository just to name a few.

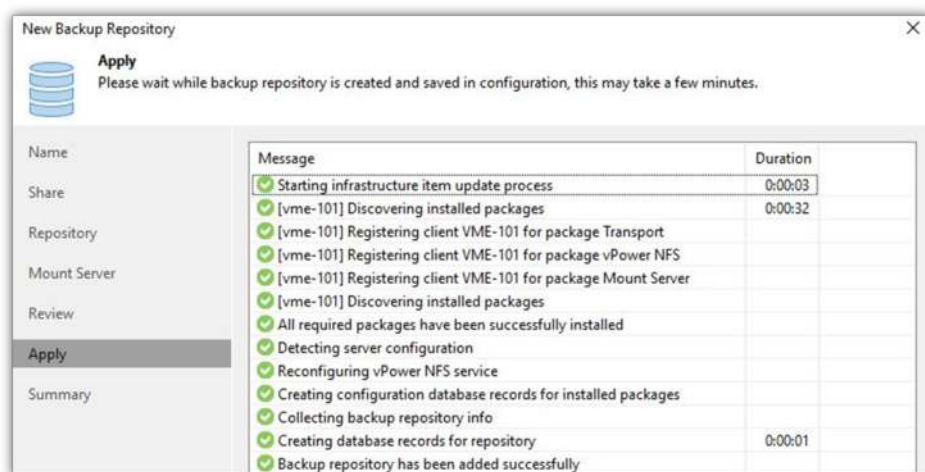


Figure 31 - Applying All settings and Creating Backup Repository

Review the **Summary** page (Figure 32) and click **Finish**. The VAST Cluster NFS backup repository is now configured and ready to be used to create a backup job.

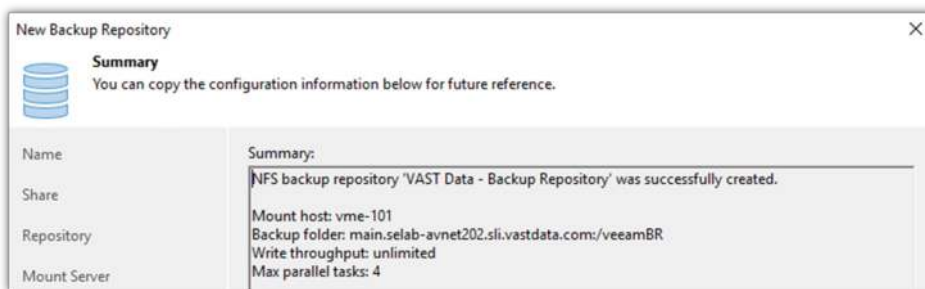


Figure 32 - Review the Summary Page

Veeam performs a rescan of hosts and servers every four hours so if the newly created backup repository is not showing a manual rescan can be performed.

VAST NFS Repository Settings

The following table is a summary of settings that are recommended when configuring the VAST cluster as an NFS backup repository.

Transport Mode	Automatic or Direct NFS (Production is on NFS)
Storage Compatibility	Enable the following backup repository settings: <ul style="list-style-type: none">Align Backup file-data blocksUse per-VM backup files
Mount Server	Enable the vPower NFS service on the mount server

Figure 33 - Backup Repository Settings



CREATING AN S3 BACKUP REPOSITORY

During the process of adding an S3 backup repository the credentials of the S3 will be needed so the creation of an S3 user discussed first in the next section.

Creating an S3 User on VAST

To create an S3 user with the VAST UI, go to the Dashboard and then **User Management** and select **Users** (Figure 34).

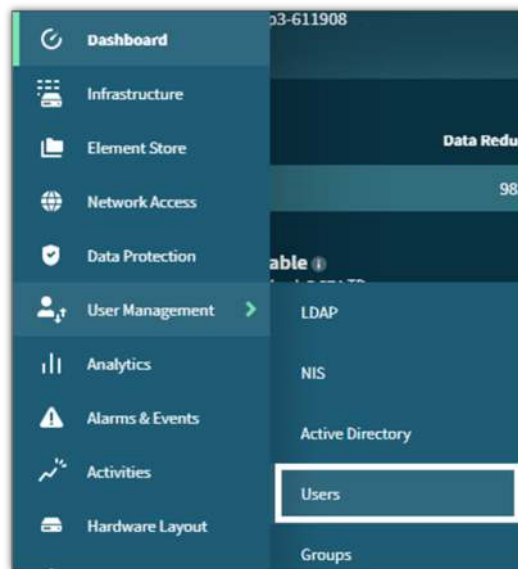


Figure 34 - Selecting Users Category

The **Add User** window appears (Figure 35) and a username is entered along with any desired UID. This particular user is granted full S3 credentials but that is not a requirement during the Veeam process of adding an S3 repository.

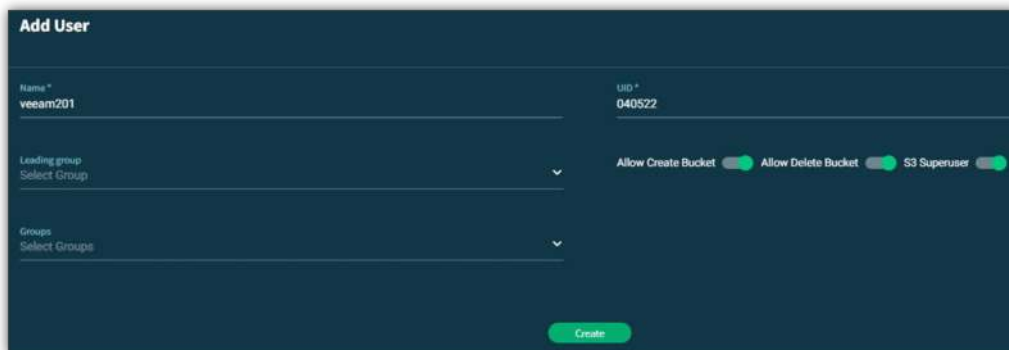


Figure 35 – Creating a New User for S3



After clicking **Create** the user will show up in the User Management window Figure 36.

User Management										
LDAP	NIS	Active Directory	Users	Groups						
Name	UID	SID	Leading Group	Leading Group GID	Groups	Group Count	Create Bucket	Delete Bucket	S3 Superuser	
veeam201	40522	S-1-111-2899...				0	Yes	Yes	Yes	

Figure 36 - User Created

Now that the user is created it needs to be edited to create and capture the active and secret keys. On the far right of the user under the action column click on the three dots and select **Edit** (Figure 37).

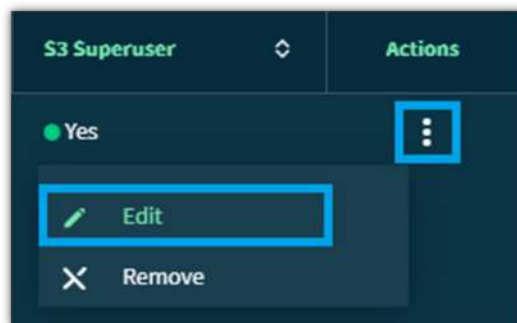


Figure 37 - Editing a User

In the **Update User** window (Figure 38) click on the **Create new key** button to create a new **Access Key**. Now copy the active key and secret key some place for a bit later. This is the only time when the secret key will be obtainable so be sure to copy it down in some safe location. Click Update to close the window.

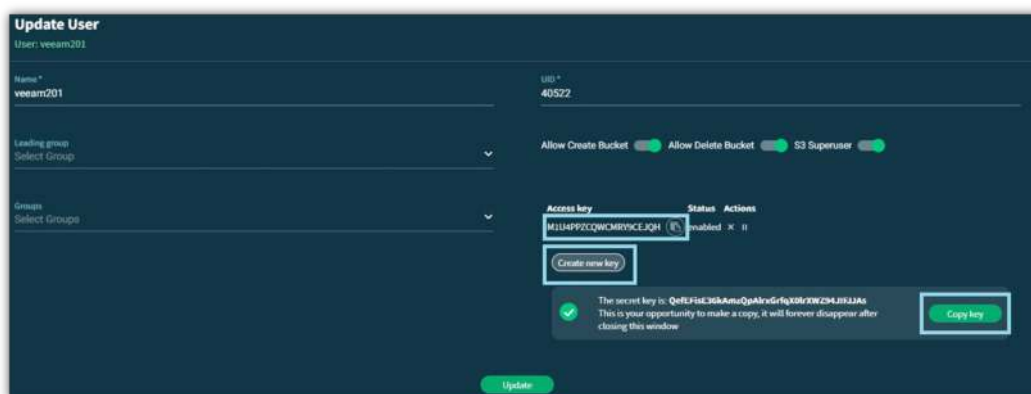


Figure 38 - Capturing Active and Secret Keys

The user will now be used in the creation of a bucket or view within the VAST UI.



Creating an S3 Bucket on VAST

Just as before, when creating a view, from the dashboard go to **Views** as shown in Figure 19 and then click Create as shown Figure 20.

This view is being configured as a bucket only using just the S3 Bucket protocol. Multiple protocols (NFS, S3) can be used on a view.

Figure 39 shows the all the settings needed to create the new bucket. Since this is a new view the create directory toggle is selected.

Add View

General S3

Path *
/veeam201

Protocols *
S3 Bucket

S3 Bucket Name *
vast201

Policy name *
s3_default_policy

Create Directory ☒

Create

Figure 39 - S3 Bucket Configuration

The user created previously needs to be given access or ownership of this view. This is done on the S3 tab as shown in Figure 40. When finished click Create from either tab.

Add View

General S3

S3 Bucket Owner *
veeam201

S3 Versioning ☐

Anonymous access ☐

Bucket Owners (Users)
Type...

Bucket Owners (Groups)
Type...

Create

Figure 40 - Adding Bucket Owner (User)



The view (bucket) now shows up in the Element Store as shown in Figure 41.



Path	Alias	Share	Bucket	Bucket Owner	Policy name	Protocols	Physical Capacity	Logical Capacity	Actions
/	/				default	NFS	4.511 TB	15.981 TB	⋮
/veeam201			vast201	veeam201	s3_default_policy	S3 BUCKET	0 Bytes	0 Bytes	⋮

Figure 41 – S3 Bucket Created

With the bucket now created it's time to create the backup repository within the Veeam UI.

Create Veeam Backup S3 Repository

Similar to creating an NFS backup repository in the Veeam UI, select the category **Backup Infrastructure** and then right click **Backup Repositories** and then select **Add backup repository** (Figure 23).

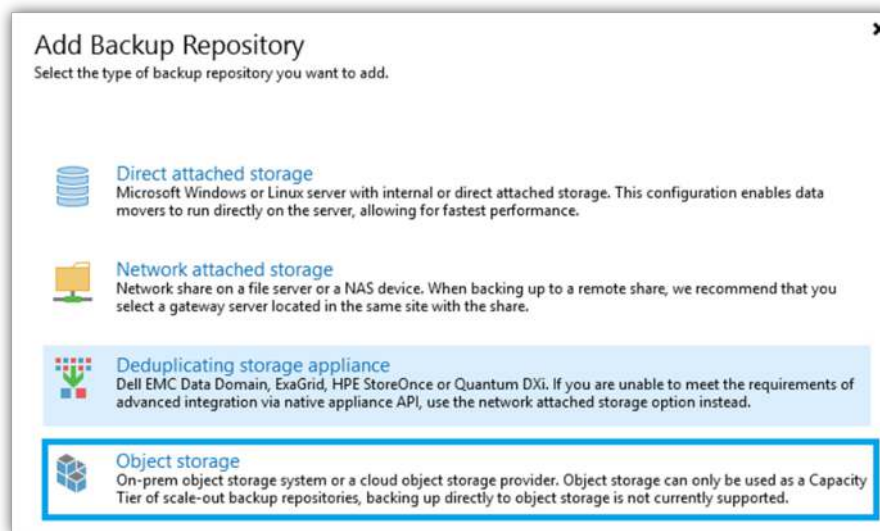


Figure 42 – Choose Object Store

This will bring up the **Add Backup Repository** window where **Object Storage** should now be selected (Figure 42). On the next screen select **S3 Compatible** (Figure 43).



Figure 43 – Select S3 Compatible



On the next window (Figure 44) fill in an appropriate name for this S3 repository and any description desired.

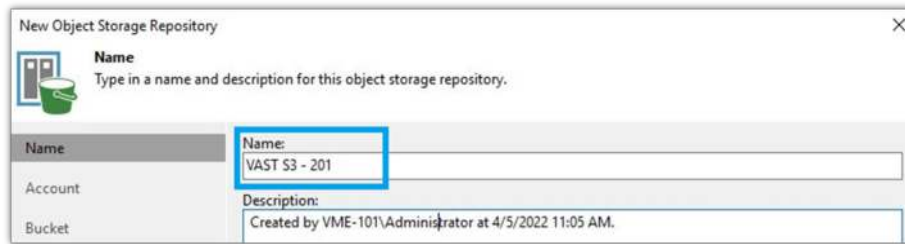


Figure 44 - Add Name to S3 Object Storage

There is an additional setting for concurrent tasks (not shown in Figure 44) which was left unchecked because there is no need to limit the number of tasks or API requests to the VAST cluster. The VAST bucket information is entered on the next screen as show in Figure 45. The service point is ascertained in the same manner as described in the section **Creating an NFS Backup Repository**.

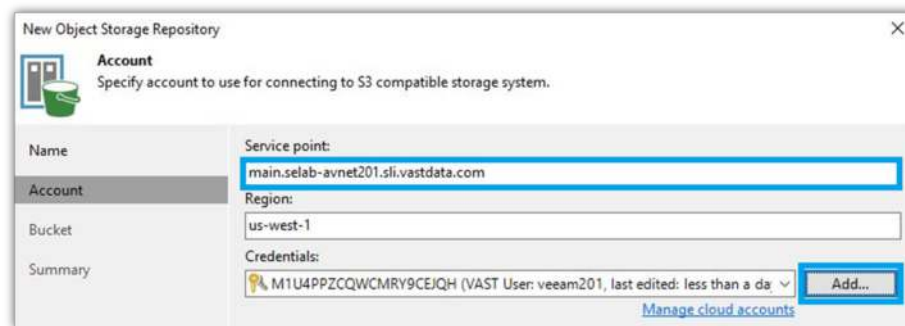


Figure 45 - Enter Service Point and Credentials

If the credentials have not been previously added then they'll need to be added by clicking the Add button and entering the access and secret keys that were created previously (Figure 46). It may be useful to add the VAST username to the description to keep track of multiple cloud accounts.



Figure 46 - Adding S3 User Credentials



Click next when all the fields have been filled in. Veeam now reaches out to the VAST cluster with the defined S3 user credentials and scans for a list of accessible buckets. Go ahead and accept the certificate warning if it appears.

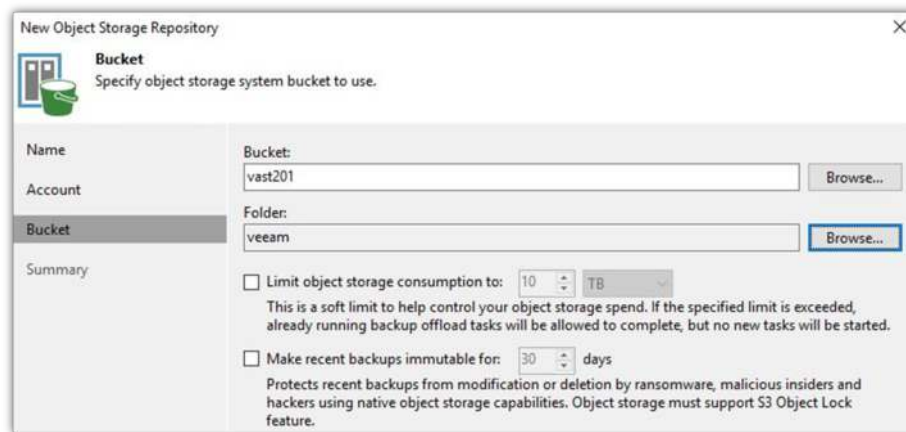


Figure 47 – Specify Bucket and Folder

The next window (Figure 47) is for bucket and folder selection. Starting with the bucket click the **Browse** button to bring up the **Select Bucket** window (Figure 48). Select the bucket (view) that was just created.




Figure 48 – Select VAST Bucket

After selecting the appropriate bucket click on the **Browse** button for the **Folder** (Figure 47). This brings up a list of folders (Figure 49) within the bucket. Either select an existing folder or create a new one.




Figure 49 – Select or Create a Folder



When all fields are filled out as in Figure 47 click next and review the **Summary** page (Figure 50).

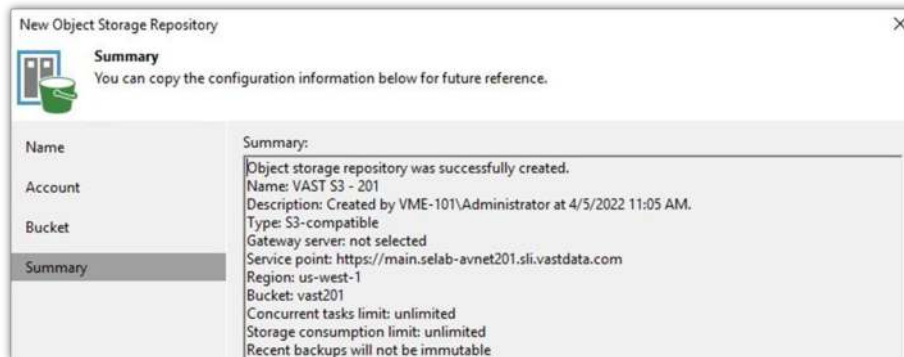


Figure 50 – Review Summary of S3 Repository

After finishing the wizard the S3 repository should appear in Veeam under Backup Repositories (Figure 51).

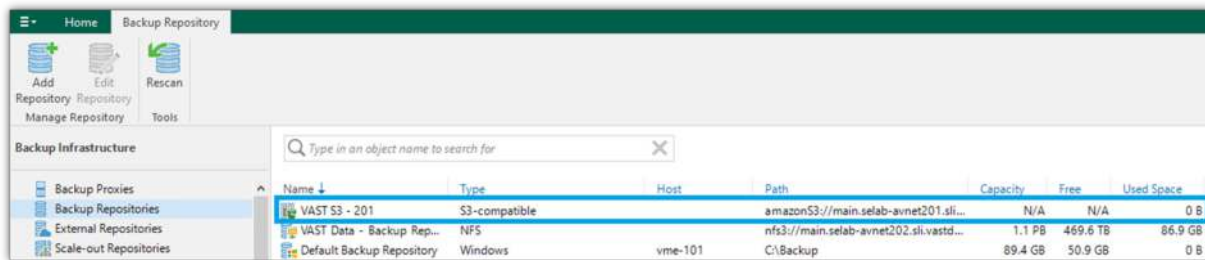


Figure 51 – VAST S3 Repository Created

Creating a Scale-Out Backup Repository

Scale out backup repository (SOBR) is a Veeam concept that combines unlimited extents of performance layer storage together along with an optional archive layer. This archive layer is where an S3 bucket from VAST participates (but could be used in the performance layer as well).

Veeam combines all components into a singular entity so that all the sub-parts are no longer usable in other backup jobs. Veeam is essentially creating a tiered storage system with each tier being separately scalable.



To create a SOBR, under **Backup Infrastructure** right click **Scale-out Repositories** and select **Add**.




Figure 52 – Select Add Scale-Out Repository

Give the SOBR a name and an optional description (Figure 53).




Figure 53 – Give the SOBR a name

The first layer of SOBR is the performance tier and Figure 54 shows one extent has already been added to that tier. In this example, VAST is doing double duty with a second flash VAST cluster being used for the performance tier. Like all tiers in a SOBR construct they can be scaled with additional extents. So, multiple extents could be added as needed however, the VAST cluster technology can scale infinitely so realistically only a single VAST extent is needed here drastically simplifying the implementation and management of SOBR.

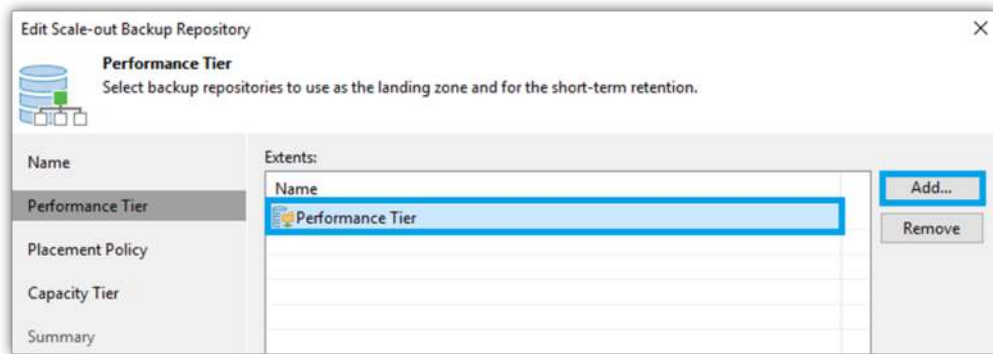


Figure 54 - Performance Tier

The Placement Policy of data onto each extent must be considered (Figure 55). VAST recommends using Data locality since as an all-flash system no additional performance is achieved spreading data across multiple extents. And again, with VAST's ability to scale performance only a single extent is ever really needed which again simplifies the choice here.

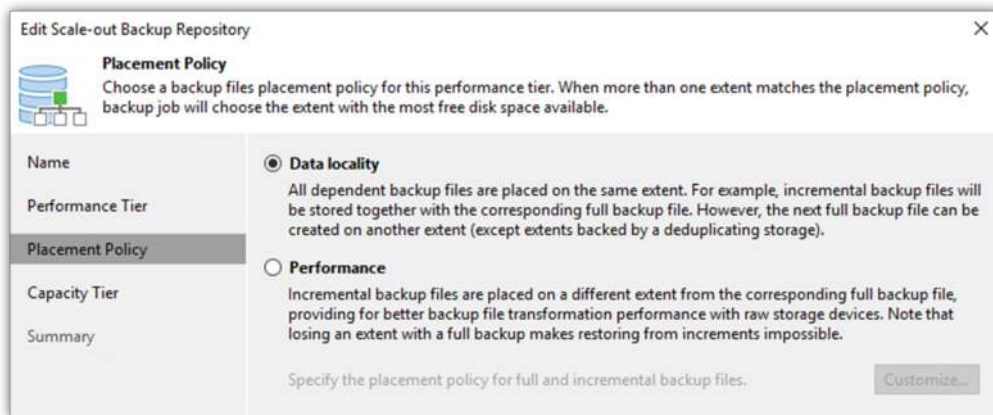


Figure 55 - Placement Policy Considerations

The next layer in a SOBR construct is the capacity layer which the redundant or disaster recovery copy of the data and is the VAST S3 repository. Figure 56 shows the **Capacity Tier** window with the top box checked which ungrays the pull down menu. Since the VAST S3 repository was already created, the item is listed and selected from the pull-down menu. If it wasn't created or if an additional capacity extent is needed then click on the Add button.

When and how often to initiate the backups to the S3 capacity is left to the user however, if possible, it is recommended to check the first box and copy to S3 as soon as the backup on the performance tier is created.

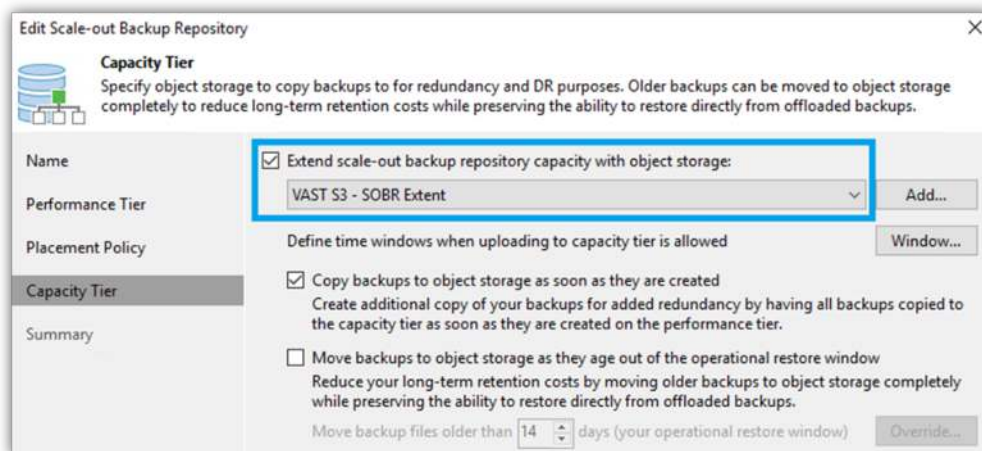


Figure 56 – Capacity Tier Options

On the summary page (not shown) simply review all settings and click Finish. After the scale-out backup repository is created it will show up as in Figure 57.

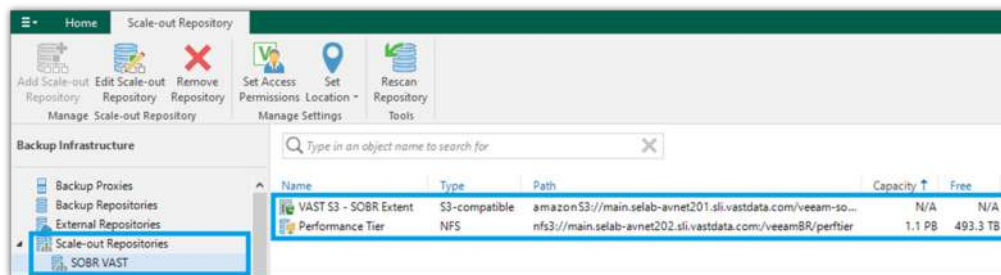


Figure 57 – SOBR Created

Notice in the right pane with the SOBR selected that both tiers are listed – performance (NFS) and capacity (S3). Now that these repositories are apart of a scale-out object repository they will no longer be available for any other function within Veeam.



CONFIGURING VEEAM BACKUP JOBS

A Veeam backup job defines how, where, and when to backup VM data. This section covers creating backup jobs for both a VAST NFS and a VAST S3 repository.

VEEAM BACKUP JOB TO A VAST NFS REPOSITORY

There are myriad ways to initiate creating a new backup job. This example will be creating a backup for a few virtual machines and the first step starts in the Home section of Veeam. Right click **Jobs** in the left pane and select **Backup** and then **Virtual Machine** (Figure 58).

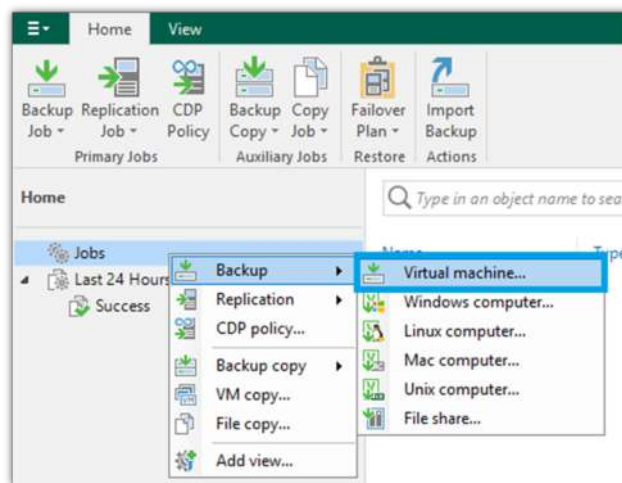


Figure 58 - Selecting Virtual Machine Backup Job

In the **Name** window give the backup job an appropriate name (Figure 59).




Figure 59 - Naming the Backup Job



In the Virtual Machines window click on the Add button (Figure 60).

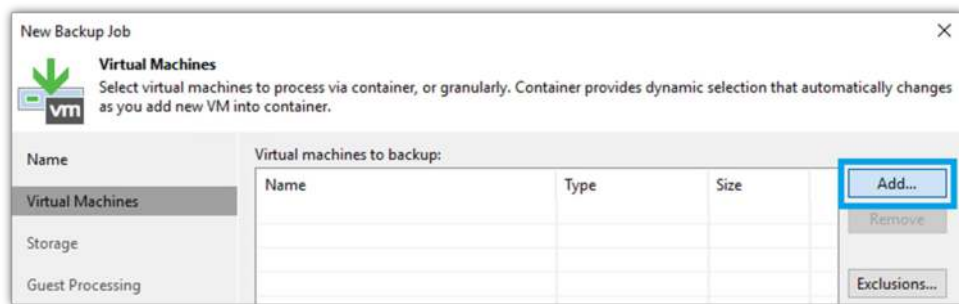


Figure 60 - Add VMs to the Backup Job

This will bring up a browse window for selecting the virtual machines (Figure 61). The virtual machines can be individually picked out of the list or a convenient search can be done to speed up selection.

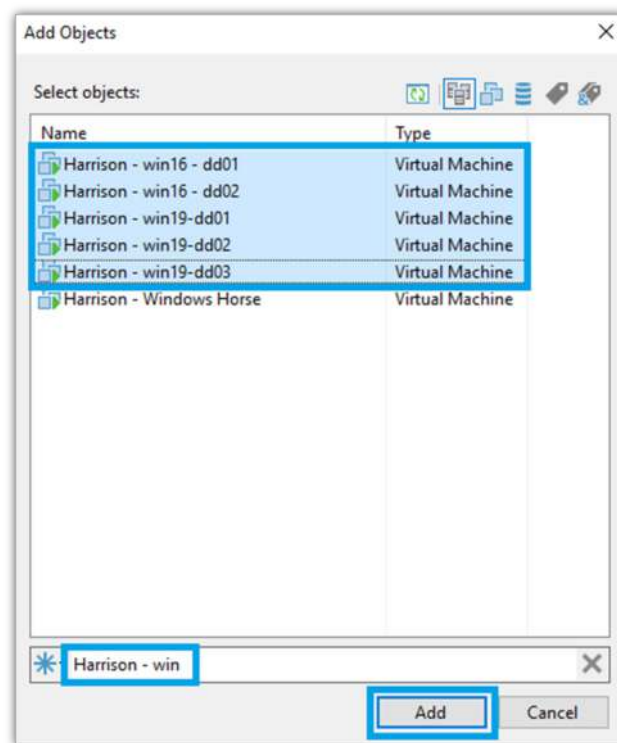


Figure 61 - Selecting the VMs with a Filter



Click Add when all of the desired virtual machines are selected. Review the list of virtual machines to be added to the backup and make any adjustments including exclusions (Figure 62).

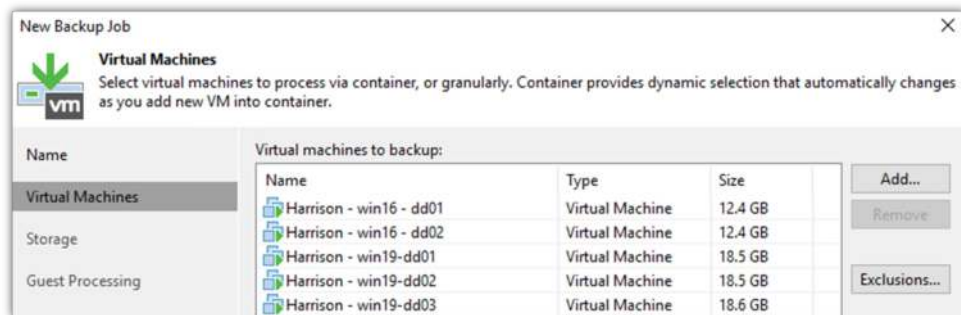


Figure 62 - VMs Selected

On the **Storage** window (Figure 63) select the appropriate backup proxy. Automatic is the recommended setting even if Direct NFS has been configured. Under the backup repository ensure it is set to the repository created in the Creating an NFS Backup Repository section and not the default backup repository.

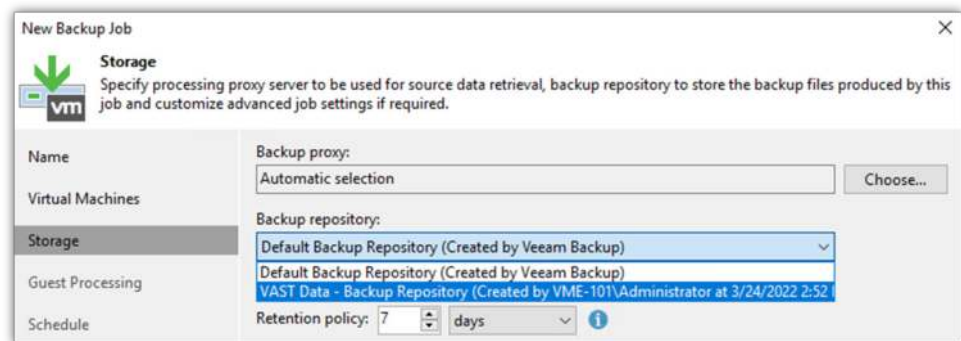


Figure 63 - Select VAST Backup Repository

Make any necessary changes on the Guest Processing window (Figure 64).

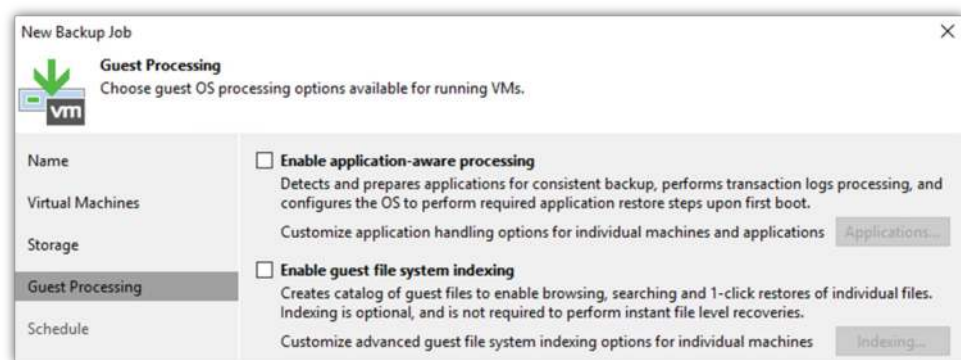
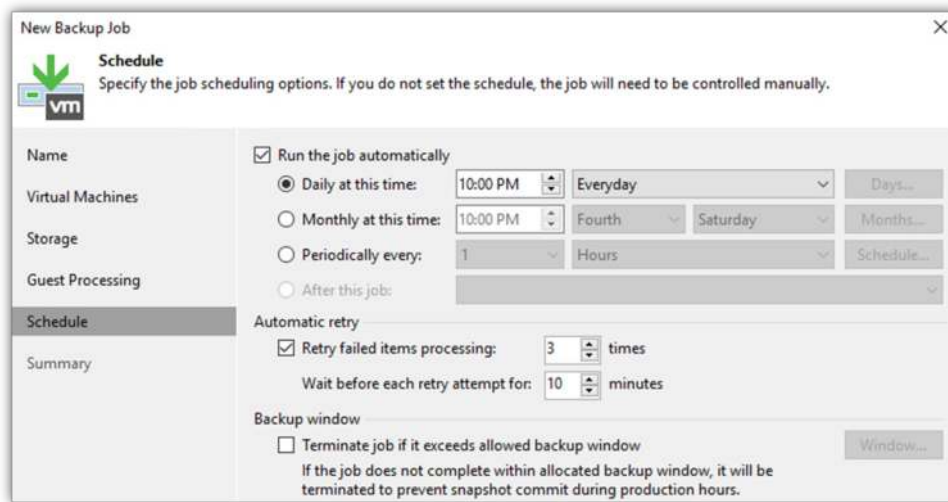


Figure 64 - Guest Processing Settings



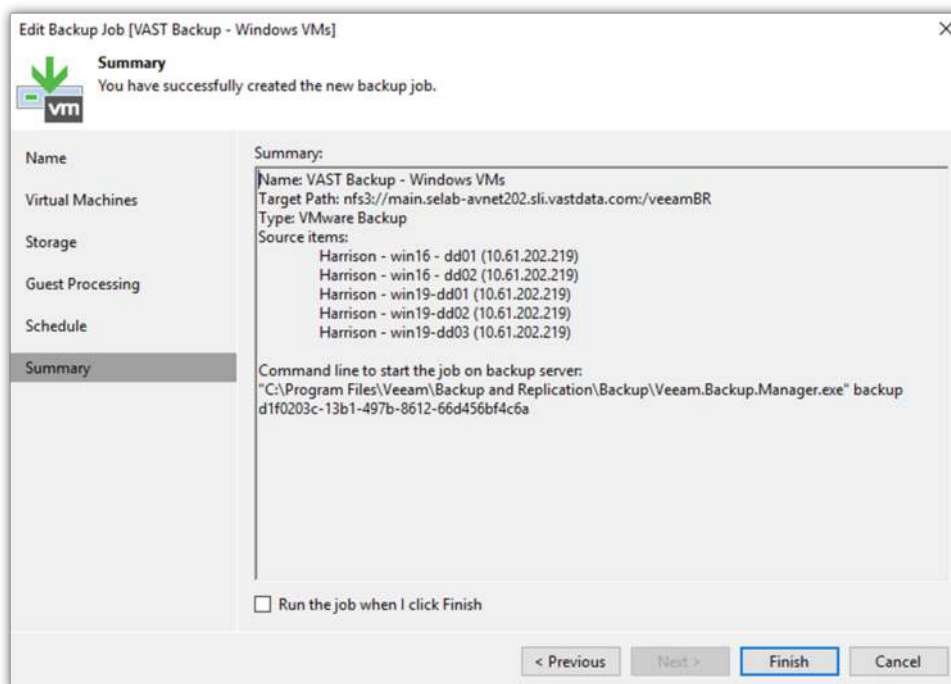
In the Schedule window (Figure 65) create an appropriate backup schedule for the desired recovery point objective (RPO).



The 'New Backup Job' window is shown with the 'Schedule' tab selected. The 'Run the job automatically' checkbox is checked. The schedule is set to 'Daily at this time: 10:00 PM' on 'Everyday'. The 'Automatic retry' section has 'Retry failed items processing' checked, set to '3 times' with a 'Wait before each retry attempt for: 10 minutes'. The 'Backup window' section has 'Terminate job if it exceeds allowed backup window' unchecked. The left sidebar shows 'Name', 'Virtual Machines', 'Storage', 'Guest Processing', 'Schedule', and 'Summary'.

Figure 65 – Scheduling of Backup

Review the configured backup job and click Finish when ready (Figure 66).



The 'Edit Backup Job [VAST Backup - Windows VMs]' window is shown with the 'Summary' tab selected. The 'Summary' section displays the job details: 'Name: VAST Backup - Windows VMs', 'Target Path: nfs3://main.selab-avnet202.sli.vastdata.com/veeamBR', 'Type: VMware Backup', and 'Source items: Harrison - win16 - dd01 (10.61.202.219), Harrison - win16 - dd02 (10.61.202.219), Harrison - win19-dd01 (10.61.202.219), Harrison - win19-dd02 (10.61.202.219), Harrison - win19-dd03 (10.61.202.219)'. The 'Command line to start the job on backup server' is shown as '"C:\Program Files\Veeam\Backup and Replication\Backup\Veeam.Backup.Manager.exe" backup d1f0203c-13b1-497b-8612-66d456bf4c6a'. The 'Run the job when I click Finish' checkbox is checked. The bottom buttons are '< Previous', 'Next >', 'Finish', and 'Cancel'.

Figure 66 – Review Summary

The backup configuration is complete and will begin at the next scheduled time or can be started manually or immediately with the **Run the job when I click Finish** is checked.



VEEAM BACKUP JOBS TO AN S3 REPOSITORY

S3 Object storage repositories can be used a couple ways within Veeam. The methods discussed here will be:

1. Archive location for data from a file share (NFS, CIFS etc)
2. Archive layer of Veeam Scale Out Backup Repository (SOBR)

Backup Job for NFS Share onto VAST S3

In this example an NFS file share was added (steps not shown) to Veeam. The data from this share will be backed up eventually to the VAST S3 repository previously created. A new backup job will be created and just as was done with the backup job for the NFS Repository right click **Jobs** in the left pane and select **Backup** but this time select **File Share** (Figure 58). Give the job a Name as before (Figure 59).

In the next window, the NFS share with the source data that was added is selected (Figure 67). The window allows for browsings so that specific files and folders can be selected for backup.




Figure 67 – Select Files to Backup

After closing the search window confirm that all the appropriate files and folders are selected (Figure 68 – full file path is not displayed).

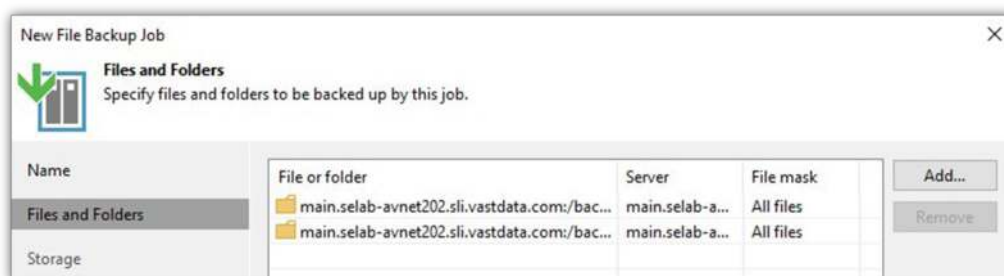


Figure 68 – Review Backup of Files and Folders



On the **Storage** window (Figure 69) there are several key tasks. The first is to select a backup repository and in this case the NFS backup repository is used again. This will be the primary backup location for the NFS file share data selected.

Now, checking the **Keep previous file versions** check box will ungray the **Archive repository** selection. This will allow the VAST S3 repository to be selected that was created in Creating an S3 Backup Repository.

New File Backup Job

Storage
Specify target backup repository and file retention policy for this job.

Name

Files and Folders

Storage

Secondary Target

Schedule

Summary

Backup repository:
VAST Data - Backup Repository (Created by VME-101\Administrator at 3/24/2022 2:52 PM.)

474 TB free of 1.10 PB [Map backup](#)

Keep all file versions for the last: 28 days

Retains recent versions of each file for the specified time period, allowing for restore of entire file shares to a point-in-time state, restore of deleted files, and restore of earlier file versions.

☒ **Keep previous file versions for:** 1 years

Archives older versions of active and permanently deleted files after they are no longer covered by the recent versions retention policy. For scalability reasons, we recommend using object storage.

Archive repository:
VAST S3 - 201 (Created by VME-101\Administrator at 4/5/2022 11:05 AM.)

Files to archive:
All [Choose...](#)

Figure 69 – Enabling the S3 Archive Repository

The **Secondary Target** window allows the user to select a secondary backup location in case multiple copies are desired. Perhaps the secondary target is on a separate VAST cluster or perhaps it's at an offsite location. Here the option is left empty.

New File Backup Job

Secondary Target
We can create additional copies of the short-term file store for redundancy, using the same or different retention policy. The data copy process will start automatically after each primary job run.

Name

Files and Folders

Storage

Secondary Target

Schedule

Summary

Secondary repositories:

Name	Capacity	Retention

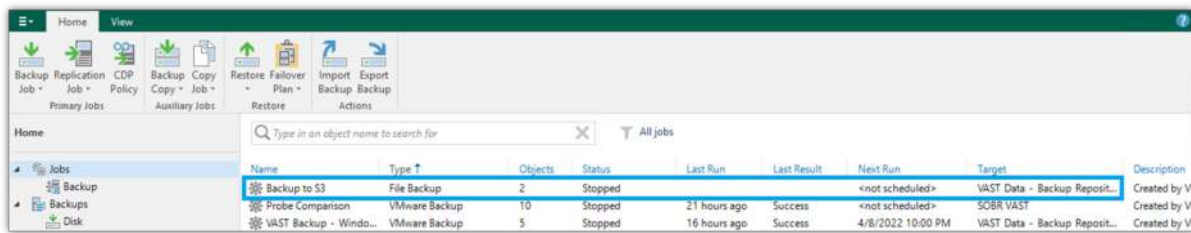
[Add...](#)
[Edit...](#)
[Remove](#)

Data duplication to each secondary repository is performed automatically. You can customize retention, encryption and copy window settings by selecting the repository and clicking Edit.

Figure 70 – Secondary Target



The Schedule screen is next (not shown) and is left to the user to define their backup frequency needs. The last window (not shown) is a summary page of all the settings. Once the backup job is completed it will show up as in Figure 71.



Name	Type	Objects	Status	Last Run	Last Result	Next Run	Target	Description
Backup to S3	File Backup	2	Stopped			<not scheduled>	VAST Data - Backup Reposit...	Created by VI
Probe Comparison	VMware Backup	10	Stopped	21 hours ago	Success	<not scheduled>	SOBR VAST	Created by VI
VAST Backup - Windo...	VMware Backup	5	Stopped	16 hours ago	Success	4/8/2022 10:00 PM	VAST Data - Backup Reposit...	Created by VI

Figure 71 – S3 Backup Job for NFS Share

Backup job for SOBR with VAST S3

The steps for creating a backup job using a SOBR is identical to any other backup job except during storage selection the SOBR is selected as the primary target. Within that construct exists the two tiers, the performance and archive S3 tier. Following the steps from Veeam Backup Job to a VAST NFS Repository up to the point of picking the backup repository (Figure 63).

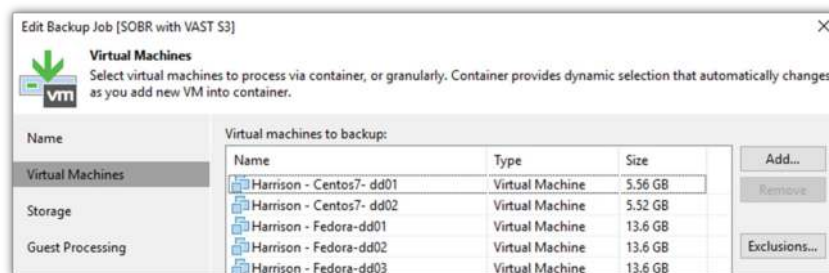


Figure 72 – Selecting VMs for SOBR Backup Job

This example will backup a handful of linux VMs (Figure 72) and more importantly will use the SOBR that was created previously (Figure 73). The rest of the settings are left to the user to determine.

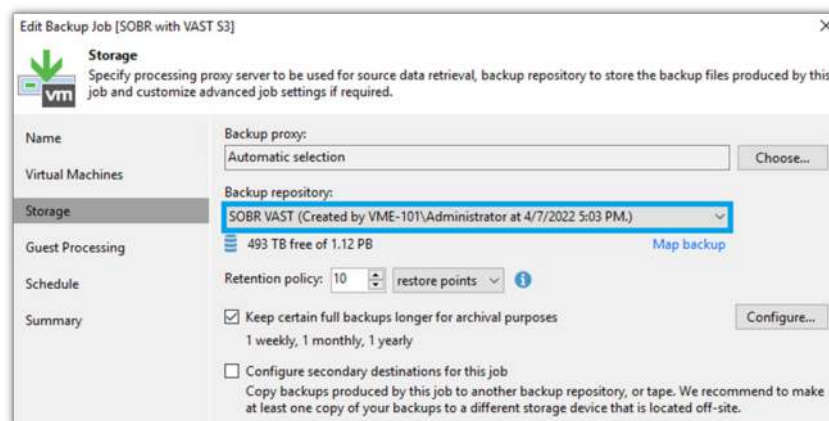


Figure 73 – SOBR Selected as Backup Repository



RESTORING FROM A VEEAM BACKUP JOB

A backup job is only as good as its ability to be restored. This section highlights the restore process on a couple of the backups that were shown in the previous sections.

RESTORING VMS FROM A VAST NFS REPOSITORY

The backup job described in section Veeam Backup Job to a VAST NFS Repository backed up a handful of virtual machines. This section will show how simple it is to restore those VMs.

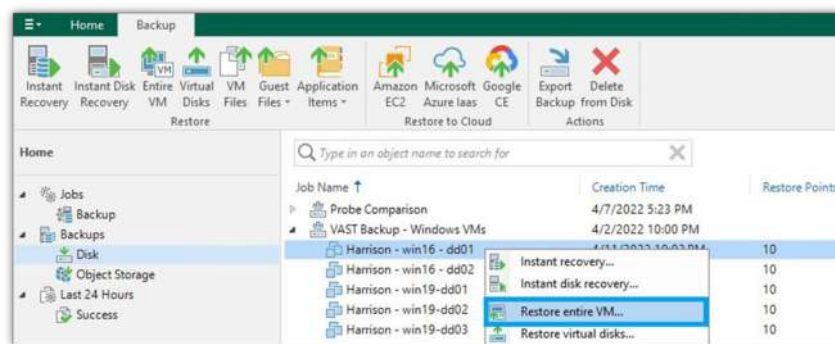


Figure 74 – Select Restore Entire VM

There are multiple ways to bring up the restore wizard. Figure 74 shows the start point as the Home section. Right clicking on the **Disk** icon reveals all of the jobs that have been run in the right window. Opening up the appropriate bckup job reveals all of the VMs that were backed up as part of that job.




Figure 75 – Selecting VMs and Restore Point



Right clicking on one VM and in this example selecting **Restore entire VM** brings up the **Full VM Restore** window shown in Figure 75. The initial screen will only show a single VM from the right click but additional VMs were added using the Add button (Multiple VMs can be selected initially).

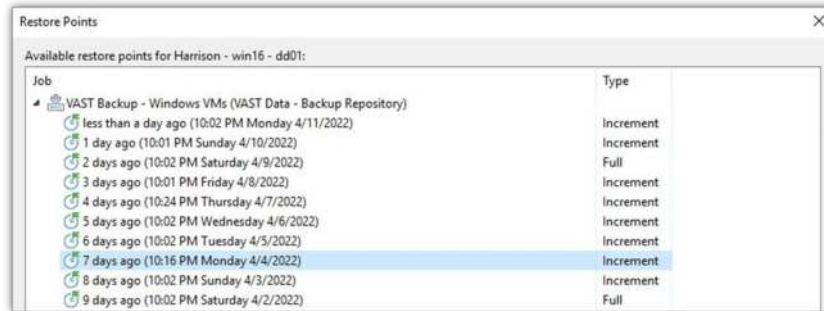


Figure 76 - Pick a Restore Point

The other configuration aspect to note is the Restore Point. For illustration purposes each VM was given a different restore point going from three to seven days back. This was done by clicking the **Point** button and selecting the point in time for recovery of that VM (Figure 76).

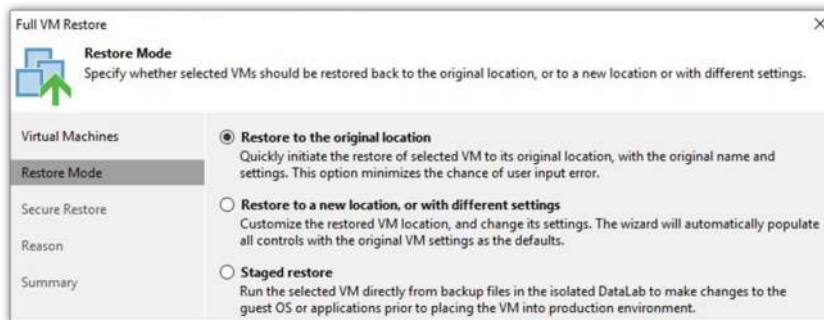


Figure 77 - Select Restore Location

The Restore Mode window (Figure 77) gives the user the ability to place the restored VMs in a multitude of places. Here the VMs are simply being restored to their original location.





The Secure Restore screen gives the user the option to scan the data during recovery for malware. This could be extremely useful if recovering from ransomware to increase confidence that a particular restore point was taken before the ransomware attack. After a few additional minor clicks and reviewing the summary the task will begin.

From the Home tab the current running restore is shown (Figure 78).

The screenshot shows the Veeam Backup & Replication console. The 'Home' tab is selected, and the 'Running (5)' filter is applied. The table below lists the current restore jobs.

Job Name	Session Type	Status	Start Time
Harrison - win19-dd03	Full VM Restore	0% completed	4/12/2022 2:52 PM
Harrison - win19-dd02	Full VM Restore	0% completed	4/12/2022 2:52 PM
Harrison - win19-dd01	Full VM Restore	0% completed	4/12/2022 2:52 PM
Harrison - win16 - dd02	Full VM Restore	15% completed	4/12/2022 2:52 PM
Harrison - win16 - dd01	Full VM Restore	99% completed	4/12/2022 2:52 PM

Figure 78 - Restores Underway

After the restore completes the restored VMs' status show up under the Success section (Figure 79).

The screenshot shows the Veeam Backup & Replication console. The 'Home' tab is selected, and the 'Success' filter is applied. The table below lists the completed restore jobs.

Job Name	Session Type	Status	Start Time	End Time
Harrison - win19-dd03	Full VM Restore	Success	4/12/2022 2:52 PM	4/12/2022 3:00 PM
Harrison - win19-dd02	Full VM Restore	Success	4/12/2022 2:52 PM	4/12/2022 2:58 PM
Harrison - win19-dd01	Full VM Restore	Success	4/12/2022 2:52 PM	4/12/2022 2:58 PM
Harrison - win16 - dd02	Full VM Restore	Success	4/12/2022 2:52 PM	4/12/2022 2:55 PM
Harrison - win16 - dd01	Full VM Restore	Success	4/12/2022 2:52 PM	4/12/2022 2:55 PM

Figure 79 - Successful Restores

RESTORING FILES TO NFS FILE SHARE

In the section, **Backup Job for NFS Share onto VAST S3**, all the files on an NFS share had been backed up to an NFS repository with an archive copy to an S3 repository. To restore files and folders back to the NFS file share is an easy task.

Starting from the Home tab select the Disk icon in the left pane and then open up the appropriate backup job (Figure 80). There are a several options to choose from for restoring files including restoring the entire file share but here, **Files and Folders** was chosen so that specific files can be restored.

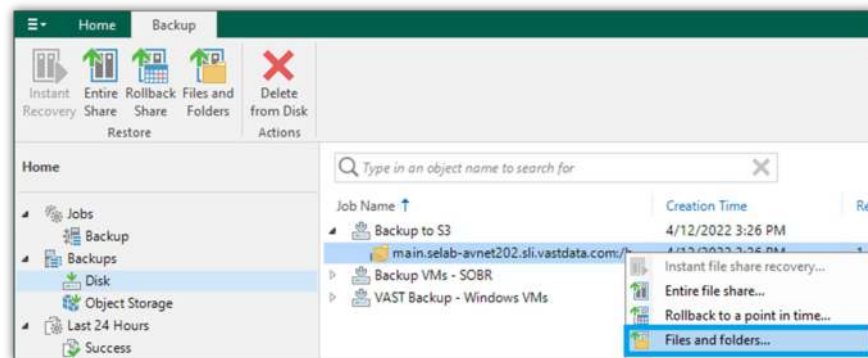


Figure 80 - Select Files and Folders for Restore

A new window will appear (Figure 81) that shows a searchable tree structure. The files to be restored are easily found, highlighted and then right-clicked. From there the files can be restored as a copy or can overwrite the existing version.

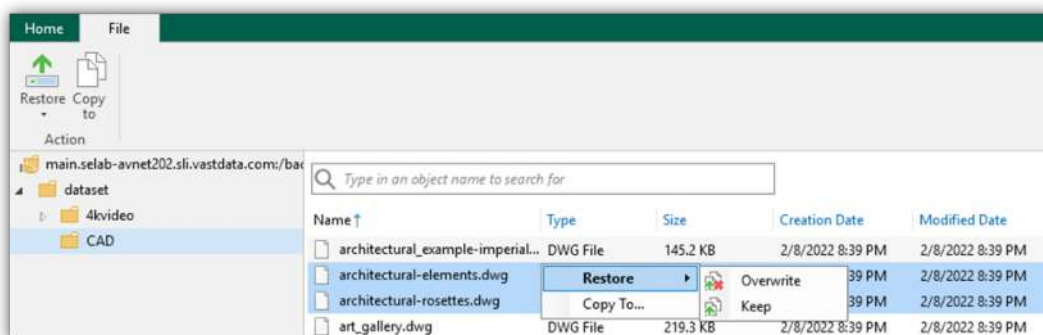


Figure 81 - Browse Folder and Select File(s)

Figure 82 simply shows a successful restore of the two files. To ensure that the restore is successful make sure proper write permissions are given to that share for the Veeam server.

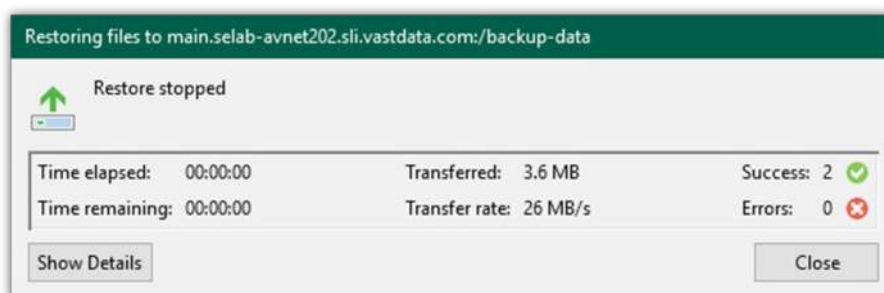


Figure 82 - Files Restored Successfully



RESTORING DATA FROM AN S3 REPOSITORY

Restoring directly from an S3 repository is slightly different than restoring from NFS repository since the role of S3 object store is restricted to archive layer. There needs to be manual intervention to the performance tier (SOBR) or primary backup location or some catastrophic event that the primary backup location is no longer available. The example here puts the performance tier of a SOBR into maintenance mode which prevents any reading and writing to that layer.

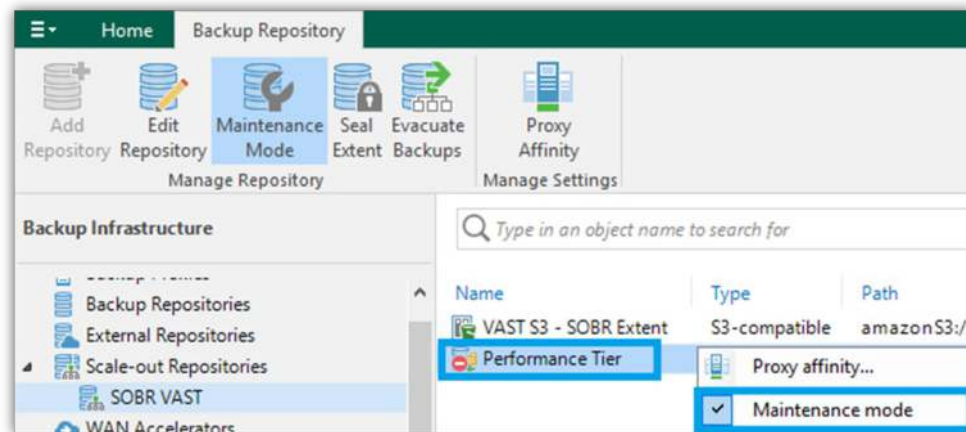


Figure 83 – Put Performance Tier in Maintenance Mode

From the **Backup Infrastructure** tab select the SOBR from the left pane. In the right pane right click on the performance tier extent and select **Maintenance Mode**. Figure 83 shows the tier already in maintenance mode as well as the right-click options. With the primary backup location offline now go to the **Home** tab and select **Object Storage** (Backup job is under Disk too) and then open up the backup job in the right pane (Figure 84). Just like the VM restore from section, **Restoring VMs From a VAST NFS Repository**, select Restore entire VM and proceed through the the steps.

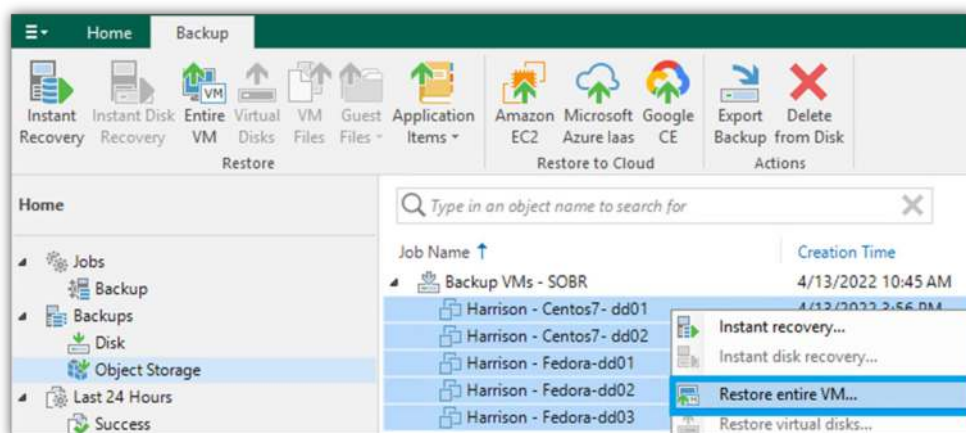


Figure 84 – Selecting VMs to Restore From S3



When the task finishes restoring each VM, the results clearly show where the data was restored from. Figure 85 highlights the fact that all the data came from the capacity or S3 tier.

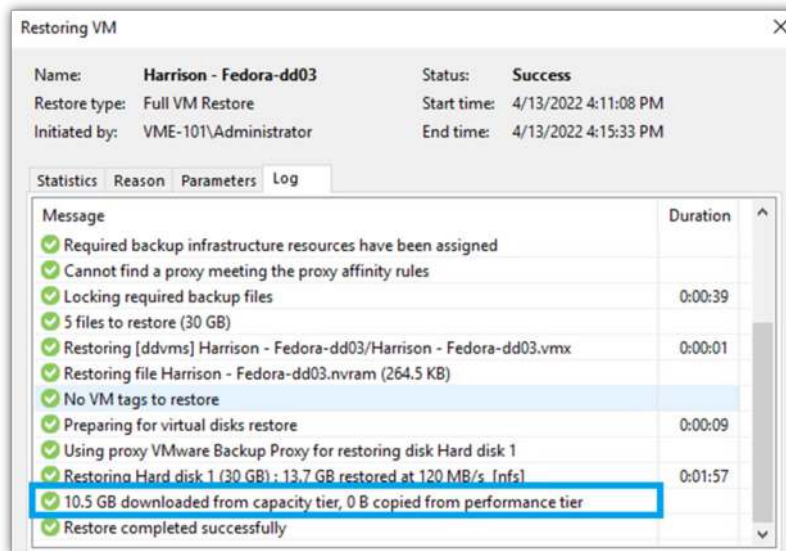


Figure 85 - VM Restored From Capacity (S3) Tier



© 2022 VAST Data, Inc. All rights reserved.

All trademarks belong to their respective owners.