

## Use IDrive® e2 to store Veeam Backups

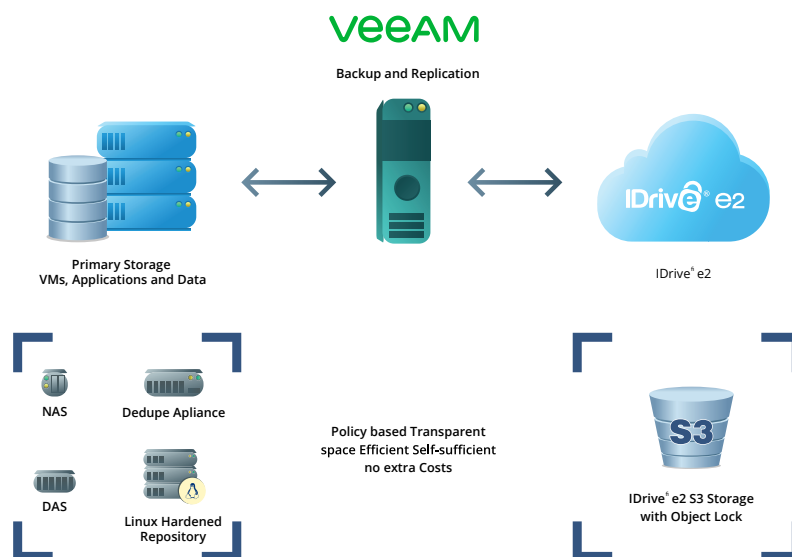
Use IDrive® e2 as Capacity Tier in integration with Veeam Scale-out backup repositories in a smooth transition of backups from Local Backup Repositories to Object Storage Repositories.

### Sections included in this article:

- ✓ Native S3 interface
- ✓ Create a Local Backup Repository (Performance Tier)
- ✓ Configure IDrive® e2 as Object Storage Repository (Capacity Tier)
- ✓ Create Scale-out Backup Repository
- ✓ New Veeam Backup Job

### Native S3 interface

Veeam v11 release has continued its native support for S3 interface for Veeam's Backup & Replication product. This native S3 interface will allow Veeam customers to use S3-compatible storage such as IDrive® e2 as a Capacity Tier.



## Note:

- ✓ The below steps are applicable for Veeam v11 and will not be accurate for Veeam's older versions.
- ✓ To implement a Scale-out backup repository (SOBR), a minimum of Enterprise or Enterprise License is required to leverage IDrive® e2 as a Capacity Tier.

To use IDrive® e2 Object Storage with Veeam Backup and Replication, it is essential to create a Scale-out backup repository that requires a Local Backup Repository (Performance Tier) and IDrive® e2 Object Storage Repository (Capacity Tier).

Local Storage Backup Repository can be created on a vast variety of Storage Systems (DAS, NAS, Dedupe Appliance, Local Storage, etc.)

In this guide, we will create a Local Performance Tier from Veeam Backup Server Local Drives (Dedicated Disk partition from the server).

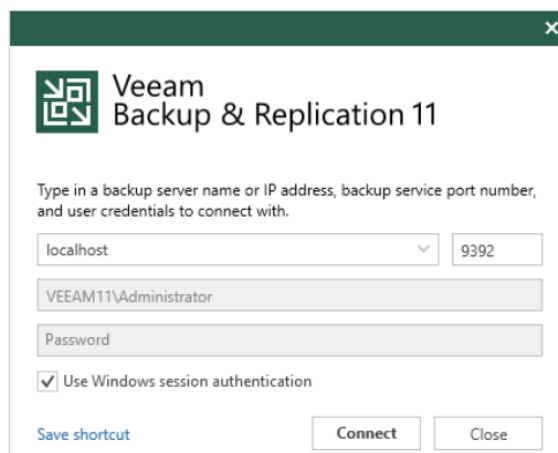
## Note:

- ✓ The steps below show how to create a Local Backup Repository (Performance Tier). Any specific environment may differ from this example.

## Create a Local Backup Repository (Performance Tier)

Backups initially will be stored to the Local Backup Repository and then copied/moved to IDrive® e2 Object Storage, also known as Capacity Tier.

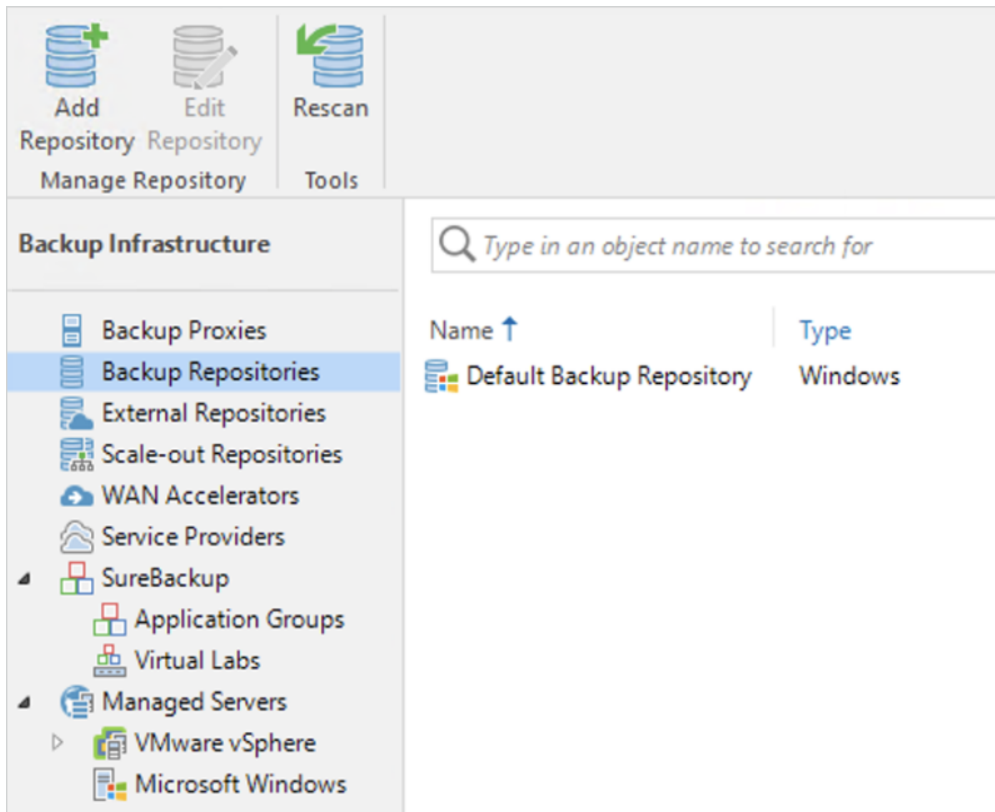
1. Login to **Veeam console**.



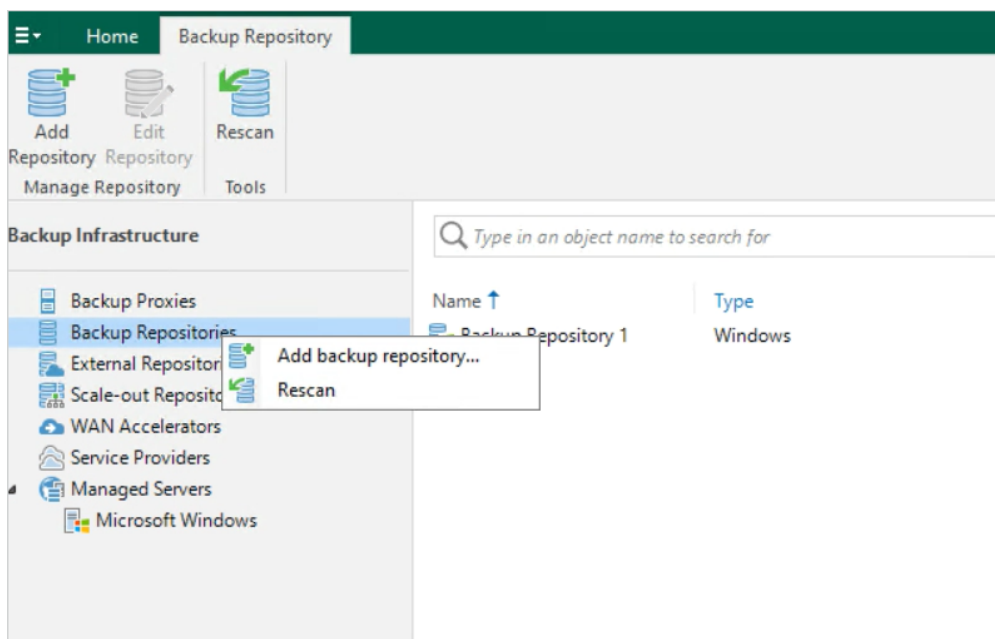
The screenshot shows the Veeam Backup & Replication 11 console login window. The window title is "Veeam Backup & Replication 11". Below the title, there is a logo and the text "Veeam Backup & Replication 11". The main area contains the following fields and options:

- A text prompt: "Type in a backup server name or IP address, backup service port number, and user credentials to connect with."
- A dropdown menu for the backup server name or IP address, currently set to "localhost".
- A text input field for the backup service port number, currently set to "9392".
- A text input field for the username, currently set to "VEEAM11\Administrator".
- A text input field for the password, currently empty.
- A checkbox labeled "Use Windows session authentication" which is checked.
- Buttons for "Save shortcut", "Connect", and "Close".

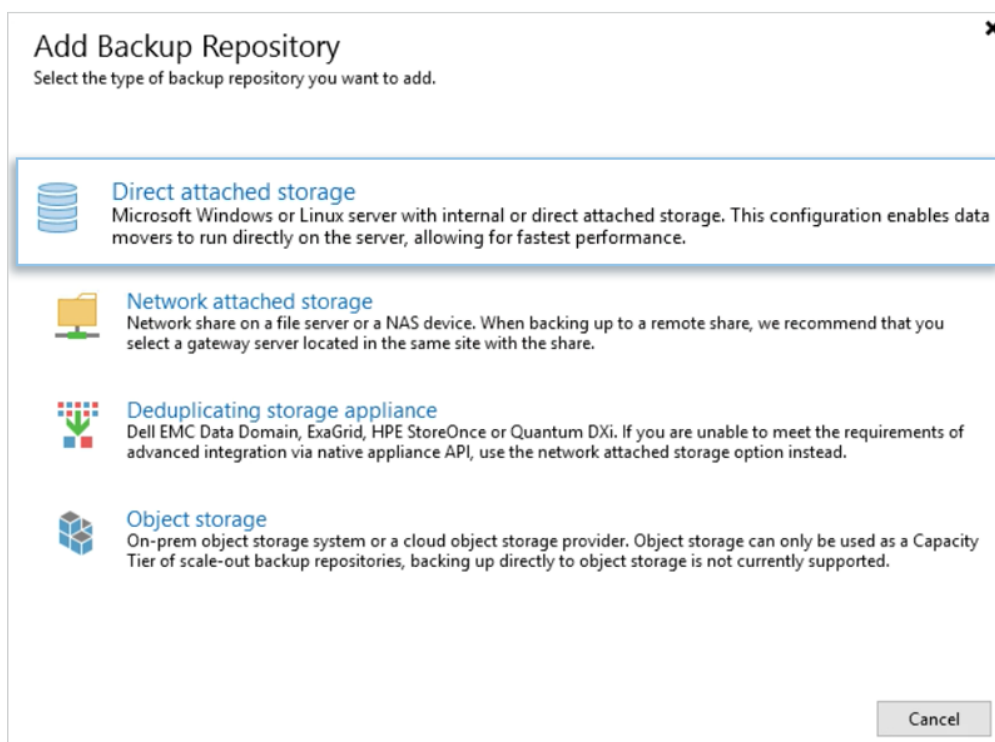
- Navigate to **Backup Infrastructure > Backup Repositories**.



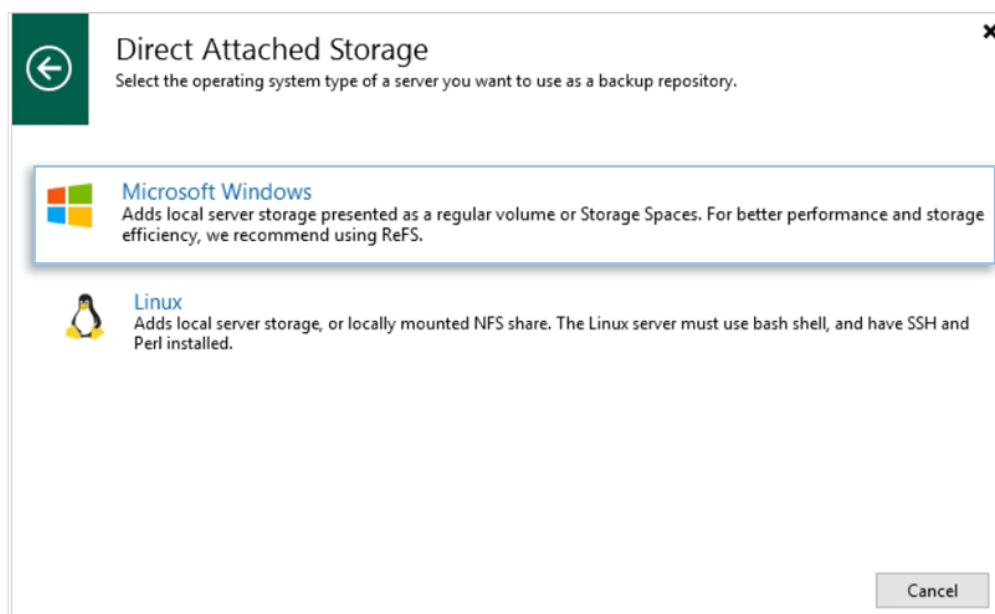
- Right-click on **Backup Repositories** and select **Add Backup Repository**.



4. Select **Direct Attached Storage** ( A NAS or Dedupe Appliance can also be selected.)



5. Select **Microsoft Windows**.



- Enter a name and description for the **Backup Repository**.

**New Backup Repository**

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

**Name:**  
IDrive® e2 Cloud Storage Repo

**Description:**  
My IDrive® e2 Repo

< Previous   **Next >**   Finish   Cancel

- Click **Populate** and chose the drive you want to use as a **Local Backup Repository**.

**New Backup Repository**

**Server**  
Choose repository server. You can select server from the list of managed servers added to the console.

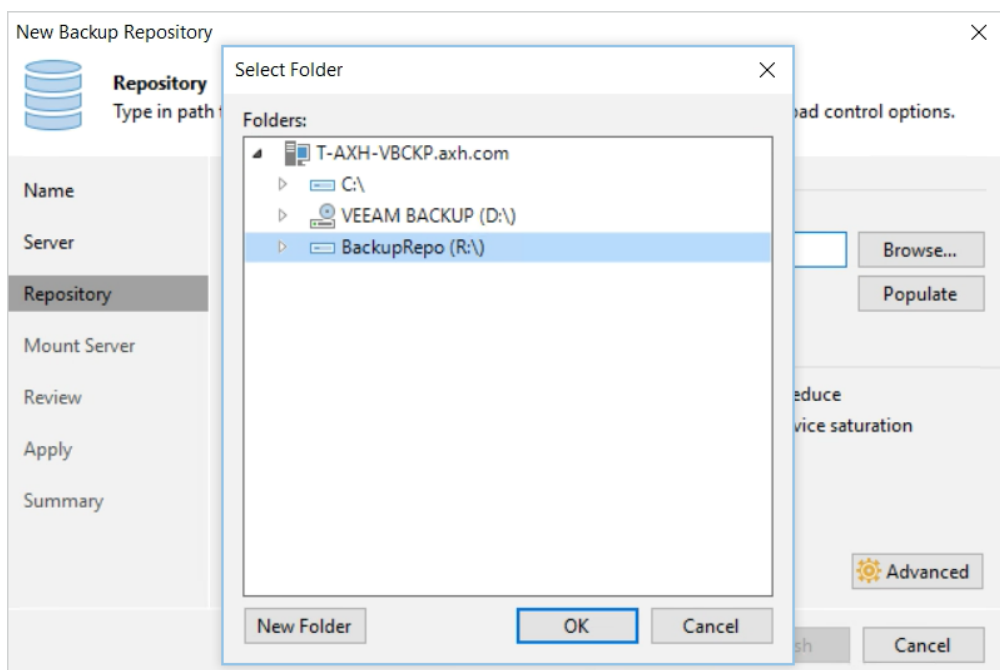
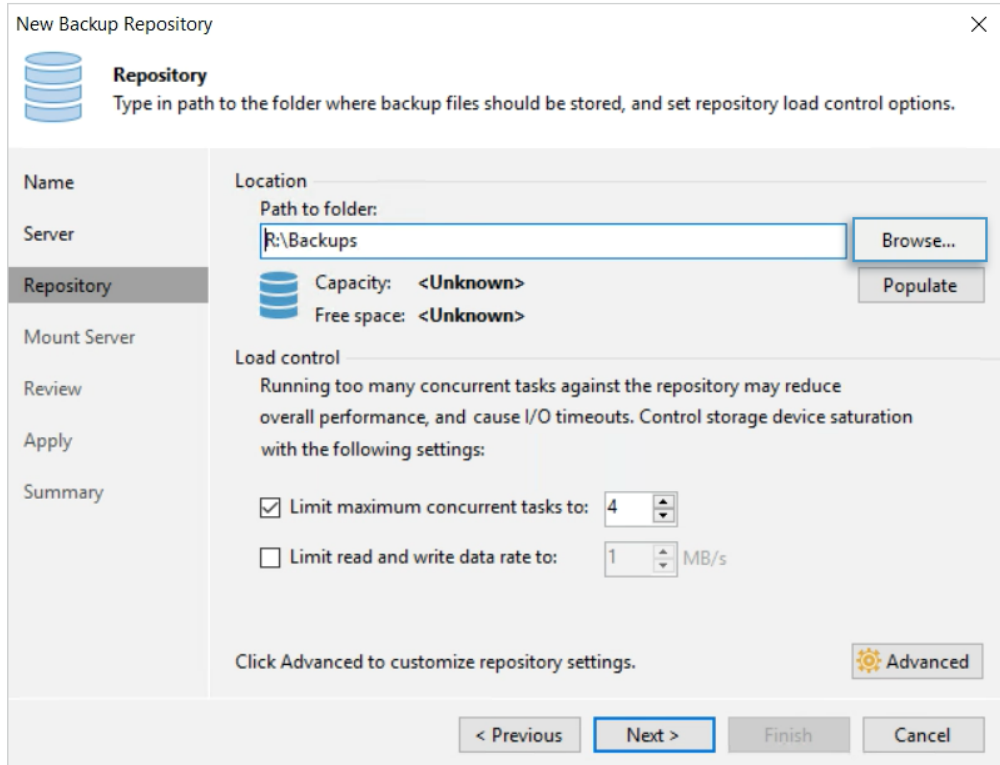
**Repository server:**  
T-AXH-VBCKP.axh.com (Backup server)   Add New...

Path	Capacity	Free
C:\	89.4 GB	47.6 GB
R:\	299.9 GB	296.7 GB

**Populate**

< Previous   **Next >**   Finish   Cancel

- Click **Browse** to select the specific Disk Partition and a Folder from the Veeam Backup Server Disks, that will be used as the target Local Backup Repository.



- While creating a disk partition on Windows, select **ReFS** file system and **64KB** Block size. This is the recommended file system for storing Veeam Backups. If the default **NTFS** file system is selected during disk partition creation, a warning will be shown on the next screen, to inform that the recommended file system is not selected.

**New Backup Repository** [Close]

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

<b>Name</b>	<b>Location</b>
<b>Server</b>	Path to folder: <input type="text" value="R:\"/> <span>Browse...</span>
<b>Repository</b>	Capacity: <Unknown> <span>Populate</span> Free space: <Unknown>
<b>Mount Server</b>	<b>Load control</b>
<b>Review</b>	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:
<b>Apply</b>	<input checked="" type="checkbox"/> Limit maximum concurrent tasks to: <input type="text" value="4"/>
<b>Summary</b>	<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

- Default settings will work fine for Mount Server. Click **Next** to continue with the configuration.

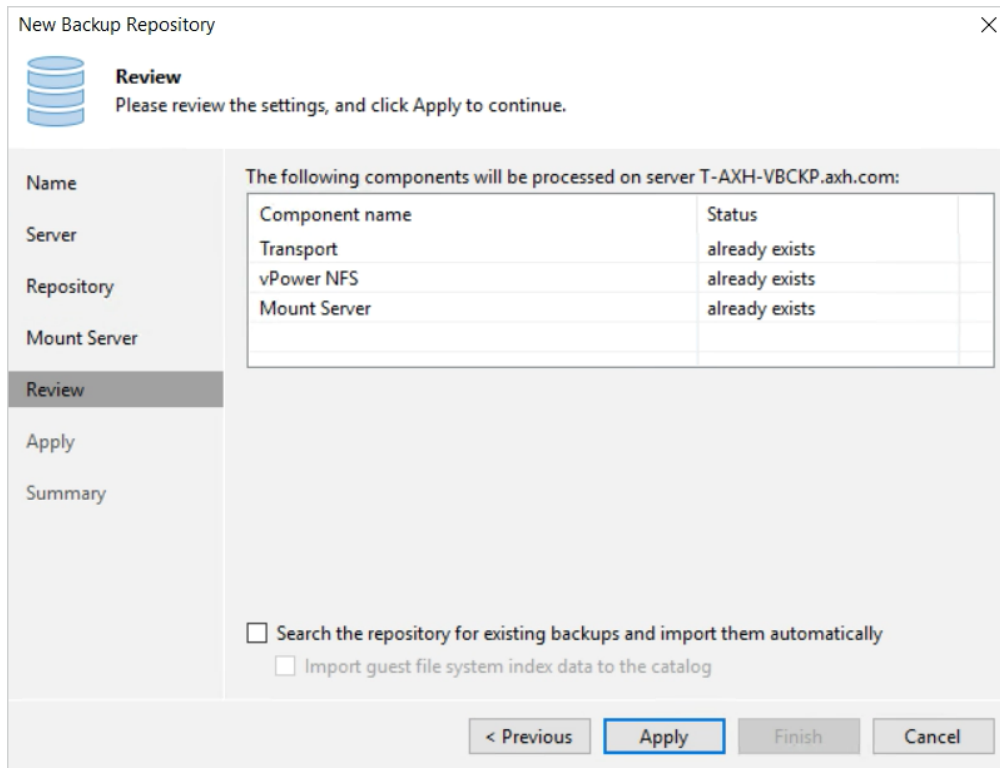
**New Backup Repository** [Close]

**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.

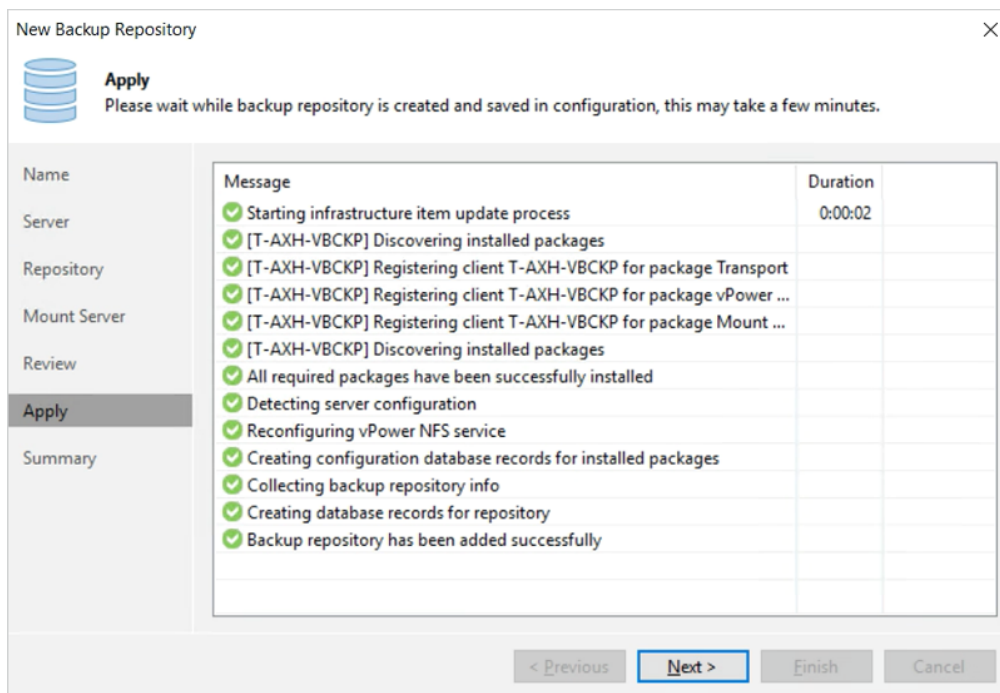
<b>Name</b>	<b>Mount server:</b>
<b>Server</b>	T-AXH-VBCKP.axh.com (Backup server) <span>Add New...</span>
<b>Repository</b>	Instant recovery write cache folder: <input type="text" value="R:\ProgramData\Veeam\Backup\IRCaché\"/> <span>Browse...</span>
<b>Mount Server</b>	Ensure that the selected volume has sufficient free disk space to store changed disk blocks of instantly recovered VMs. We recommend placing write cache on an SSD drive.
<b>Review</b>	<input checked="" type="checkbox"/> Enable vPower NFS service on the mount server (recommended) <span>Ports...</span>
<b>Apply</b>	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.
<b>Summary</b>	

< Previous **Next >** Finish Cancel

11. Review the configuration settings and click **Apply**.

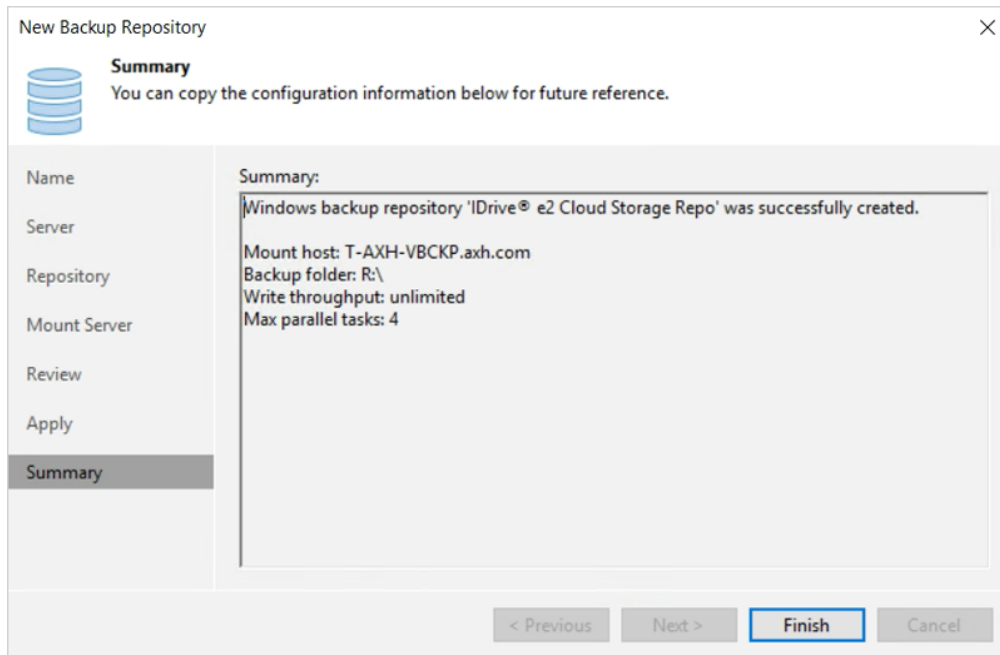


12. Click **Next** to continue.



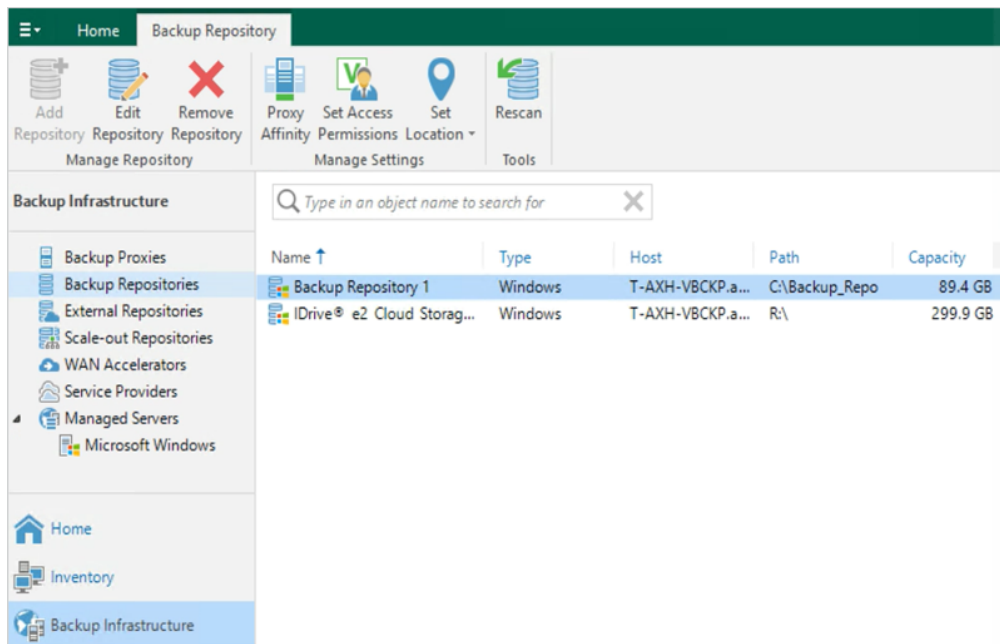


13. Copy the configuration information for future reference and click **Finish**.



14. Choose **No** when asked to change the configuration backup location.

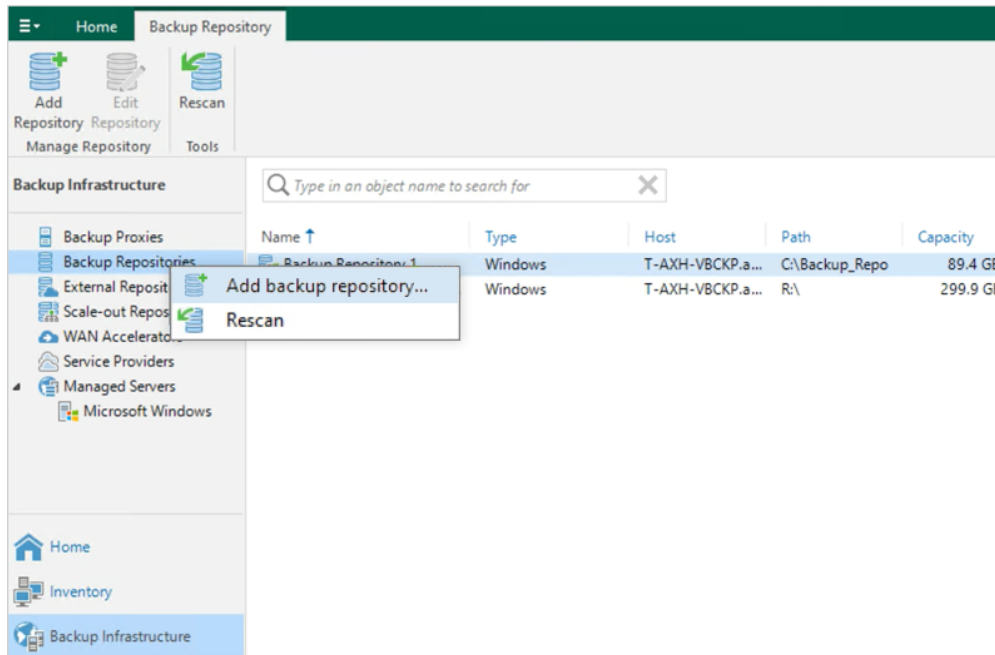
15. A new Backup Repository will be listed under the **Backup Repositories** tab.



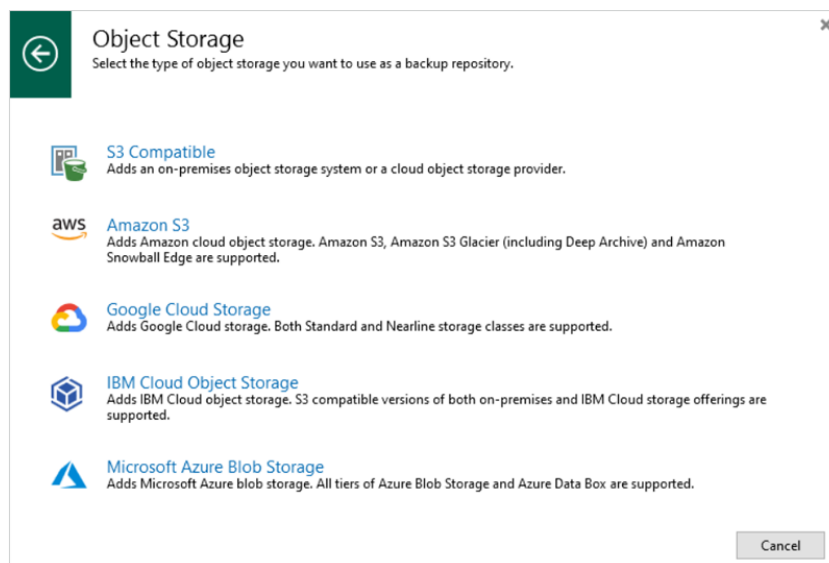
## Configure IDrive® e2 as Object Storage Repository (Capacity Tier)

Follow the steps below to configure IDrive® e2 as an Object Storage Repository:

1. Navigate to **Backup Infrastructure > Backup Repositories** and select **Add Backup Repository**.



2. Select **Object Storage** and choose **S3 compatible**.



3. Name the Object Storage Repository and click **Next**.

New Object Storage Repository

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

**Name:**  
IDrive® e2

**Description:**

Limit concurrent tasks to: 2

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.

< Previous   **Next >**   Finish   Cancel

4. Click **Add** and enter the IDrive® e2 **Access Key** and **Secret Key** and click **OK**.  
[Learn more about Access Keys](#)

Credentials

**Access key:** XXXXXXXXXXXXXXXXXXXXXXXX

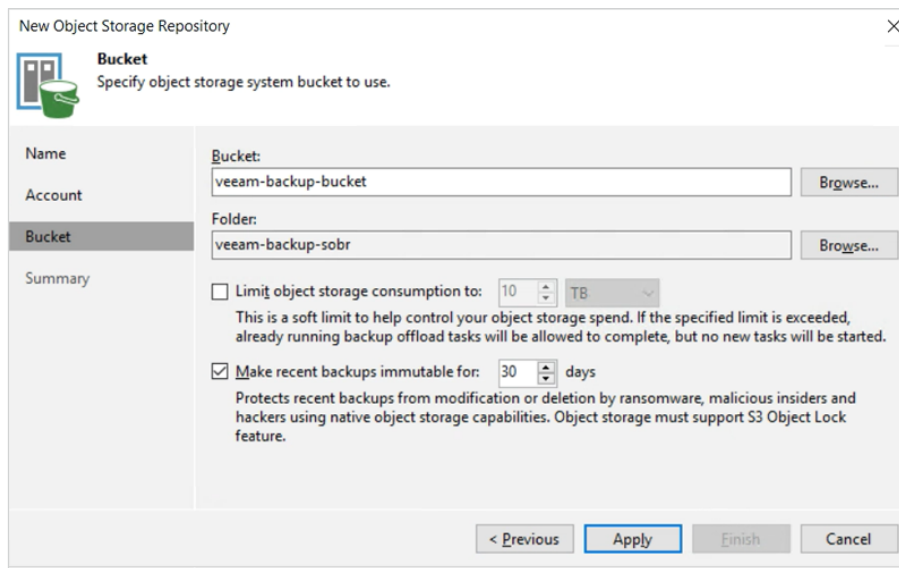
**Secret key:** ●●●●●●●●●●●●●●●●●●●●●●●●●●●●

**Description:**

**OK**   Cancel

- Enter Service Point URL: l4g4.ch11.idrivee2-2.com, provide Region: Chicago, and click Next. [Learn more about regions and endpoints.](#)

- Click **Browse** and select bucket name, folder name, and click **Next**.

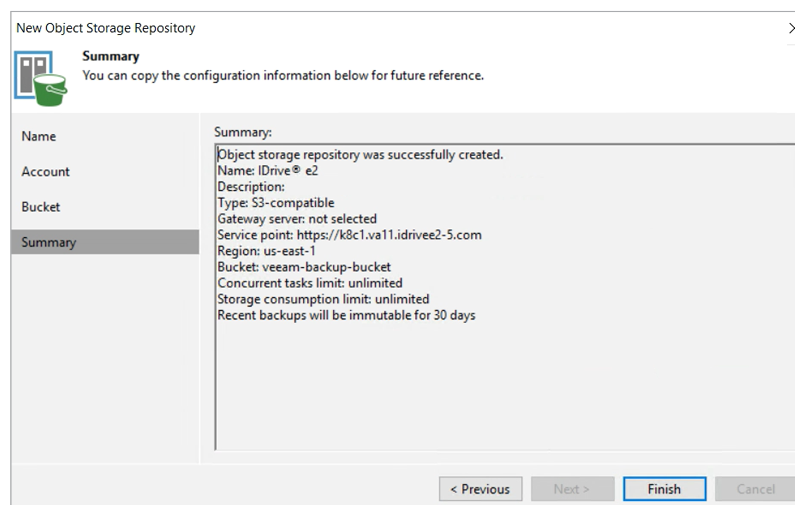


Note: To make recent backups immutable; you need to enable 'Object Lock' during bucket creation. [Learn more about Object Lock.](#)

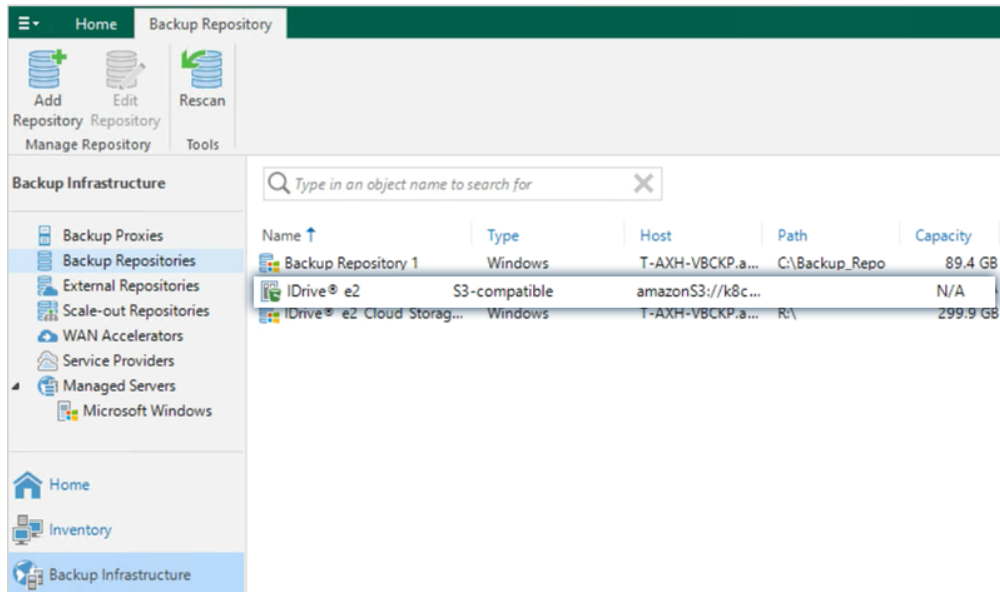
7. Deselect the option for Retention as that option will disallow using Immutability while configuring S3 Object Storage on Veeam Backup Server.

The number of days that are selected here for Immutability, will apply to backups stored on IDrive® e2 Cloud Storage. Before that period of time, backups are subject to retention lock, and thus cannot be deleted/modified or encrypted by no manual or automatic process (Nobody can change/delete/encrypt those files).

8. Copy the configuration information for future reference and click **Finish**.

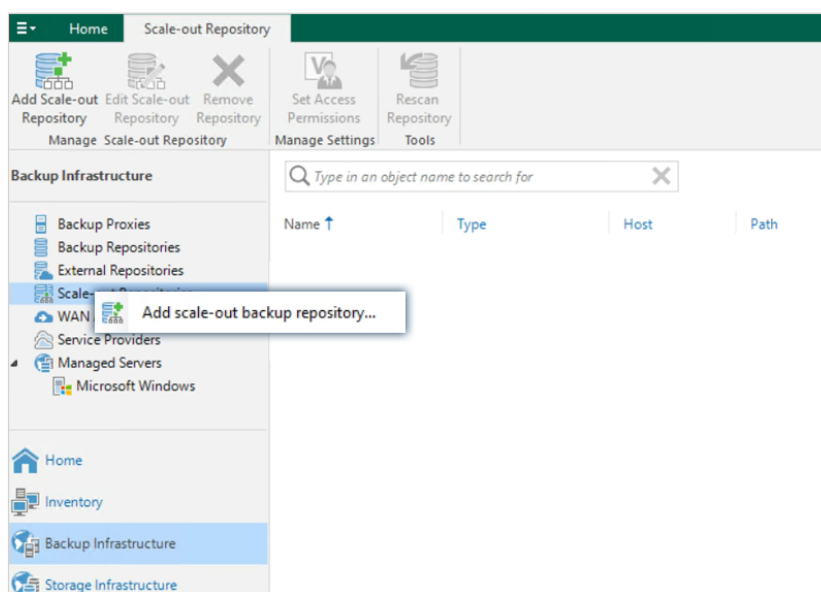


- Navigate to **Backup Infrastructure > Backup Repositories**. The recently added Object Storage Repository will be listed.

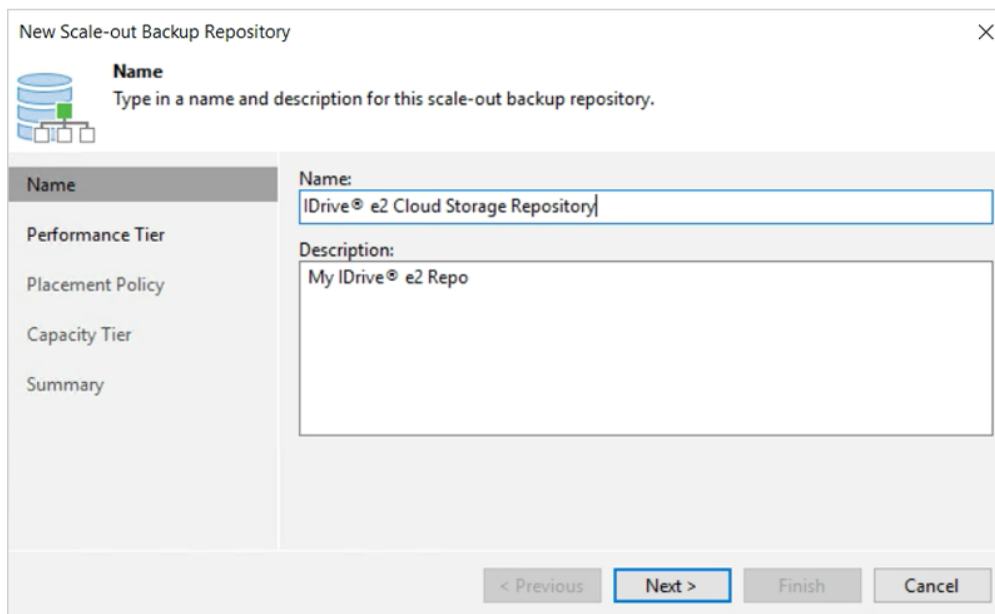


## Create Scale-out Backup Repository

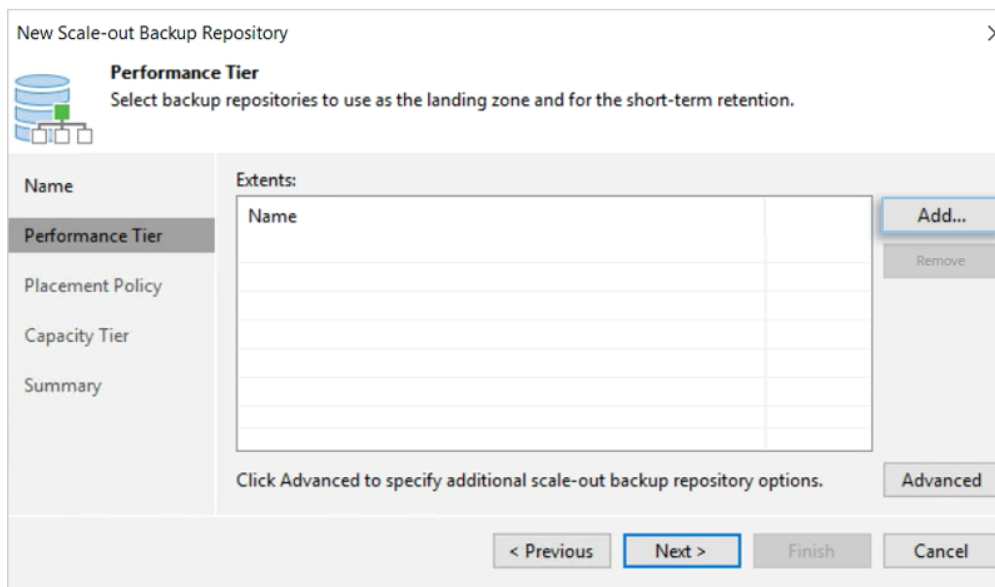
- Navigate to **Backup Infrastructure**, right-click on **Scale-out Repository** and select **Add scale-out backup repository**.

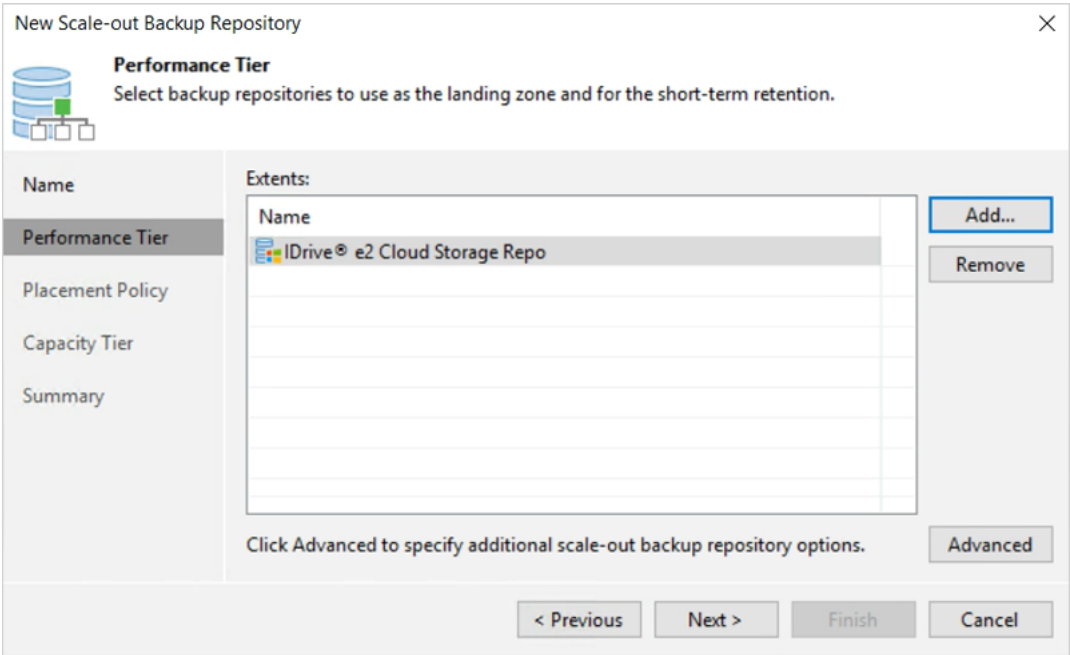
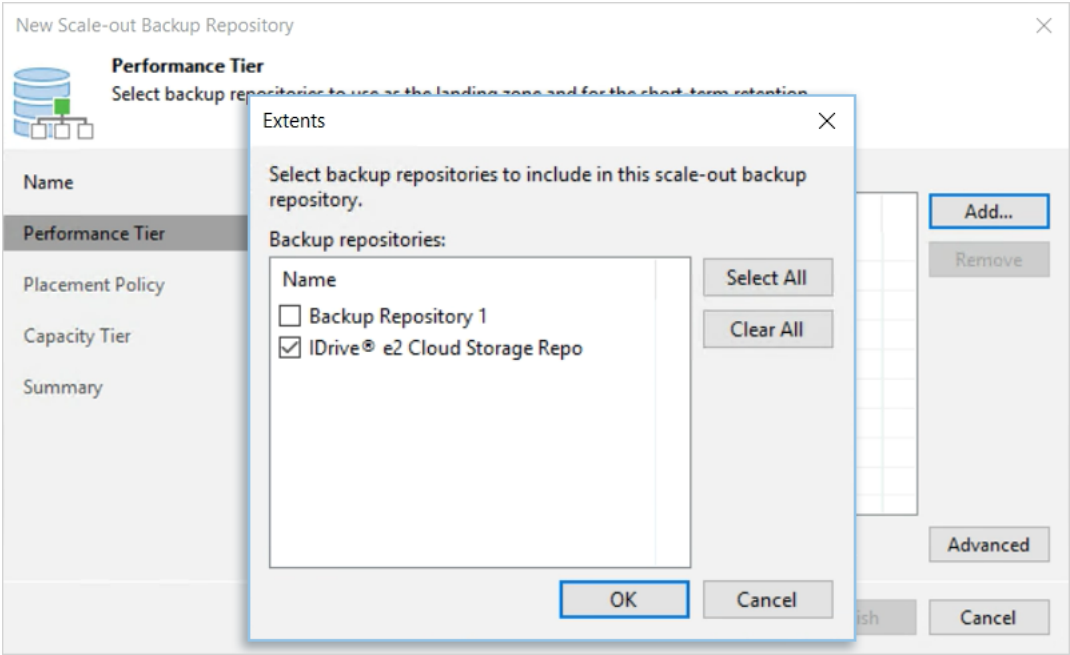


2. Provide a name and description for SOBR and click **Next**.



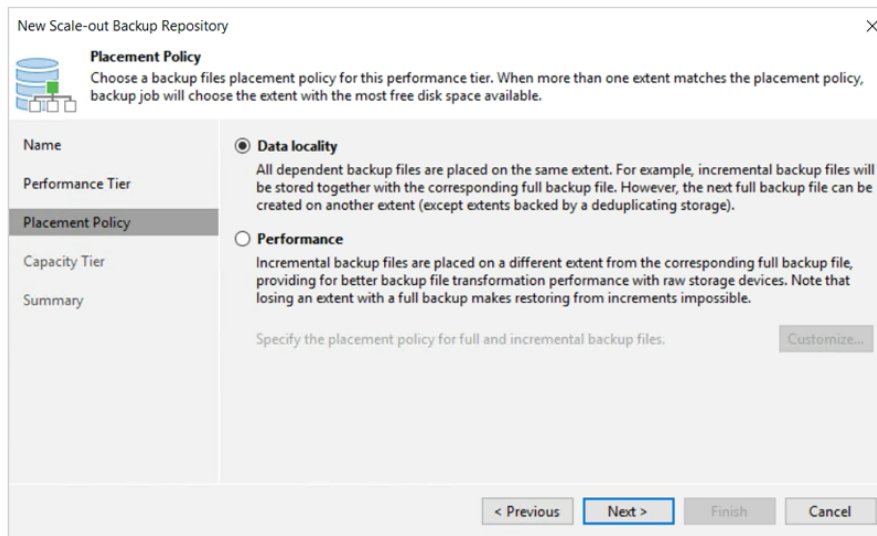
3. On the next screen under Performance Tier, click **Add**. Under the **Extents** pop-up screen, select **Local Backup Repository** created on the previous task, as follows:



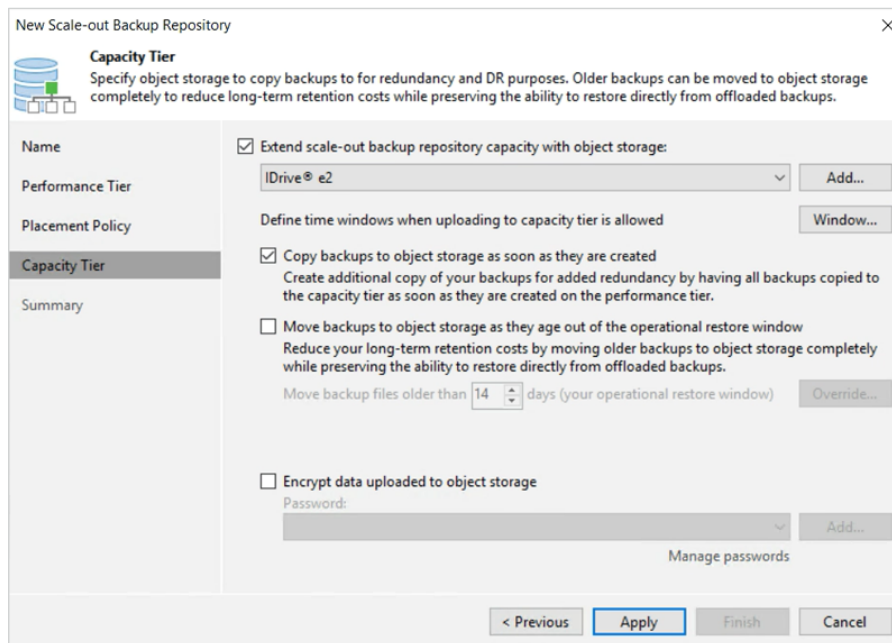




4. Default values for **Placement Policy** should be selected.

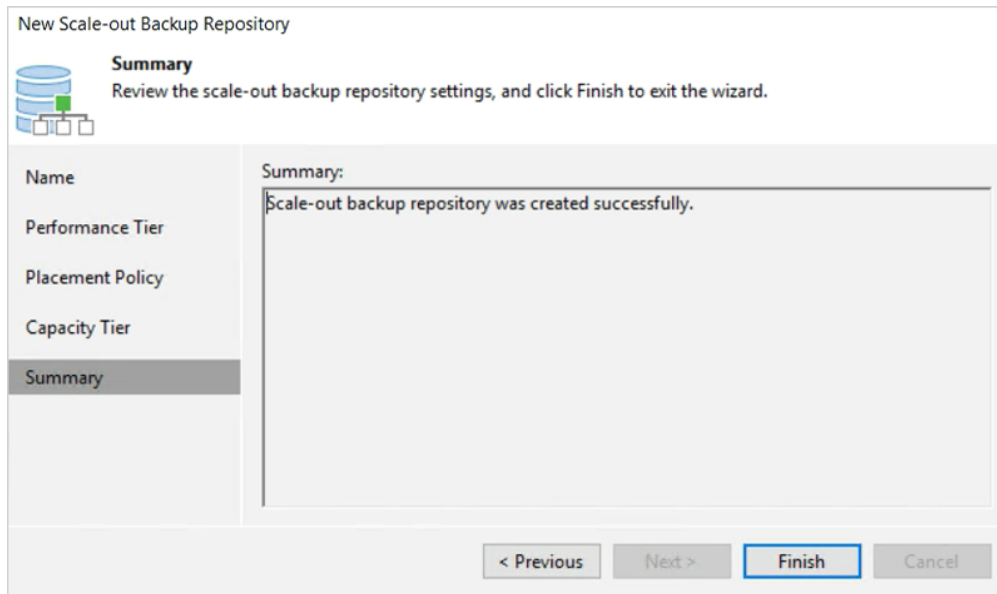


5. Under the Capacity Tier tab, the following options need to be checked/selected.

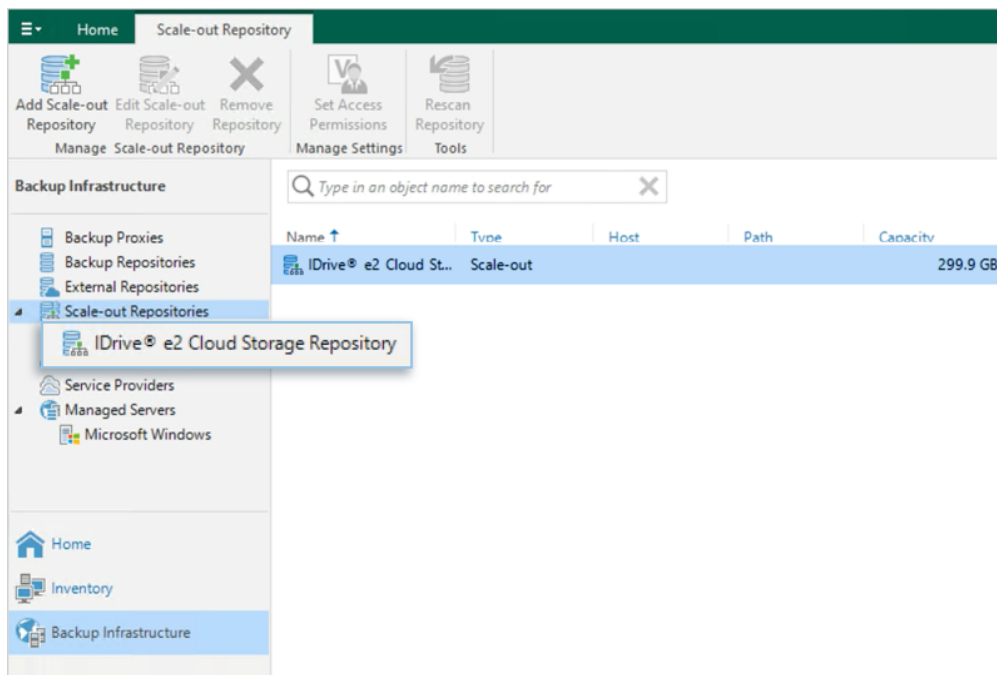


6. Check **Extend scale-out backup repository capacity with object storage** and choose previously created IDrive® e2 Object Storage as Capacity Tier Repo.
7. Check **Copy backups to object storage as they are created**. This will copy all the backups that are created on Performance Tier, immediately to IDrive® e2 Object Storage.

8. Keep the other options unchecked for now.
9. Click **Apply** to continue.
10. Review the settings and click **Finish**.

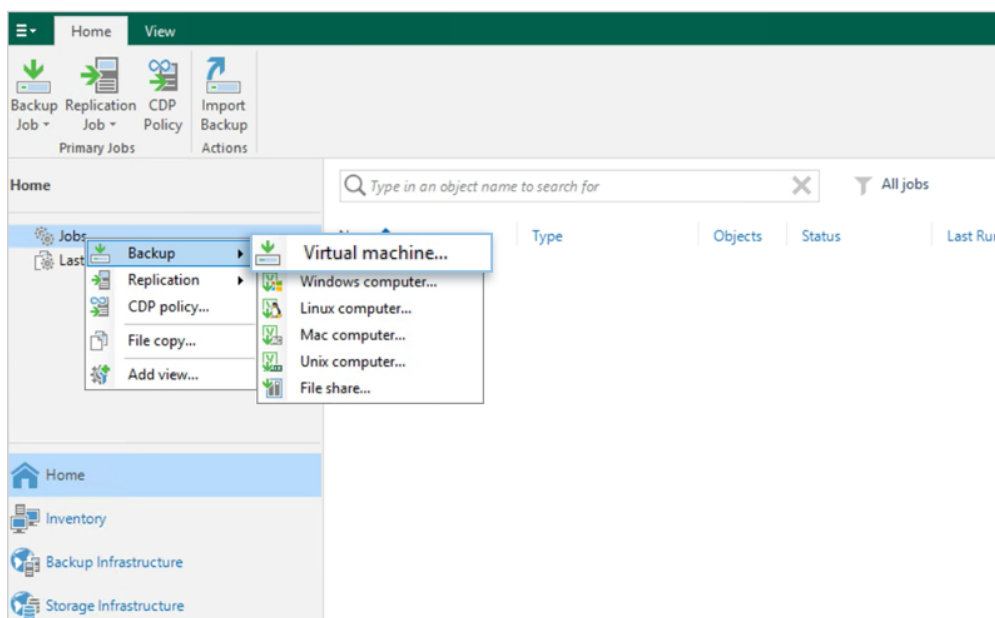


5. A newly created Scale-out backup repository will be listed under the Backup Infrastructure tab.



## New Veeam Backup Job

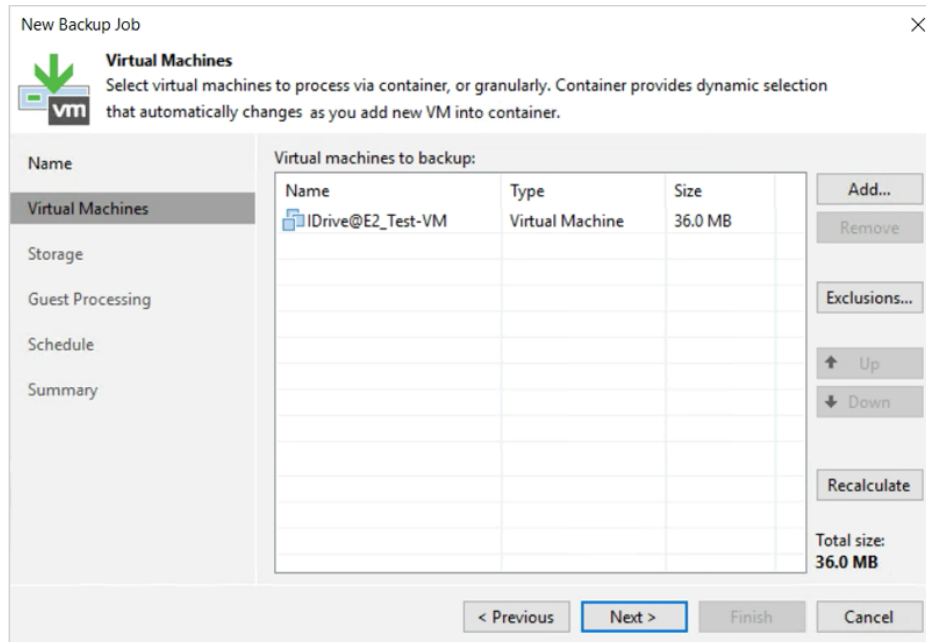
1. Navigate to **Home > Job > Backup > Virtual machine.**



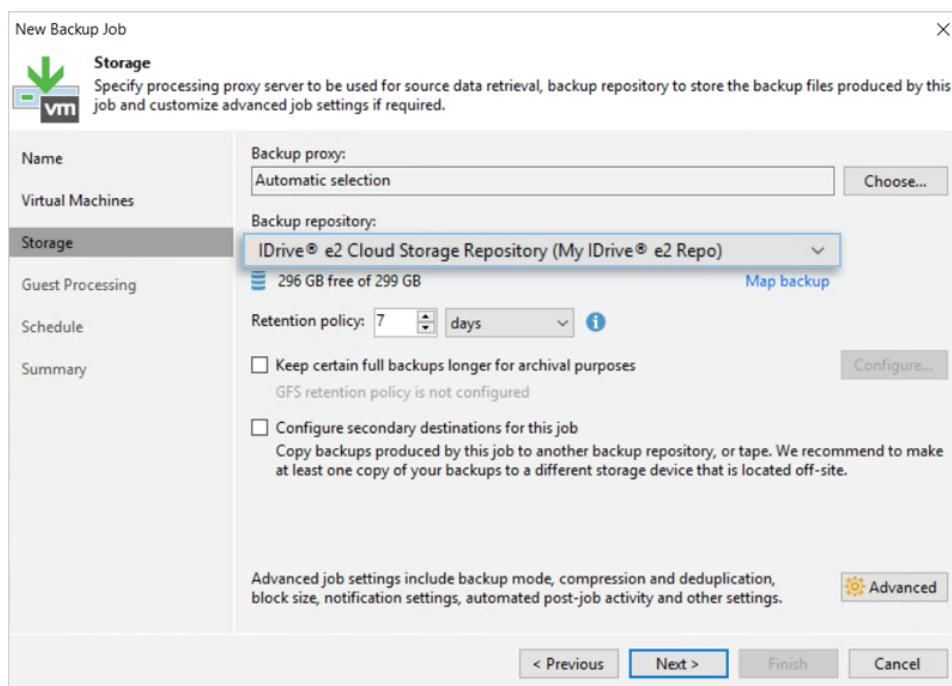
2. Provide a name and description for the backup job and click **Next.**

The 'New Backup Job' dialog box is shown. It has a title bar with a close button. Below the title bar, there is a green arrow icon and a 'vm' icon. The main section is titled 'Name' and contains the text 'Type in a name and description for this backup job.' Below this, there are two input fields: 'Name:' with the value 'IDrive® e2 First Backup Job' and 'Description:' with the value 'IDrive® e2 First Backup Job'. At the bottom, there is a checkbox for 'High priority' which is unchecked. Below the checkbox, there is a note: '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.' At the bottom right, there are four buttons: '< Previous', 'Next >', 'Finish', and 'Cancel'.

3. A test VM created for this guide needs to be selected by clicking **Add** and choose the **virtual machine** to be backed up, and then click **Next**.



4. Select a previously created Scale-out backup repository, as a target for storing backups of the selected VM. Keep the other options unchecked and click **Next**.



5. Guest processing options can be left as they are as well for this example and then click **Next**.

**New Backup Job** [Close]

**Guest Processing**  
Choose guest OS processing options available for running VMs.

**Name**  **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.  
Customize application handling options for individual machines and applications Applications...

**Virtual Machines**  **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.  
Customize advanced guest file system indexing options for individual machines Indexing...

**Storage**

**Guest Processing**

**Schedule**

**Summary**

Guest interaction proxy:  
Automatic selection Choose...

Guest OS credentials:  
 Add...

**Manage accounts**  
Customize guest OS credentials for individual machines and operating systems Credentials...

Verify network connectivity and credentials for each machine included in the job Test Now

< Previous **Next >** Finish Cancel

6. Also, Schedule options can be left default and click **Apply**.

**New Backup Job** [Close]

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

**Name**  **Run the job automatically**

**Virtual Machines**  **Daily at this time:** 10:00 PM  Everyday  Days...

**Storage**  **Monthly at this time:** 10:00 PM  Fourth  Saturday  Months...

**Guest Processing**  **Periodically every:** 1  Hours  Schedule...

**Schedule**  **After this job:**

**Summary**

**Automatic retry**

**Retry failed items processing:** 3  times

Wait before each retry attempt for: 10  minutes

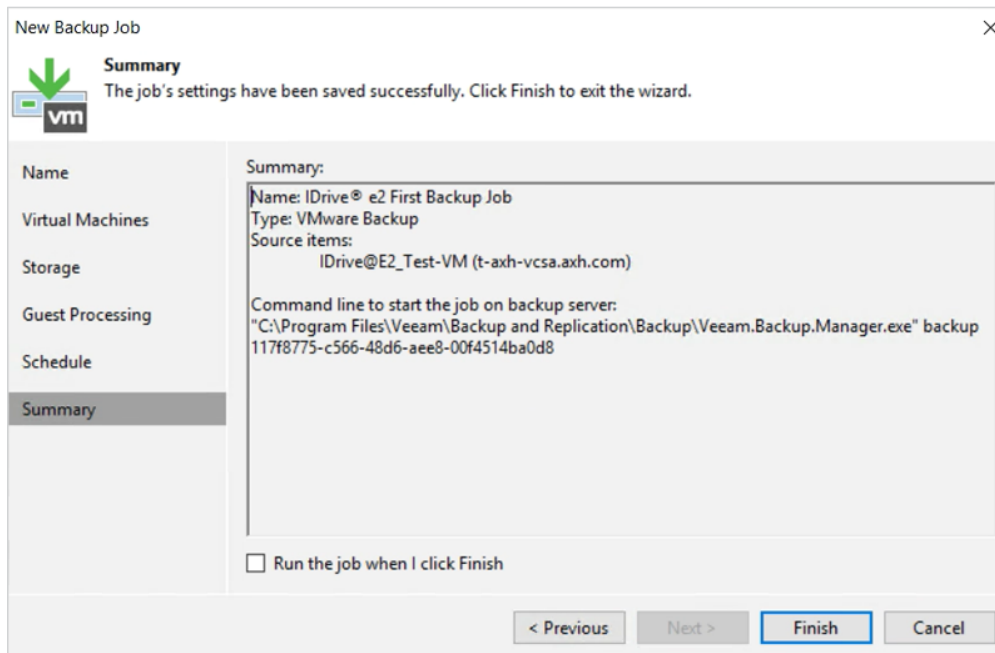
**Backup window**

**Terminate job if it exceeds allowed backup window** Window...

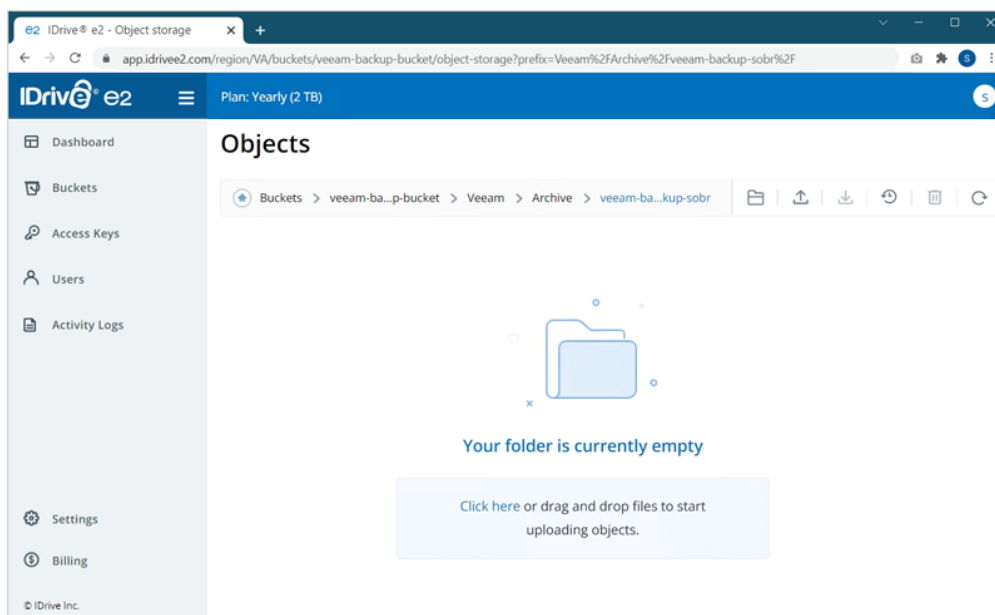
If the job does not complete within allocated backup window, it will be terminated to prevent snapshot commit during production hours.

< Previous **Apply** Finish Cancel

- In the Summary tab, click **Finish**.

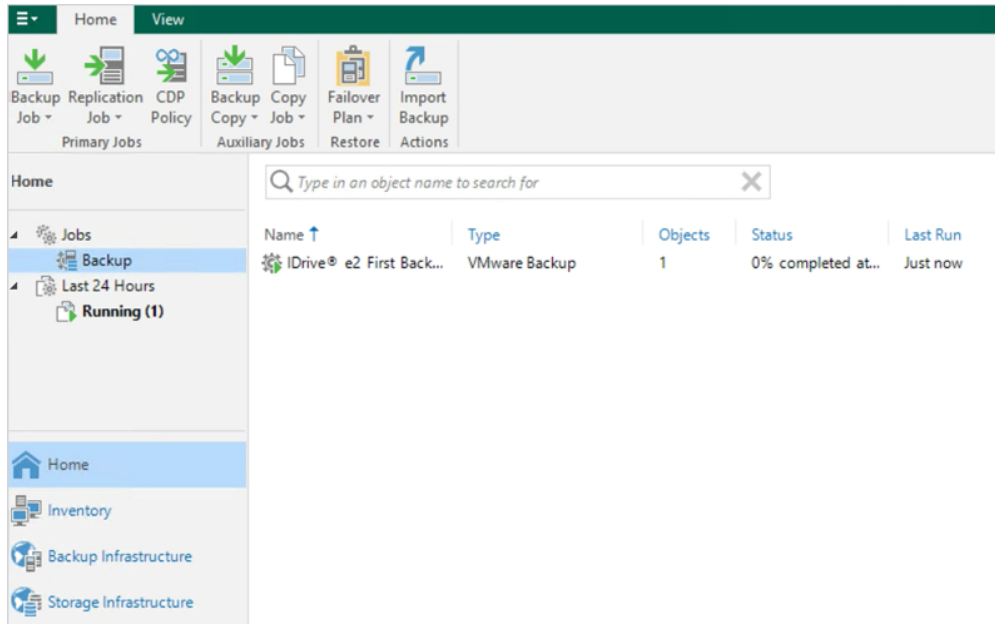


- Before starting the backup job, look at the IDrive® e2 Cloud Storage Bucket for potential files on it. The bucket created for this example purpose looks like as in the picture below:

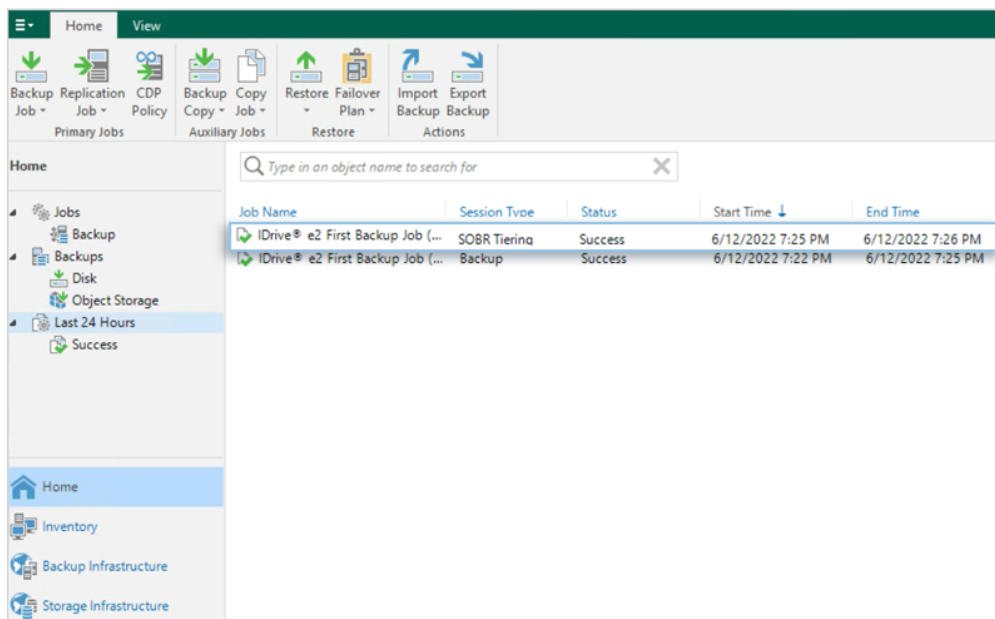


Currently, there are no backups/files stored on it. Once the backup is completed, it will be copied to Cloud Storage and the same folder.

- As Backup Job starts, the VM gets backed up on the Local Backup Repository and copied on the IDrive® e2 Cloud Storage Bucket.



- Once the Backup Job completes, an offload process will start that will move backup files instantly on IDrive® e2 Cloud Storage.



11. Check IDrive® e2 Cloud Storage bucket to ensure the files are copied during the process.

