

Setting Up Quest® QoreStor™ with Veeam®
Backup & Replication™

Technical White Paper

Quest Engineering

February 2023



© 2023 Quest Software Inc.

ALL RIGHTS RESERVED.

THIS WHITE PAPER IS FOR INFORMATIONAL PURPOSES ONLY, AND MAY CONTAIN TYPOGRAPHICAL ERRORS AND TECHNICAL INACCURACIES. THE CONTENT IS PROVIDED AS IS, WITHOUT EXPRESS OR IMPLIED WARRANTIES OF ANY KIND

This guide contains proprietary information protected by copyright. The software described in this guide is furnished under a software license or nondisclosure agreement. This software may be used or copied only in accordance with the terms of the applicable agreement. No part of this guide may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying and recording for any purpose other than the purchaser's personal use without the written permission of Quest Software Inc.

The information in this document is provided in connection with Quest Software products. No license, express or implied, by estoppel or otherwise, to any intellectual property right is granted by this document or in connection with the sale of Quest Software products. EXCEPT AS SET FORTH IN THE TERMS AND CONDITIONS AS SPECIFIED IN THE LICENSE AGREEMENT FOR THIS PRODUCT, QUEST SOFTWARE ASSUMES NO LIABILITY WHATSOEVER AND DISCLAIMS ANY EXPRESS, IMPLIED, OR STATUTORY WARRANTY RELATING TO ITS PRODUCTS INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTY OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE, OR NON-INFRINGEMENT. IN NO EVENT SHALL QUEST SOFTWARE BE LIABLE FOR ANY DIRECT, INDIRECT, CONSEQUENTIAL, PUNITIVE, SPECIAL OR INCIDENTAL DAMAGES (INCLUDING, WITHOUT LIMITATION, DAMAGES FOR LOSS OF PROFITS, BUSINESS INTERRUPTION OR LOSS OF INFORMATION) ARISING OUT OF THE USE OR INABILITY TO USE THIS DOCUMENT, EVEN IF QUEST SOFTWARE HAS BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES. Quest Software makes no representations or warranties with respect to the accuracy or completeness of the contents of this document and reserves the right to make changes to specifications and product descriptions at any time without notice. Quest Software does not make any commitment to update the information contained in this document.

If you have any questions regarding your potential use of this material, contact:

Quest Software Inc.

Attn: LEGAL Dept

4 Polaris Way

Aliso Viejo, CA 92656

Refer to our Web site (<https://www.quest.com>) for regional and international office information.

Patents

Quest Software is proud of our advanced technology. Patents and pending patents may apply to this product. For the most current information about applicable patents for this product, please visit our website at <https://www.quest.com/legal>.

Trademarks

Quest, the Quest logo, and Join the Innovation are trademarks and registered trademarks of Quest Software Inc. For a complete list of Quest marks, visit <https://www.quest.com/legal/trademark-information.aspx>. Veeam® and Veeam Backup & Replication™ are registered trademarks or trademarks of Veeam Software. All other trademarks and registered trademarks are property of their respective owners.

Legend

 **CAUTION:** A CAUTION icon indicates potential damage to hardware or loss of data if instructions are not followed.

 **IMPORTANT, NOTE, TIP, MOBILE, or VIDEO:** An information icon indicates supporting information.

Setting Up Quest® QoreStor™ with Veeam® Backup & Replication™

Updated – February 17, 2023

Contents

Configuring QoreStor as a CIFS/NFS Repository.....	6
Creating a CIFS container for use with Veeam.....	6
Adding the QoreStor CIFS container as a repository in Veeam	8
Creating a NFS container for use with Veeam.....	14
Adding the QoreStor NFS container as a repository in Veeam	16
Configuring Rapid CIFS for Veeam.....	21
Windows prerequisites.....	22
Installing Rapid CIFS on a Veeam Windows Proxy.....	22
Creating a backup job with the QoreStor system as target.....	25
Setting up QoreStor system replication.....	31
Creating a CIFS/NFS replication session	31
Restoring from the replication target.....	33
Using QoreStor as a Veeam Scale-Out Capacity Tier via Object Container(S3)	39
Creating an Object Container(S3) in QoreStor.....	40
Adding the QoreStor Object Container(S3) as a repository in Veeam	42
Adding the Object Container(S3) as a capacity tier to a Scale-Out repository	49
Using Instant Recovery with QoreStor	52
Instant Recovery with ESX.....	52
Enabling Instant Recovery with ESX.....	52
Performing Instant Recovery for ESX	53
Instant Recovery with Hyper-V Server	58
Enabling Instant Recovery with Hyper-V.....	58
Performing Instant Recovery for Hyper-V.....	58
Finalizing Instant Recovery.....	64
Migrating VM to production.....	64
Terminating the Instant VM Recovery Session.....	64
QoreStor and Veeam Fast Clone for Hyper-V 2016 backups or Data Copy.....	65
Requirements of Fast Clone.....	65

Configuring a new Fast Clone Repository	66
Reconfiguring an Existing QoreStor Repository for Fast Clone.....	68
Performance Tier	70
Setting up Performance Tier with QoreStor	71
Optimizing Performance Tier via Sync Always option.....	73
Cloud/Archive Tier.....	74
Cloud Tier	74
Important Considerations for Cloud Tier with Veeam.....	74
Setting up Cloud Tier.....	76
Archive Tier	79
Important Considerations for Cloud Tier with Veeam.....	79
Setting up Archive Tier.....	79
Setting up the QoreStor system cleaner	82
Monitoring deduplication, compression and performance	84

Executive Summary

This paper provides information about how to set up Quest® QoreStor™ as a backup target for Veeam® Backup & Replication™ software.

For additional information, see the QoreStor documentation and other data management application best practices whitepapers for your specific QoreStor version at:

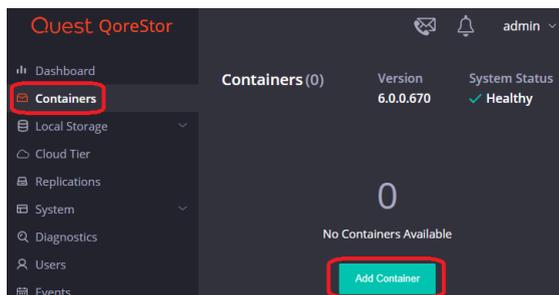
<https://support.quest.com/qorestor/>

i | **NOTE:** The QoreStor and Veeam screenshots used in this document may vary slightly, depending on the QoreStor and Veeam versions you are using.

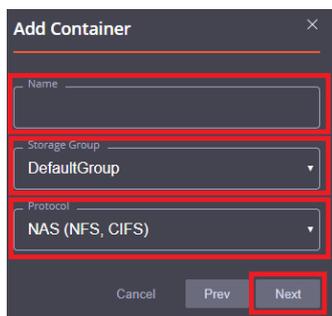
Configuring QoreStor as a CIFS/NFS Repository

Creating a CIFS container for use with Veeam

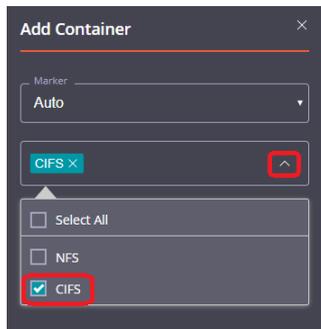
- 1 Select the **Containers** tab, then click **Add container**.



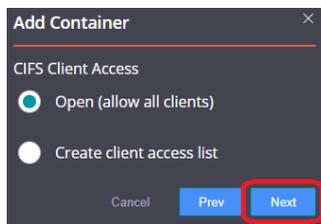
- 2 Enter a container **Name**, select a **Storage Group**, or leave the **DefaultGroup** option selected, and select **NAS (NFS, CIFS)** from the **Protocol** dropdown menu. Click **Next**.

The 'Add Container' dialog box is shown. It has three input fields: 'Name' (empty), 'Storage Group' (set to 'DefaultGroup'), and 'Protocol' (set to 'NAS (NFS, CIFS)'). At the bottom, there are three buttons: 'Cancel', 'Prev', and 'Next', with 'Next' highlighted by a red box.

- 3 Click the dropdown on the **Protocols** field then select the check mark for **CIFS**. Leave **Marker Type** on **Auto**, then click **Next**.

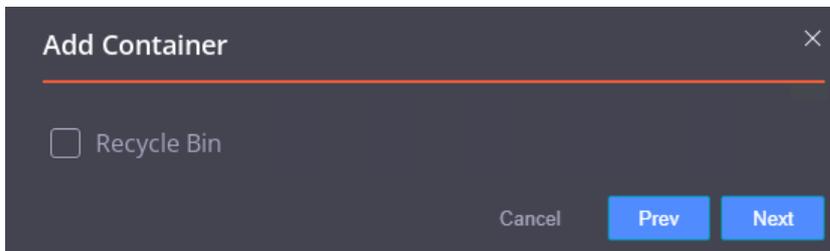


- 4 Fill in the **CIFS Client Access** options if needed then click **Next**.

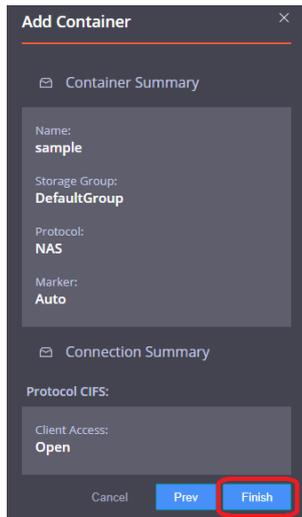


i **NOTE:** For improved security, Quest recommends adding IP addresses for only Veeam servers/proxies.

- 5 On this page, the Recycle Bin feature may be enabled, please check the user guide for more information. Click **next**.



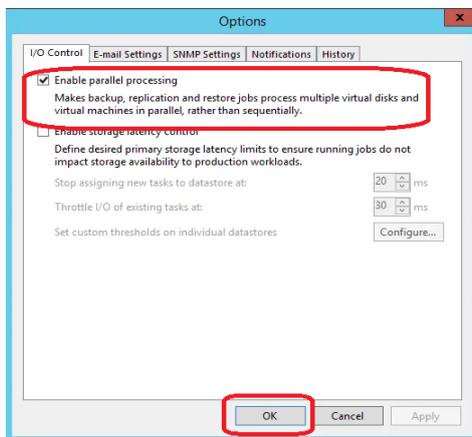
- 6 Confirm the settings and click **Finish**. Confirm that the container is added.



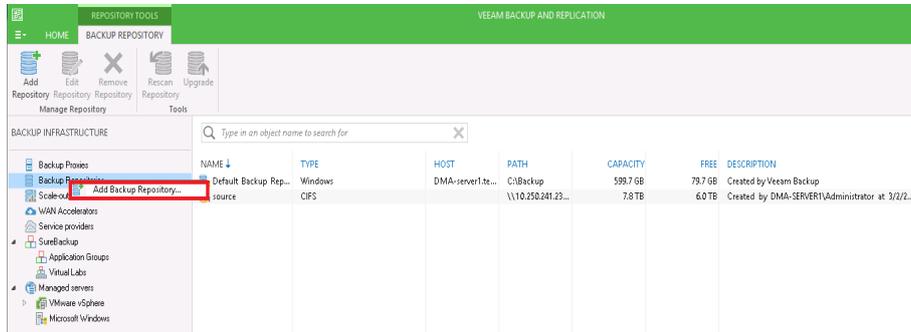
Adding the QoreStor CIFS container as a repository in Veeam

- CAUTION:** To maximize the QoreStor and Veeam deduplication savings and performance, Quest recommends using the exact settings in this guide for all the data being backed up. The backup data will change format completely when backup settings are changed. Hence, to get accurate savings numbers, all the data should be backed up with the same settings.

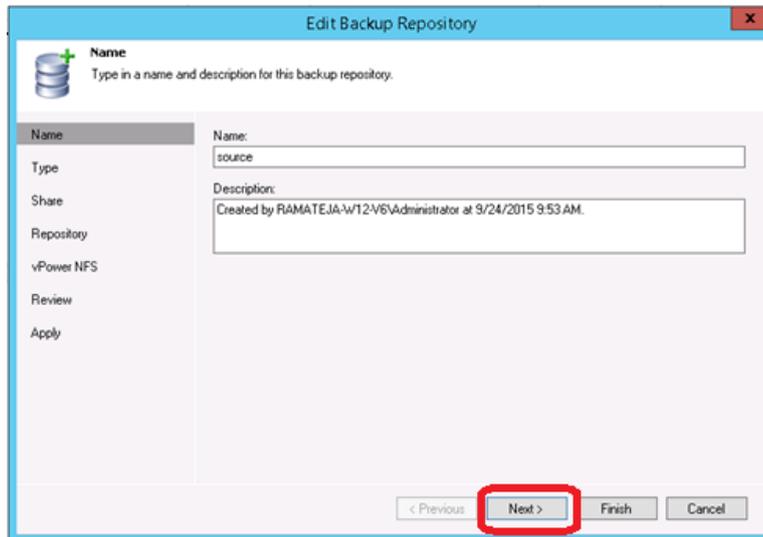
- 1 Open the Veeam Backup & Replication console.
- 2 If using Veeam 9.5 U3 or lower, select the dropdown **Menu** and click **General Options**.
- 3 Check the **Enable parallel processing** option in the I/O Control tab and click **OK**. This option will be missing in Veeam 9.5 U4 and higher as it's automatically enabled by default.



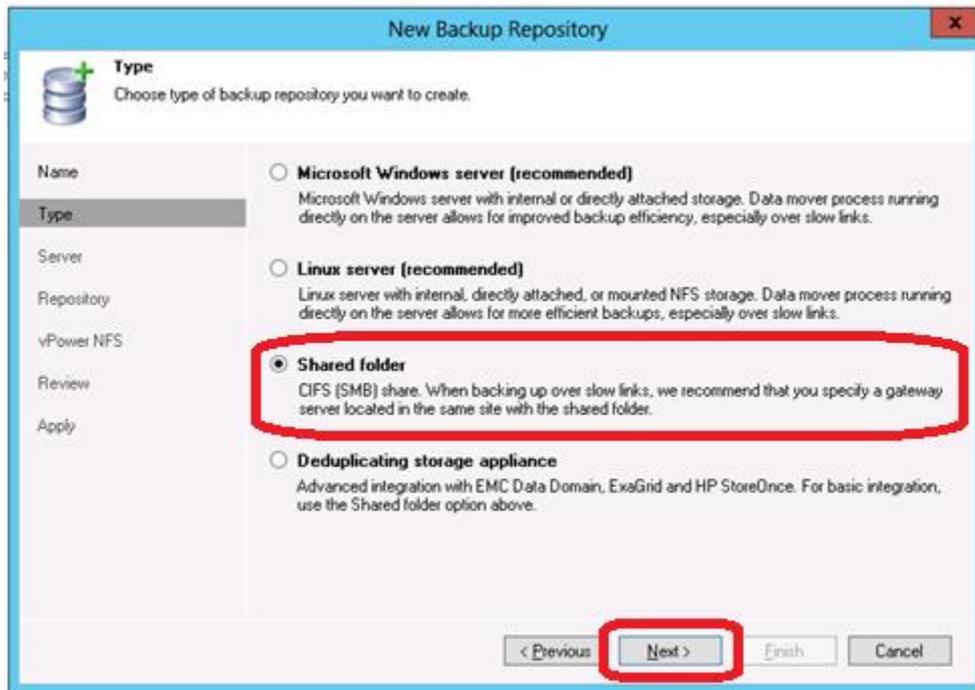
- 4 In the **Backup Infrastructure** section, right-click **Backup Repositories**, and select **Add Backup Repository**.



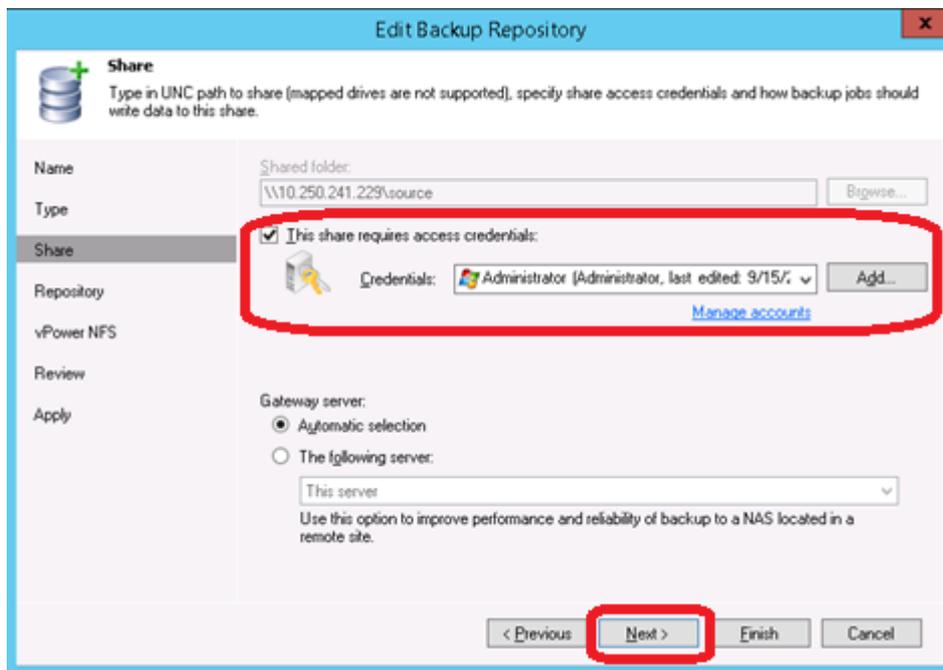
- 5 Enter a name for the QoreStor container repository and click **Next**.



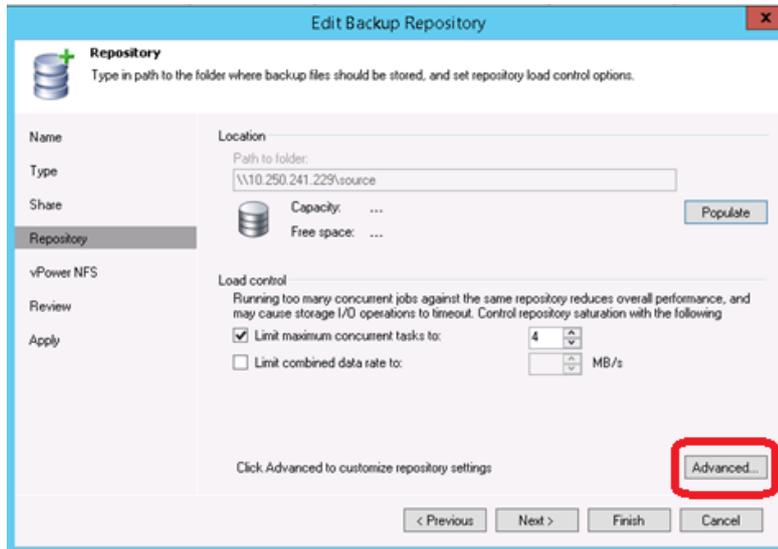
- 6 Select **Shared folder** as the type of backup repository, and click **Next**.



- 7 In the **Shared folder** field, enter the QoreStor container share UNC path (or TCP/IP address to replace hostname), select the **Gateway Server**, and click **Next**.



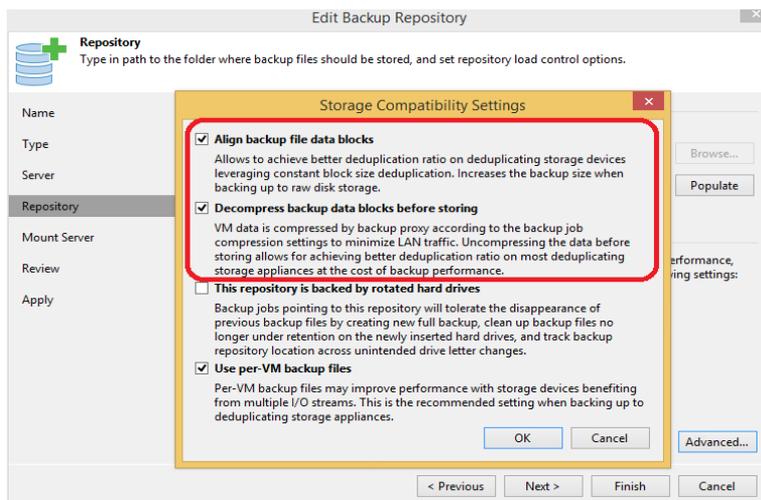
- 8 Customize the repository settings by clicking **Advanced**.



i **NOTE:** Please check the QoreStor Interoperability Guide for the maximum concurrent jobs supported for CIFS/NFS. The maximum concurrent tasks also depend upon the number of CPU cores of Veeam Servers or proxies.

- 9 Check the **Decompress backup data blocks before storing** and **Align backup file data blocks** options:

i **NOTE:** Deselecting the Decompress backup data blocks before storing or the Align backup files data blocks option can negatively impact your overall storage savings and performance. It is especially not recommended to switch these settings after data has been written to QoreStor.

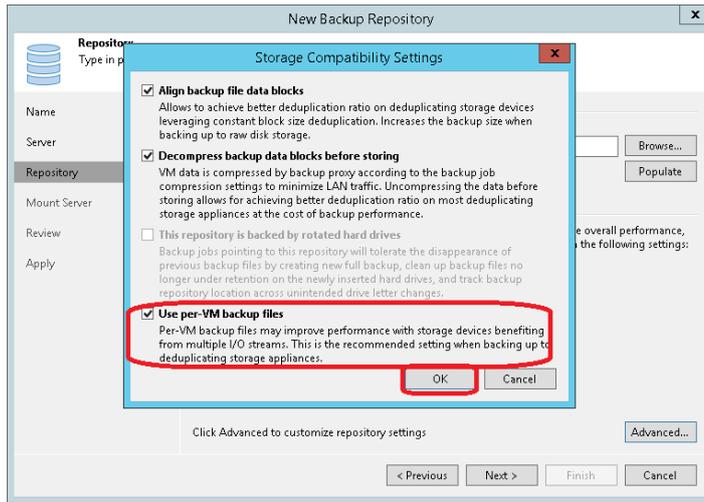




Warning: It is not recommended to change the setting for option Align backup file data blocks after backups are taken as it will impact the deduplication savings for future backups.

10 Check the **Use Per-VM Backup Files** option and click **OK**:

The Per-VM backup file option causes a per-restore point backup file to be created. In other words, this causes each VM's restore point to be placed in a dedicated backup file.



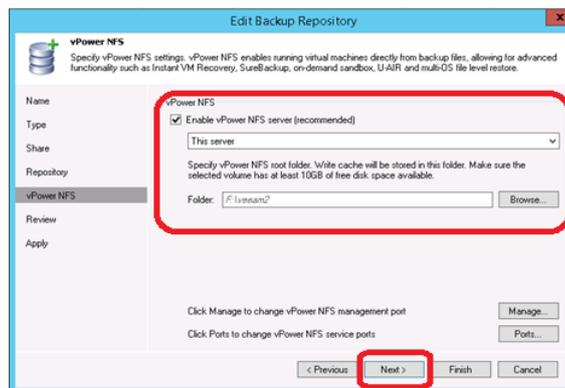
Warning: Make sure to enable the **Enable parallel data processing** option in step 3 if using **Veeam 9.5 U3 or below**



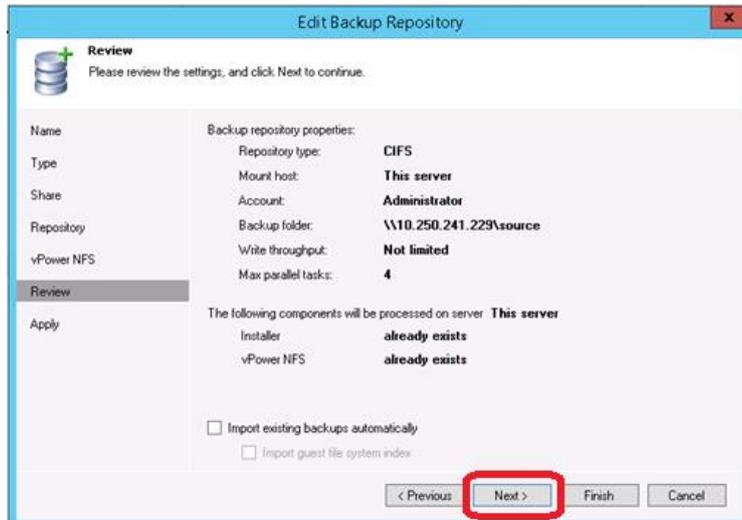
NOTE: This enables multiple write streams within a single job with parallel processing enabled. Enabling multiple streams dramatically improves overall job backup performance. So it is recommended to use per-VM backup files options for better backup throughput.

11 Click **Next**.

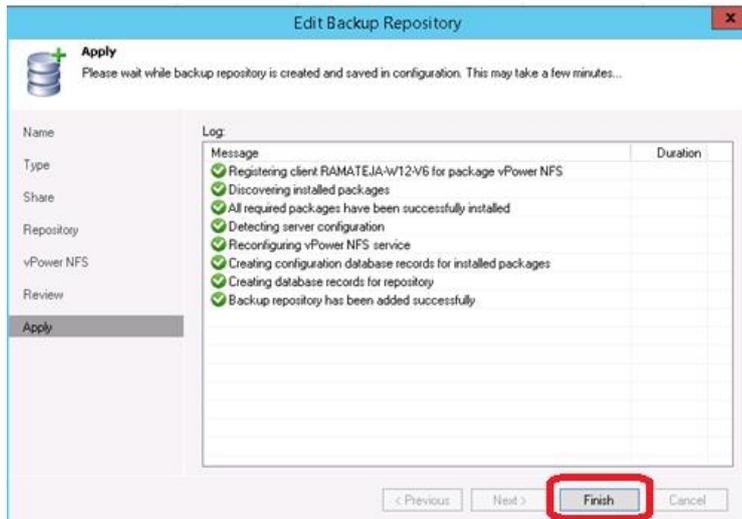
12 If you wish to use the Instant Recovery feature, enable the **vPower NFS** setting.



13 On the review page, verify the settings, and click **Next** to apply changes.

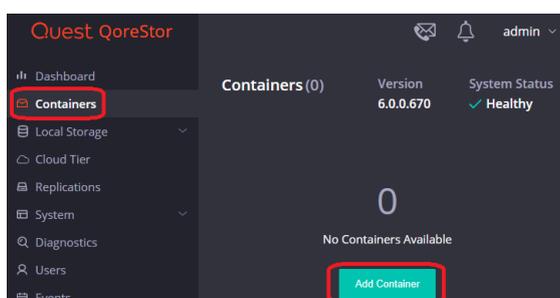


14 Click **Finish**.

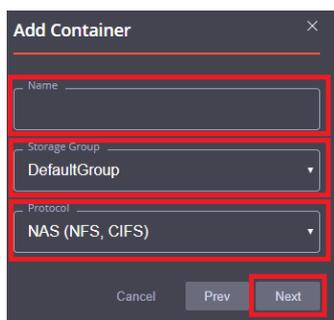


Creating an NFS container for use with Veeam

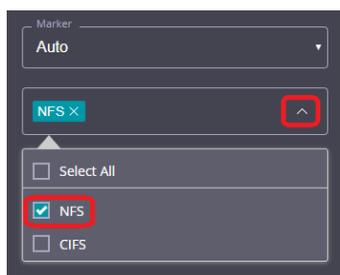
- 1 Select the **Containers** tab, then click **Add container**.



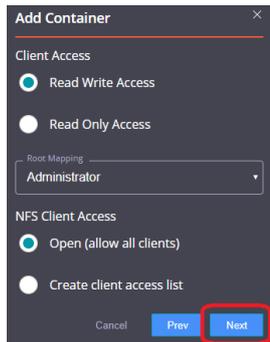
- 2 Enter a container **Name**, select a **Storage Group** or leave the **DefaultGroup** option selected, and select **NAS (NFS, CIFS)** from the **Protocol** dropdown menu. Click **Next**.



- 3 Click the dropdown on the **Access Protocols** field then select the check mark for **NFS**. Leave **Marker Type** on **Auto**, then click **Next**.

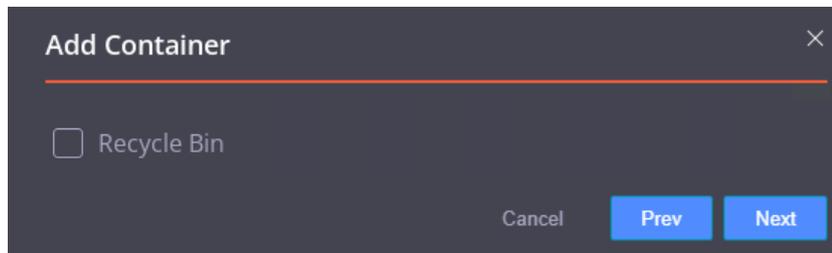


- 4 Fill in the **NFS Client Access** options if need then click **Next**.

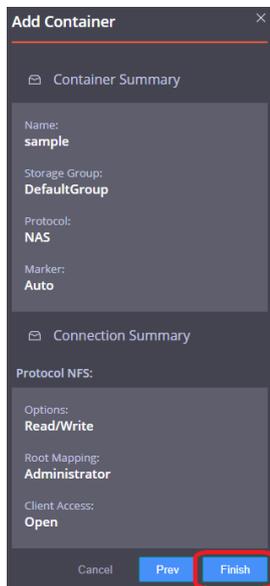


i **NOTE:** For improved security, Quest recommends adding IP addresses for only Veeam servers/proxies

- 5 On this page, the Recycle Bin feature may be enabled, please check the user guide for more information. Click **Next**.



- 6 Confirm the settings and click **Finish**. Confirm that the container is added.



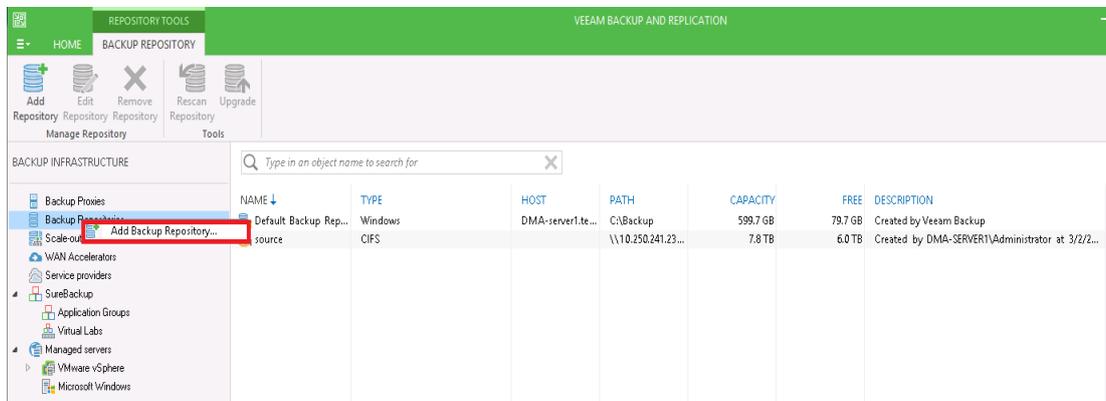
Adding the QoreStor NFS container as a repository in Veeam

NOTE: The Veeam Server is supported on Windows only. To configure an NFS container from QoreStor as a backup repository a Linux server where the NFS container would be mounted is required.

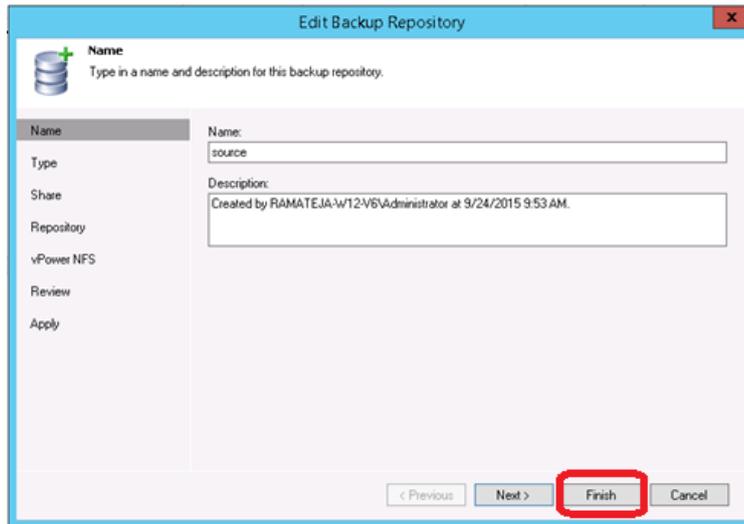
CAUTION: To maximize the QoreStor and Veeam deduplication savings and performance, Quest recommends using the exact settings in this guide for all the data being backed up.

The backup data will change format completely when backup settings are changed. Hence, to get accurate savings numbers, all the data should be backed up with the same settings.

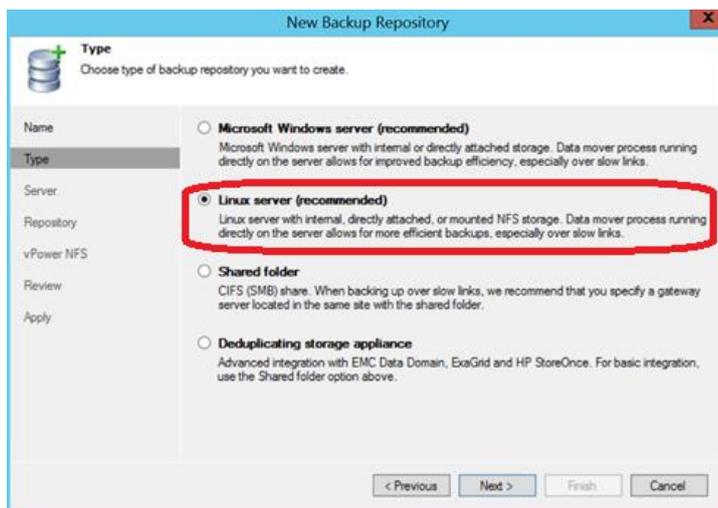
- 1 Open the Veeam Backup & Replication console.
- 2 If using Veeam 9.5 U3 or lower, select the dropdown **Menu** and click **General Options**
- 3 Check the **Enable parallel processing** option in the I/O Control tab and click **OK**. This option will be missing in Veeam 9.5 U4 and higher as it's automatically enabled by default.
- 4 In the **Backup Infrastructure** section, right-click **Backup Repositories**, and select **Add Backup Repository**.



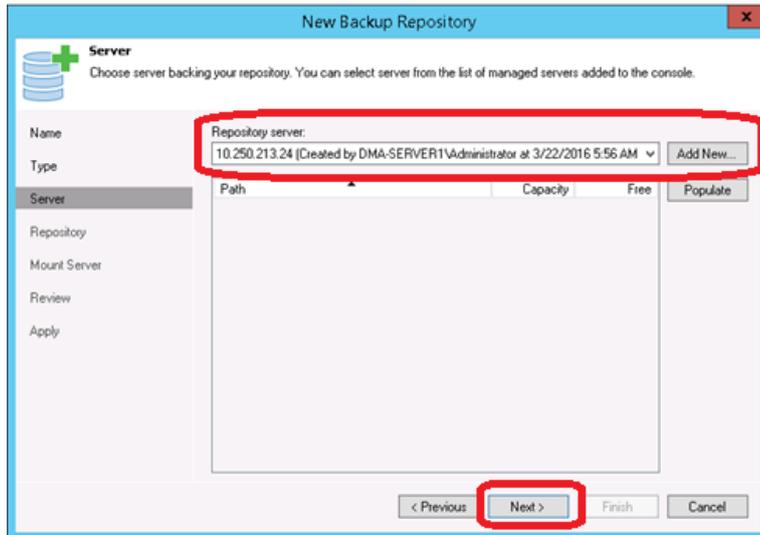
- 5 Enter a name for the QoreStor container repository and click **Next**.



- 6 Select **Linux Server (recommended)** as the type of backup repository, then click **Next**



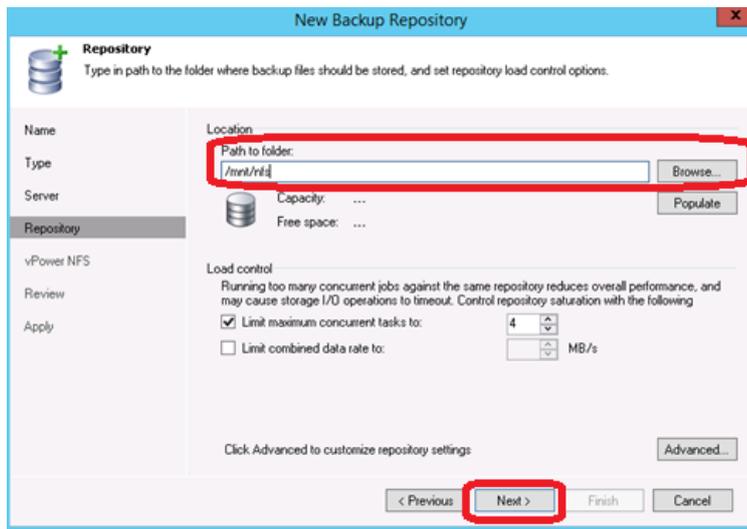
- 7 Add **the New Repository server (Linux)** or select the server from the list if added already.



- Mount the QoreStor NFS Container onto a Linux Server.

```
[root@r320-sys-41 ~]# mkdir /mnt/nfs
[root@r320-sys-41 ~]# mount -t nfs 6300-07:/containers/sample /mnt/nfs
[root@r320-sys-41 ~]# █
```

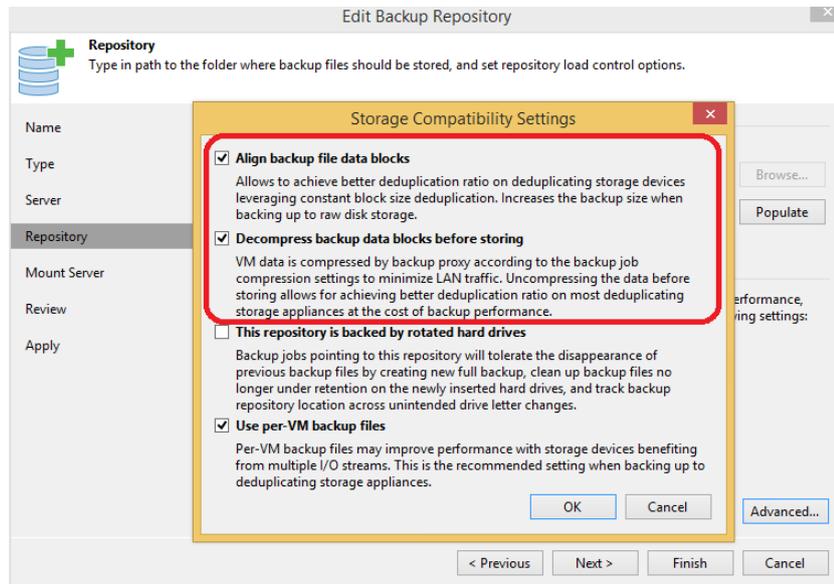
- Enter the container mount path. Then customize the repository settings by clicking the **Advanced** button.



i **NOTE:** Please check the QoreStor Interoperability Guide for the maximum concurrent jobs supported for CIFS/NFS. The maximum concurrent tasks also depend upon the number of CPU cores of Veeam Servers or proxies.

- 10 Check the **Decompress backup data blocks before storing** and **Align backup file data blocks** options:

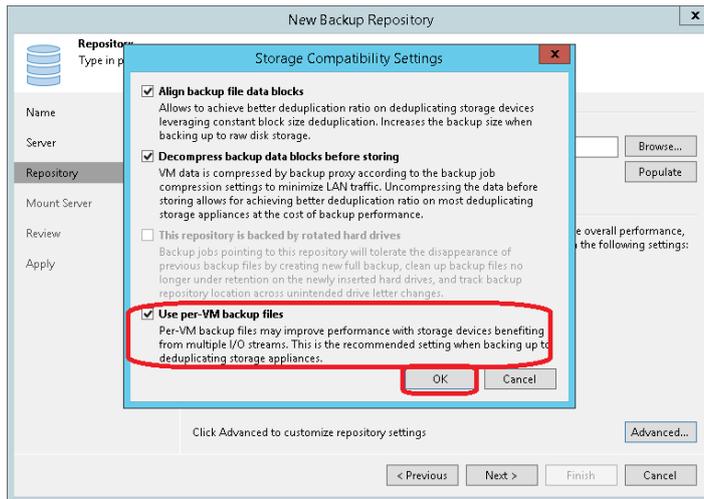
i **NOTE:** Deselecting the Decompress backup data blocks before storing or the Align backup files data blocks option can negatively impact your overall storage savings and performance. It is especially not recommended to switch these settings after data has been written to QoreStor.



! **Warning:** It is not recommended to change the setting for option **Align backup file data blocks** after backups are taken as it will impact the deduplication savings for future backups.

- 11 Check the **Use Per-VM Backup Files** option and Click **OK**:

The Per-VM backup file option causes a per-restore point backup file to be created. In other words, this causes each VM's restore point to be placed in a dedicated backup file.

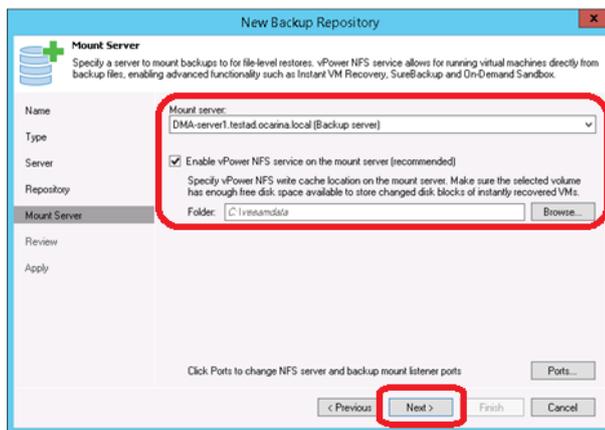


Warning: Make sure to enable the Enable parallel data processing option in step 3 if using Veeam 9.5 U3 or below

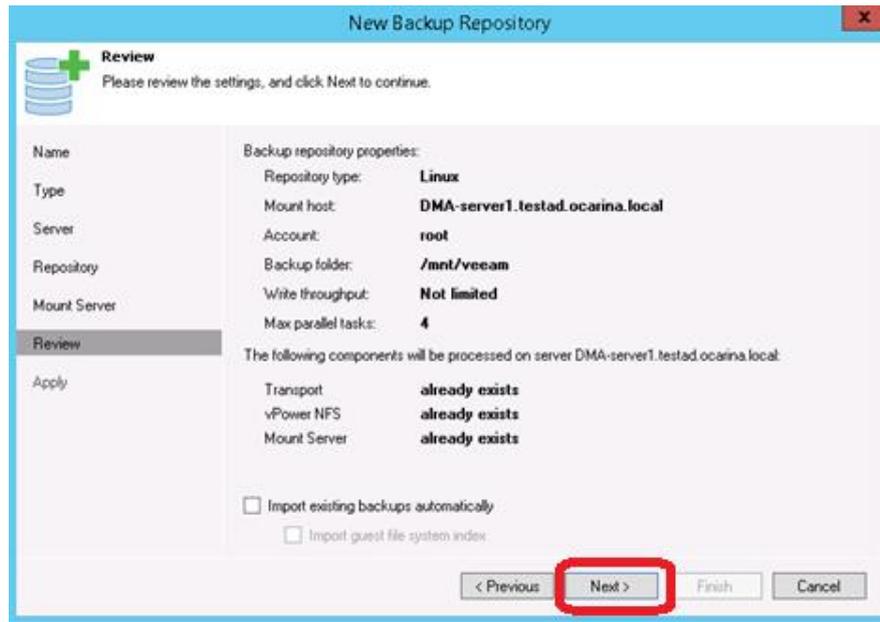
NOTE: This enables multiple write streams within a single job with parallel processing enabled. Enabling multiple streams dramatically improves overall job backup performance. So it is recommended to use per-VM backup files options for better backup throughput.

12 Click **Next**.

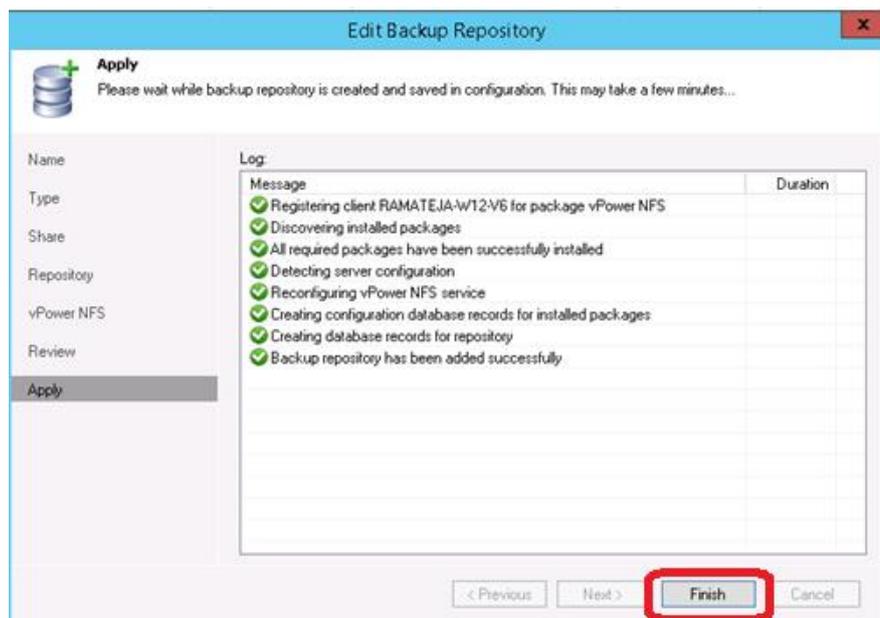
13 Optionally, if you wish to use the Instant Recovery feature, enable the **vPower NFS** service setting.



14 On the review page, verify the settings, and click **Next** to apply changes.



15 Click **Finish**.



Configuring Rapid CIFS for Veeam

Rapid CIFS is a Quest-developed protocol that accelerates writes to CIFS shares on the QoreStor system. This is done by only sending unique data to the appliance. This usually causes significant network savings and even sometimes performance boosts.

Windows prerequisites

- The Media Agent OS must be the 64-bit version of Windows 2008 R2, 2012/R2, 2016, or 2019.



NOTE: For the accelerator to work properly, the backup traffic must go directly to the QoreStor system. For Veeam, you should install RDCIFS on the Veeam Proxy pushing the data. Install location can depend on the transport mode used. For network mode, it is installed on the Veeam server itself. For HotAdd mode it needs to be installed on the HotAdd proxy in the virtual environment. For SAN mode it needs to be installed on the Veeam Server/Proxy which has direct access to the SAN storage. For Off-Host it needs to be installed on the Veeam Proxy pushing the data, for On-host it should be installed on the Hyper-V server or cluster being backed up.

Installing Rapid CIFS on a Veeam Windows Proxy

The Secure Connect feature is a set of client and server components that creates a secure channel for QoreStor communication with WAN-connected clients that is also resilient to WAN outages. This is generally only suggested for use over WAN.

Follow these steps to install Rapid CIFS.

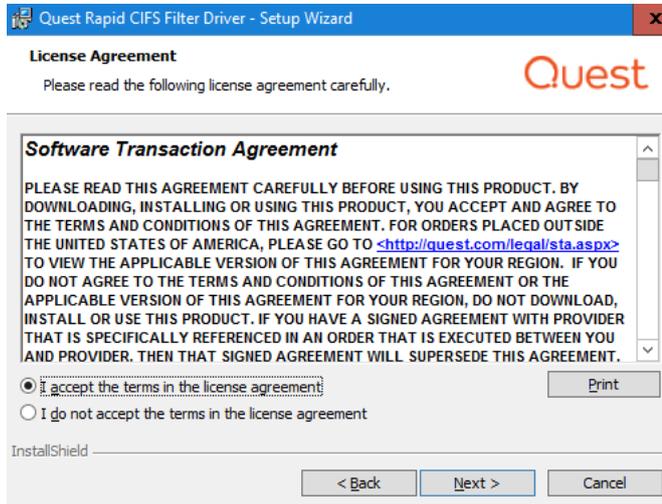


NOTE: Rapid CIFS should only be installed on a Veeam server or Proxy.

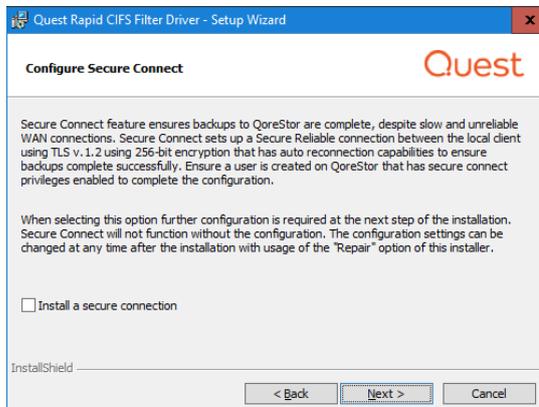
- 1 Download the MSI to the Server/Proxy by doing the following:
 - a Go to support.quest.com/qorestor/ and select your version.
- 2 On the support page for your product, click **Software Downloads**.
- 3 For the RDCIFS plugin for your QoreStor version, click the **Download** icon to download the installer package (.exe file).
- 4 Run the EXE and follow the instructions in the installation wizard as shown in the screenshots below. Click **Next** on the first screen.



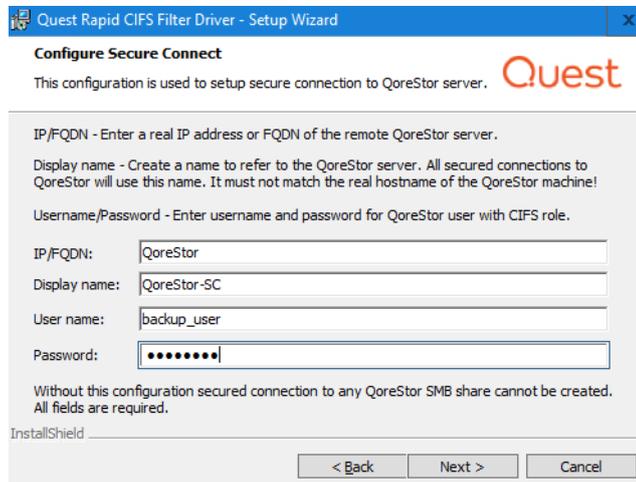
- 5 Read and accept the license agreement to proceed. Click **Next** when ready.



- 6 If installing with secure connections for WAN use check the securely connect box. Click **Next**.

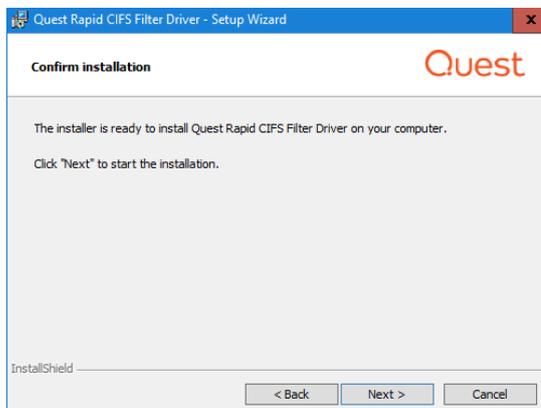


- a If installing with a secure connection insert the **IP/FQDN**. The **Display Name** field will auto-populate from the **IP/FQDN** field. The default **username** and **password** are *backup_user* and *St0r@ge!* (With a zero in place of the letter O).



- i** **NOTE:** When accessing the share from this server use the **Display Name** when accessing the share to leverage Secure connect. I.E //QoreStor-SC/share
Use the normal IP/FQDN to access WITHOUT a secure connection.

- 7 Click **Next**.



- 8 After the installation finishes click **Finish**. You can optionally verify that the “**rdcifsfd**” driver is loaded automatically; this can be checked by using the command **fltmc**.

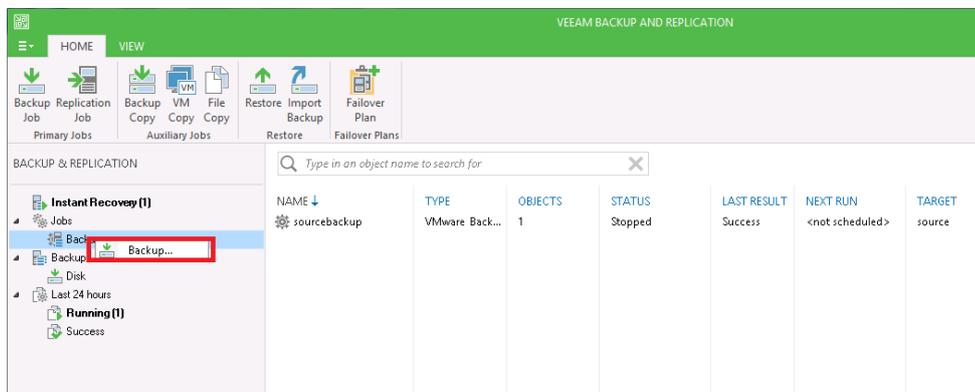
```

Administrator: Command Prompt
Microsoft Windows [Version 6.2.9200]
(c) 2012 Microsoft Corporation. All rights reserved.
C:\Users\Administrator>fltmc
Filter Name                               Num Instances  Altitude  Frame
-----
rdcifsfd                                  1             301600    0
luaflt                                    1             135000    0
npsvcctrl                                  1             46000     0
C:\Users\Administrator>_

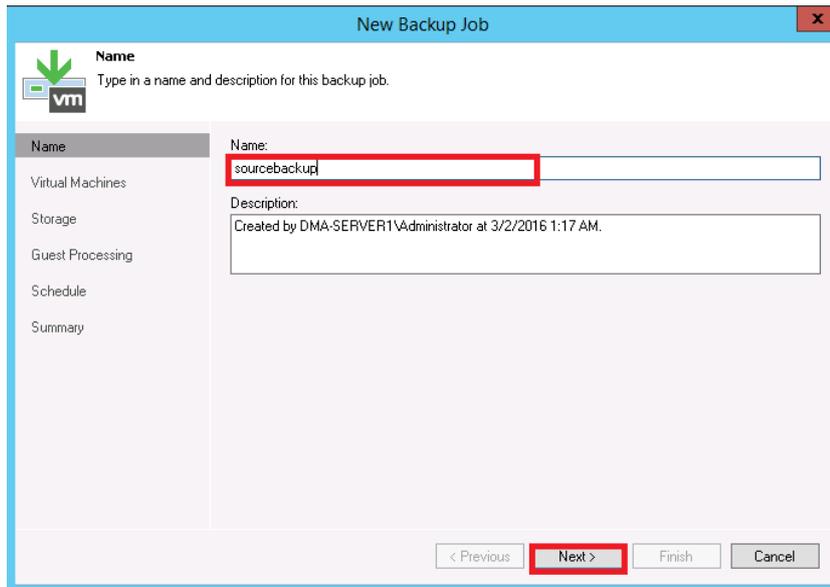
```

Creating a backup job with the QoreStor system as a target

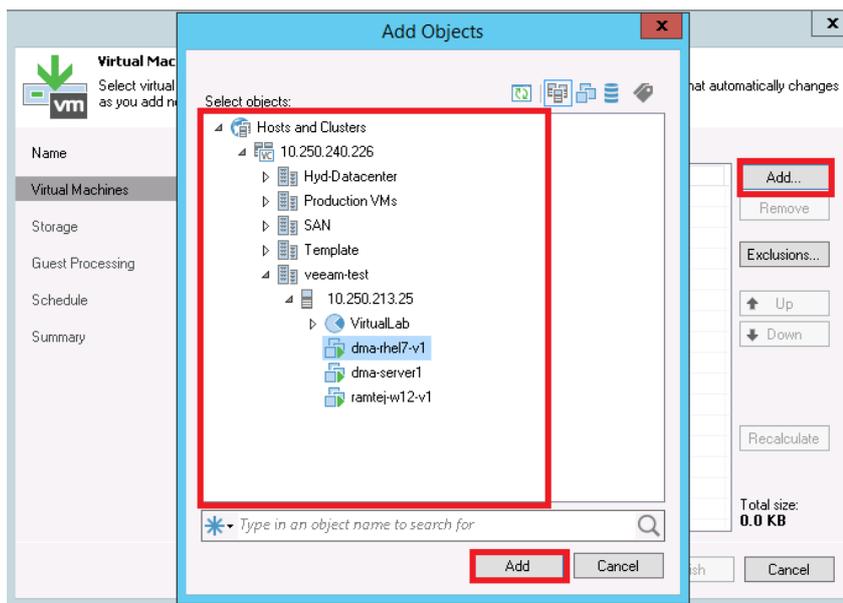
- 1 On the **Backup & Replication** menu, go to **Jobs > Backup**, and right-click **Backup** to create a new backup job.



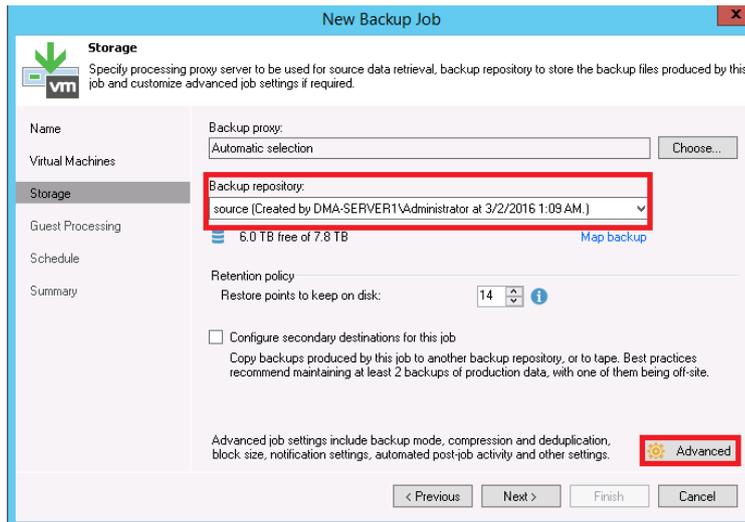
- 2 Provide the backup job name and click **Next**.



- 3 Select one or more virtual machines, data stores, resource pools, vApps, SCVMM clusters, etc. for backup.



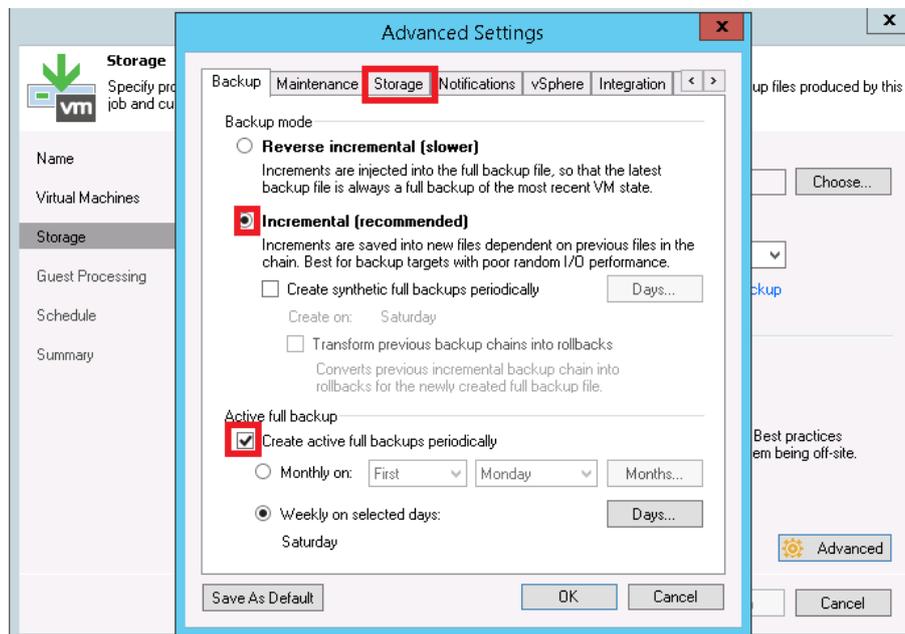
- 4 Select the QoreStor container share as the Backup Repository for this job and click **Advanced**.



- 5 On the **Backup** tab, make sure **Incremental** and **Create active full backups periodically** are selected. Set the active full schedule to whatever is needed.



NOTE: It is recommended to enable Active Full backups once a week with a Veeam Ready Archive QoreStor instance. The active full backup produces a full backup of a VM just as if it was running for the first time. The Synthetic full backup option is only suggested to be used with a Veeam Ready Repository QoreStor instance due to read performance requirements during the synthetic operation.

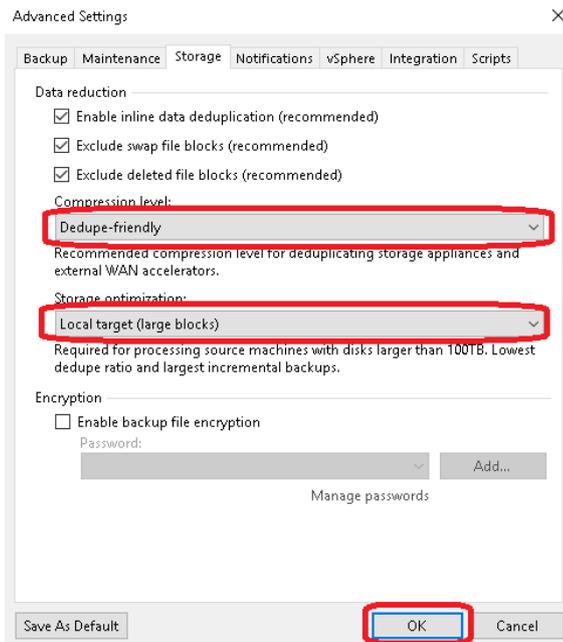


Warning: For information on configuring Fast Clone options for Hyper-V 2016 ReFS VM's please review the Fast Clone section of this document



Warning: Veeam generally recommends against very long retention combined with infrequent active or synthetic full backups. Generally speaking a full should be run at least once a month but contacting Veeam for their recommendation is suggested.

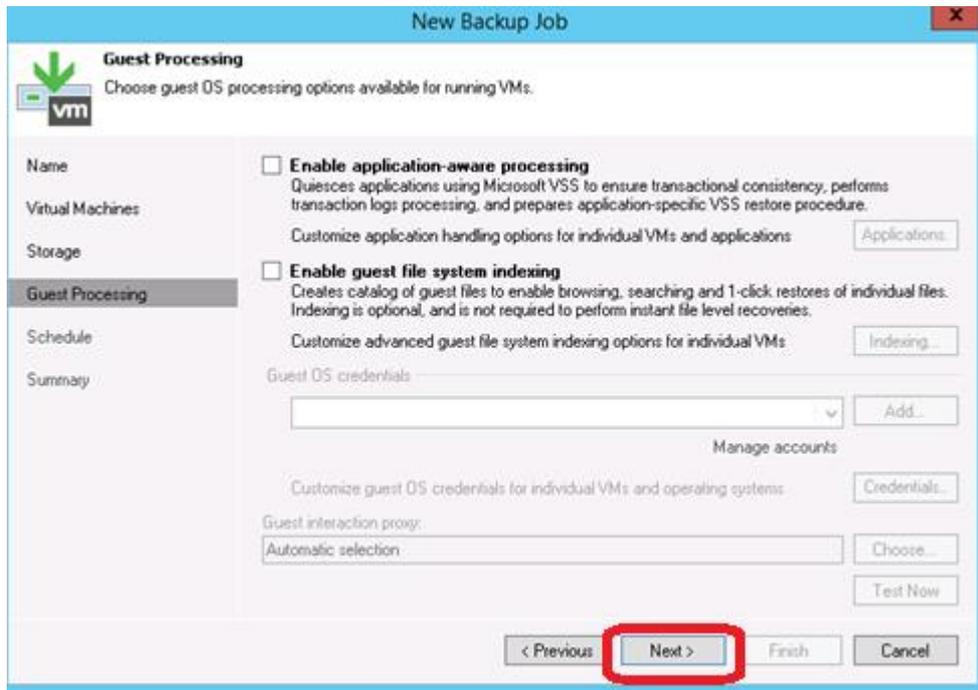
- 6 On the Storage tab, do the following:
 - a Under Deduplication, select Enable inline data deduplication.
 - b Under Compression, set the Level to Dedupe-Friendly.
 - c Under Storage optimizations, set Optimization to Local target (large blocks).



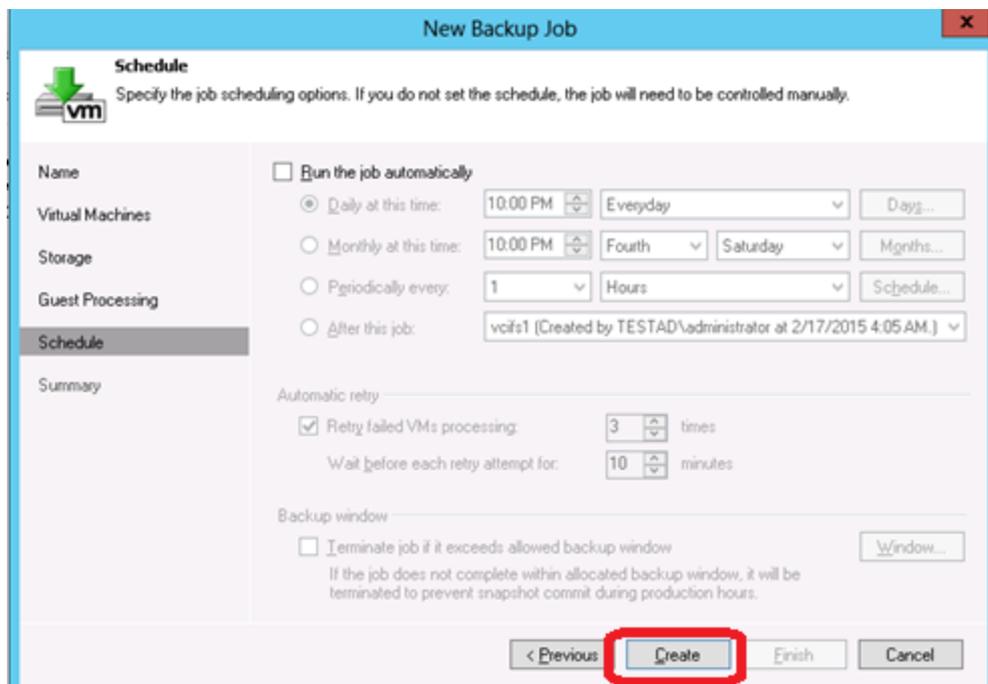
NOTE: For the best balance between backup performance and deduplication savings it is recommended to choose these options for all of the backup jobs written to QoreStor.

Normally, Quest recommends turning off encryption, compression, and deduplication in all data management applications. However, with Veeam, Quest recommends enabling deduplication. This is because Veeam runs deduplication at larger block sizes, and deduplication of these large blocks does not heavily impact QoreStor duplication results. In addition, this reduces network bandwidth utilization when Veeam sends data to the QoreStor system, this benefits the backup performance overall.

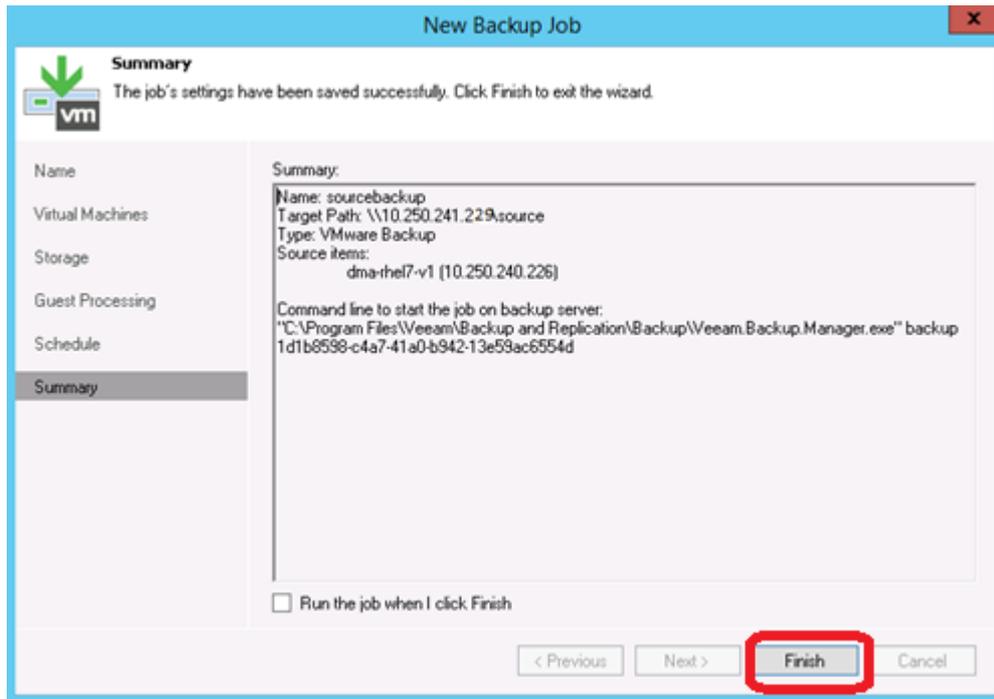
7 Enable any optional settings required by your workflow and click **Next**.



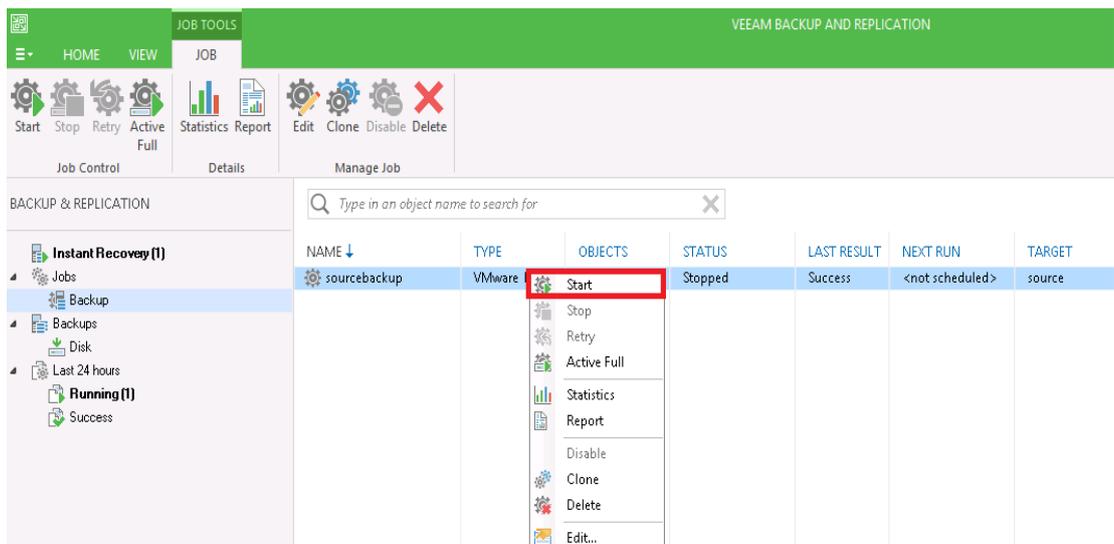
8 Schedule the backup and click **Create**.



9 Click **Finish**.



10 To Run the Backup manually, right-click the backup job configured and select **Start**.

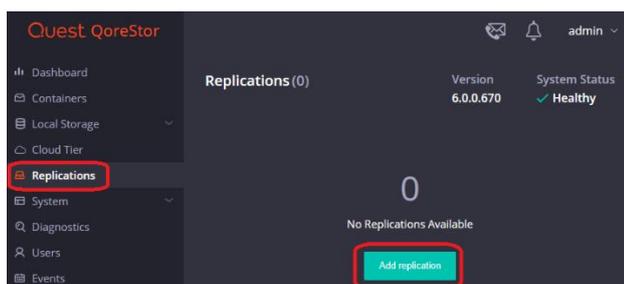


Setting up QoreStor system replication

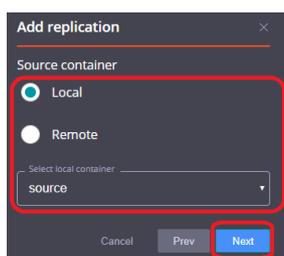
i **NOTE:** For the steps in this procedure, assume QS1 is the replication source QoreStor system, and QS2 is the replication target QoreStor system. 'source' is the replication source container, and 'target' is the replication target container.

Creating a CIFS/NFS replication session

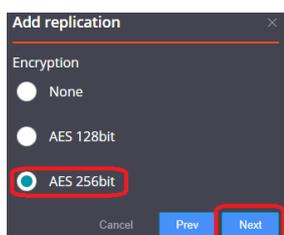
- 1 Create a source container on the source QoreStor system.
- 2 Create a target container on the target QoreStor system.
- 3 On the source QoreStor system, go to the **Replications** Tab. Click the **Add replication** button.



- 4 Select the source Container for Replication and click **Next**.



- 5 Select the **Encryption** type for the Source Container and click **Next**.



- 6 Enter the target QoreStor systems-related information then click **Retrieve Remote Containers**. Select a target container from the populated list and click **Next**.

Add replication

Target container

Local

Remote

Username
admin

Password

Remote Machine
qspl-6300-47.systest.ocarina.local

Retrieve Containers

Select remote container
target

Cancel Prev Next

- 7 Specify any **Bandwidth Limitations** needed in MBps, and leave 0 for unlimited bandwidth. Click **Next**.

Add replication

Limits

Bandwidth Limit for the peer machine [MBps]
0

Cancel Prev Next

- 8 Verify the Summary and click **Finish**.

Add replication

Source container

Source
Local

Container
source

Encryption

Algorithm
AES 256

Limits

Bandwidth Limit [MBps]
Unlimited

Target

Source
Remote

Container
target

User
admin

Password

Remote Machine Address
qspl-6300-47.systest.ocarina.local

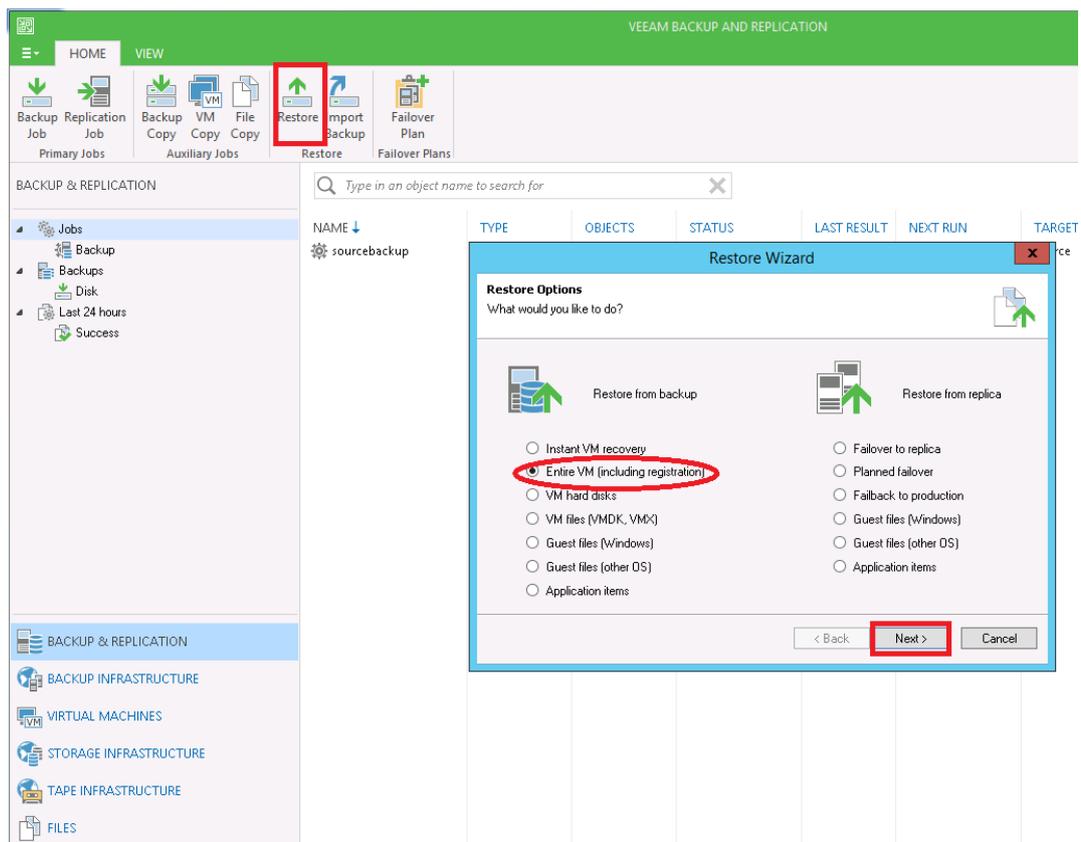
Cancel Prev Finish

- 9 Check replication is added successfully and confirm the replication details.

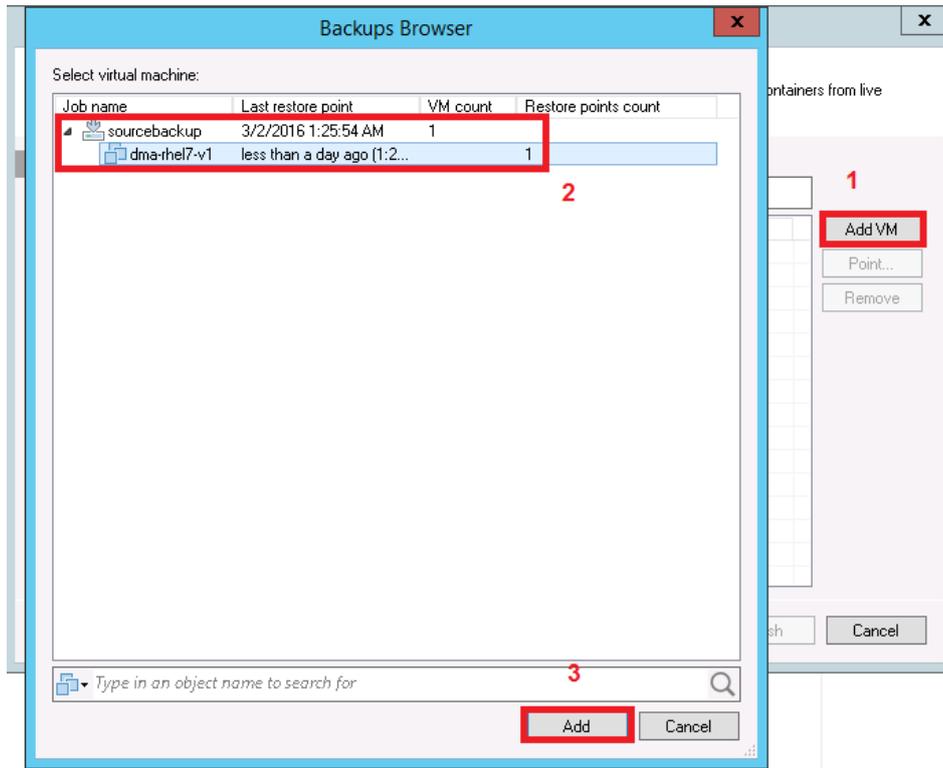
Restoring from the replication target

i **NOTE:** Before restoring from the target QoreStor system, make sure that the replication session state is INSYNC on the QoreStor system GUI Replication menu. Stop or Delete the replication session, and make sure that the target QoreStor system container has the CIFS/NFS connection(s) enabled.

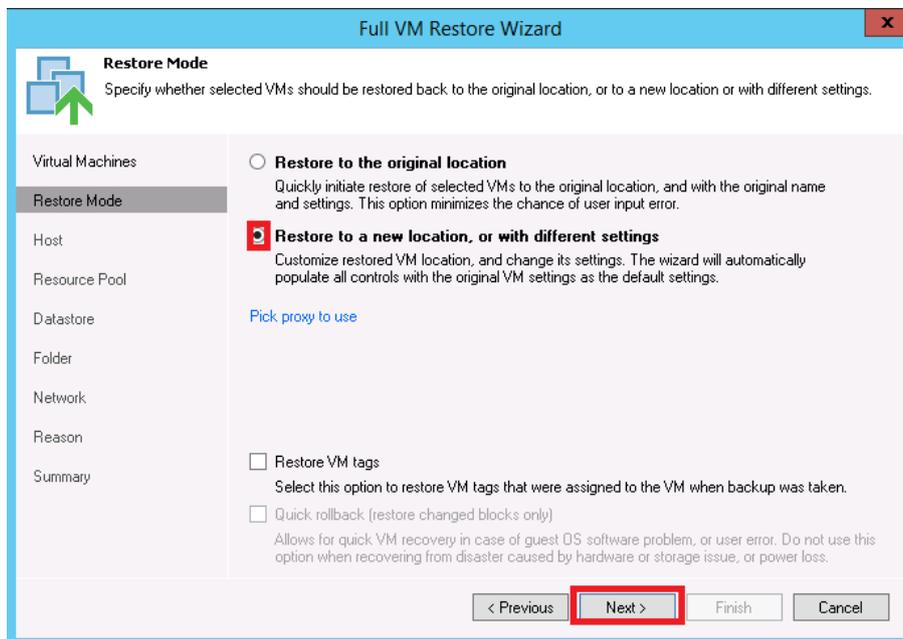
- 1 Add the target QoreStor system container to the Veeam repository. For instructions, see the above sections [Creating a CIFS container for use with Veeam](#) or [Creating an NFS container for use with Veeam](#).
- 2 Update all backup jobs that use the source QoreStor system container as a repository and change them to use the target QoreStor container as the backup repository.
- 3 Under **Backup & Replication**, click **Restore** to create a restore job. Select the appropriate restore from the backup option.



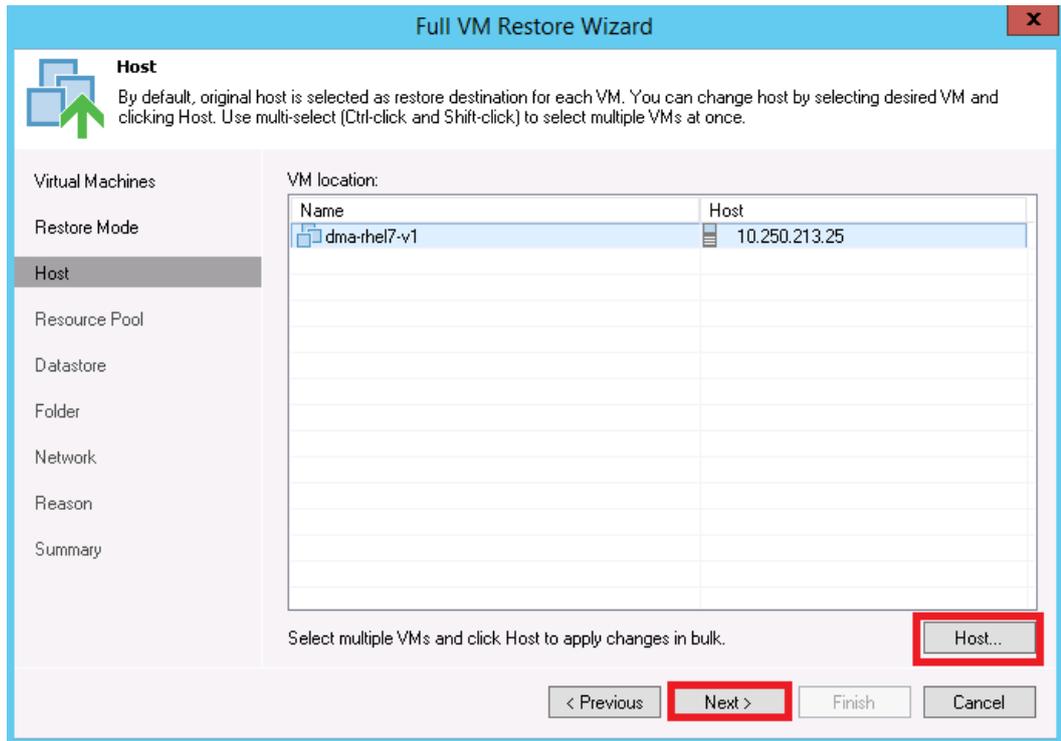
- 4 Click Add VM and select **From backup**. Select the VM to be restored and click **Add**.



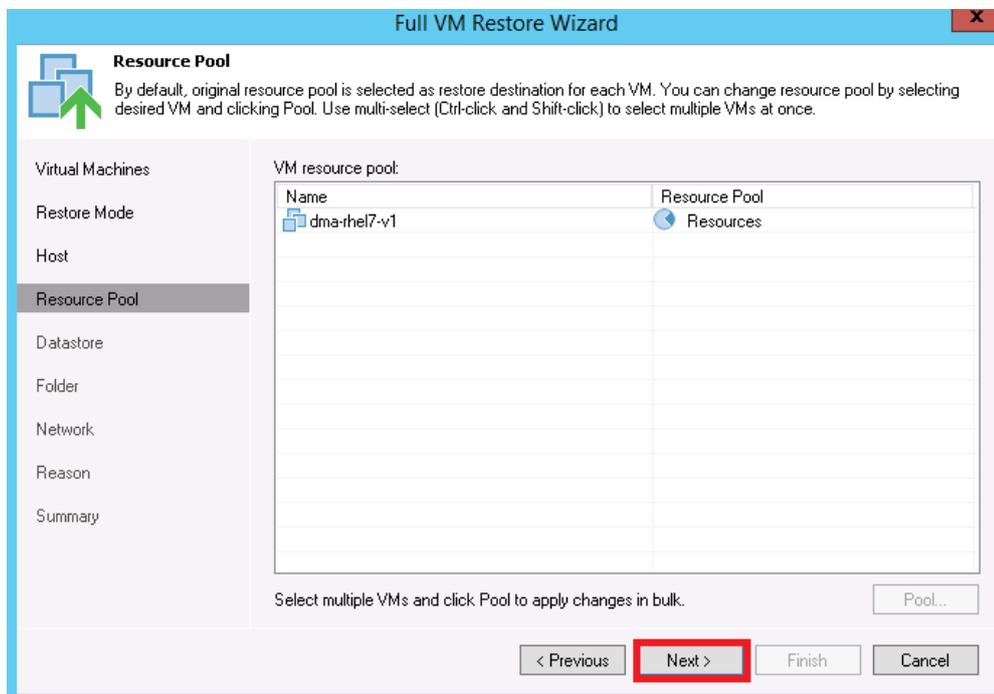
- 5 Select the **Restore Mode** and click **Next**.



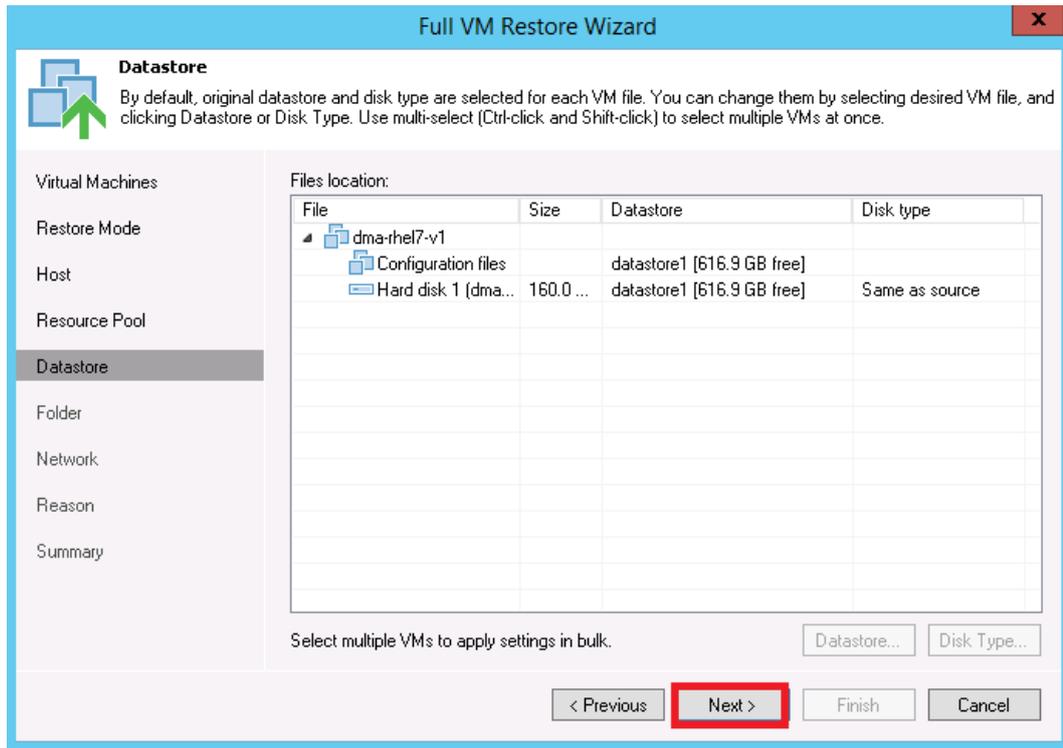
- 6 Provide the Host details as per requirement and click **Next**.



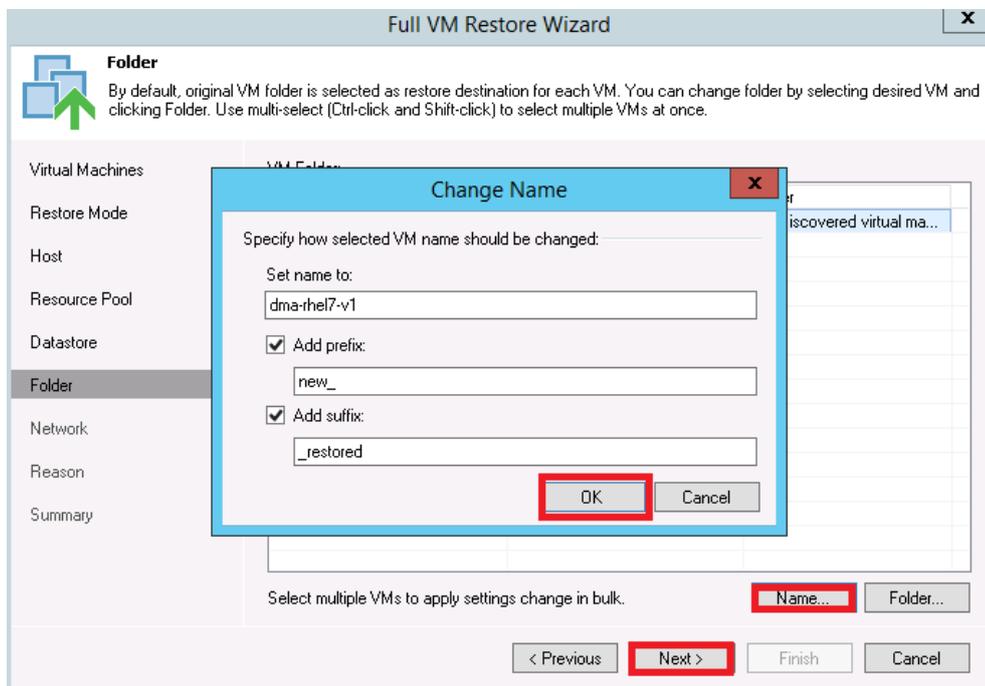
- 7 Select the resource pool and click **Next**.



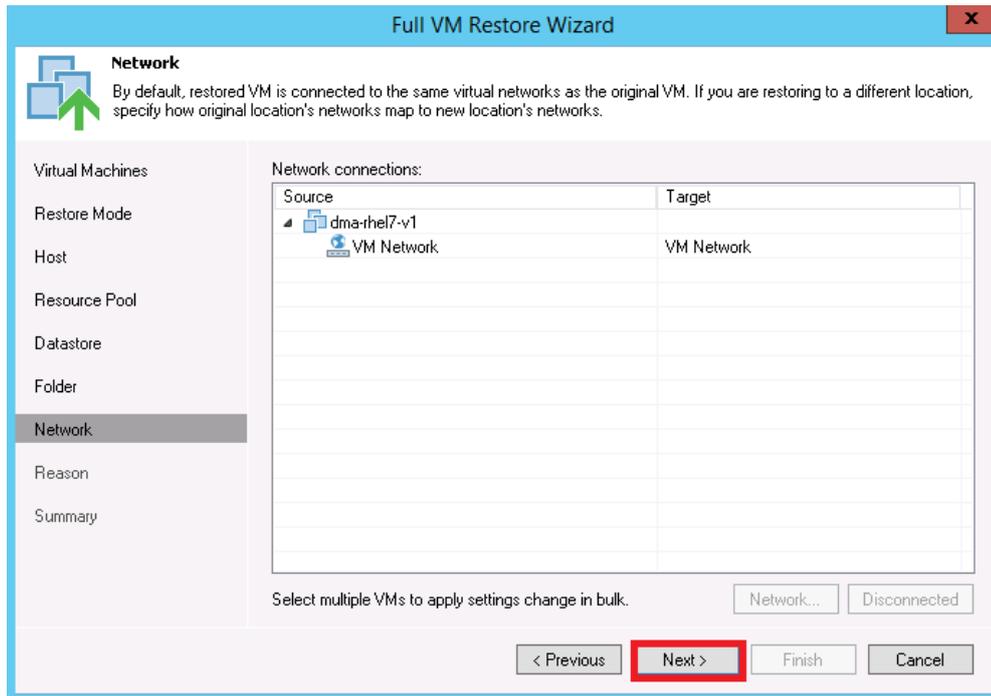
- 8 Select the data store, and disk type and click **Next**.



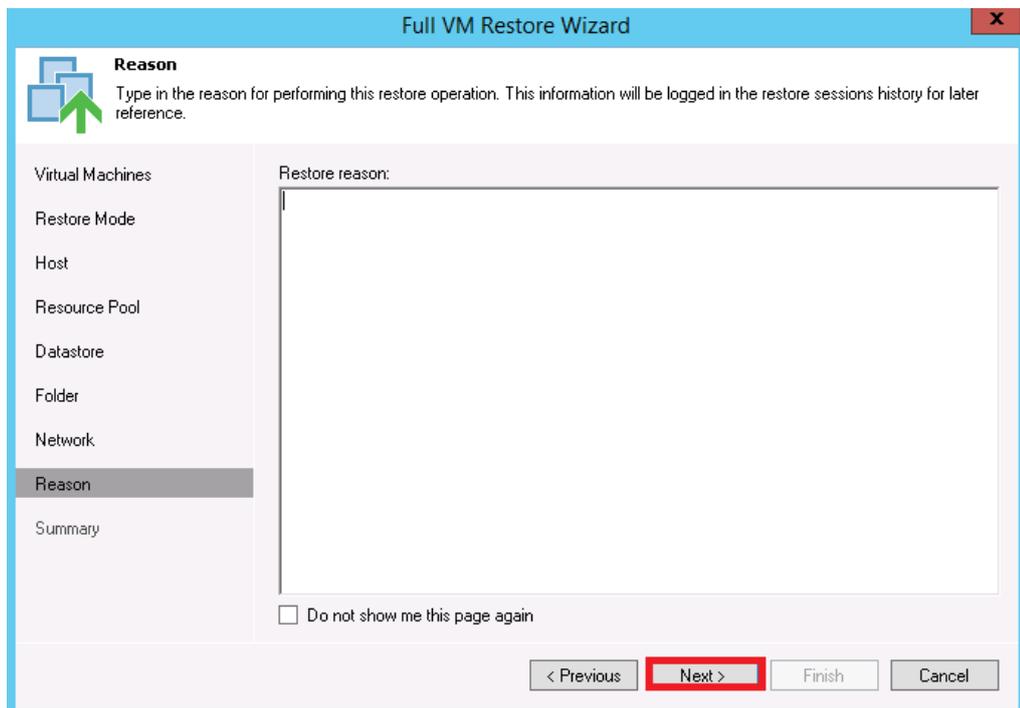
- 9 Provide the new name for the restored VM and click **Next**.



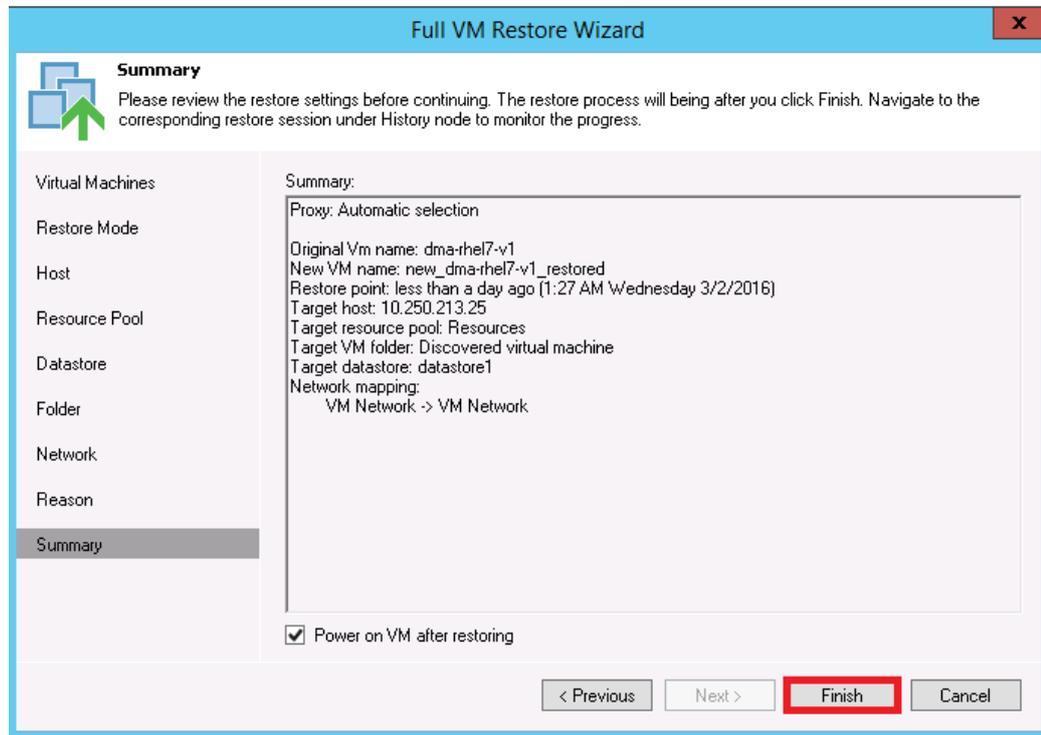
10 Select the network location and click **Next**.



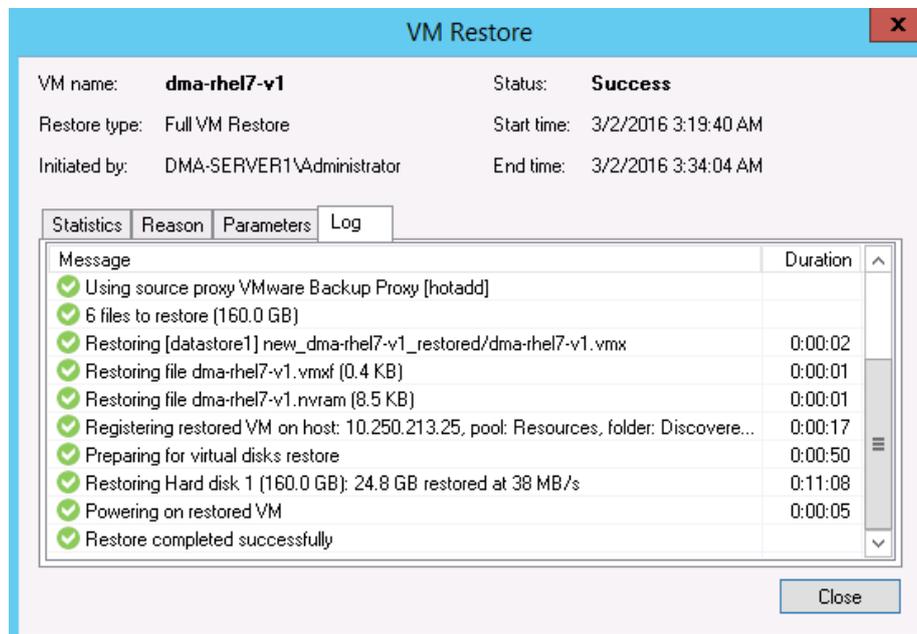
11 Provide the reason for the restoration.



12 Click **Finish**.



13 After the restore job has been created, you can run the job and monitor it from the **Backup & Replication** menu.



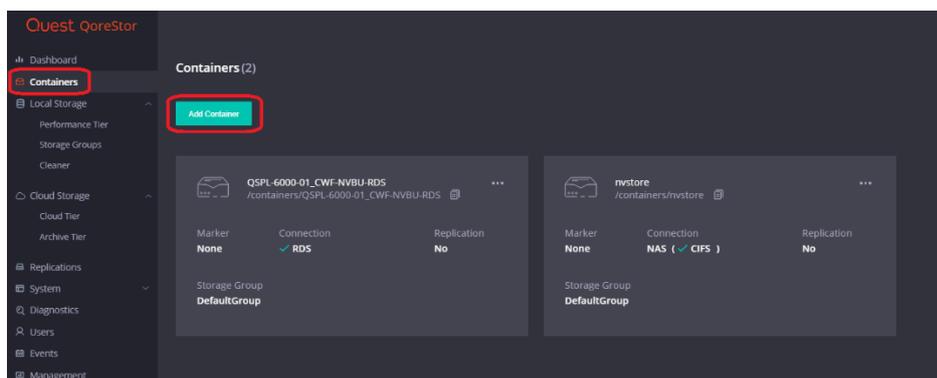
Using QoreStor as a Veeam Scale-Out Capacity Tier via Object Container(S3)

Scale-Out Repositories are a Veeam feature that allows you to transition data from one repository to another via policies defined in Veeam. This could be used with the QoreStor performance tier to move data into a slower QoreStor tier or with spindled disk-to-tier initial backups to QoreStor. In this section, we will cover using the new Object Container QoreStor feature to allow Veeam to write via S3 to QoreStor as a scale-out capacity tier.

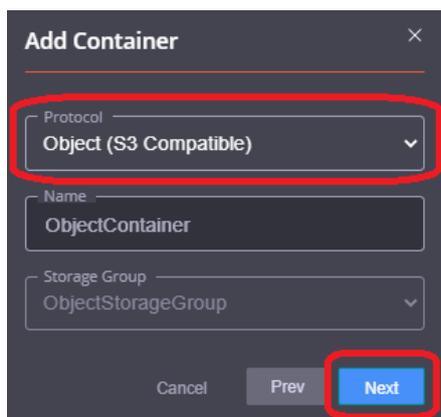
Scale-Out repositories work by first creating basic repositories. Then you create a scale-out repository adding the initial performance tiers and capacity tiers already added as basic repositories.

Creating an Object Container(S3) in QoreStor

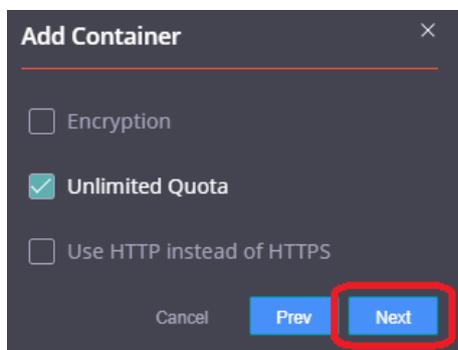
1 From the QoreStor UI select **Containers** then click **Add Container**.



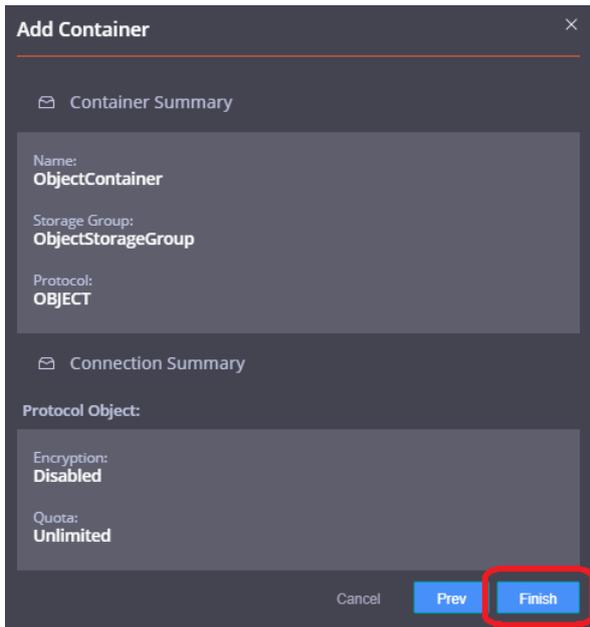
2 Select the **Protocol** dropdown and set it to **Object (S3 Compatible)**. Click **Next**.



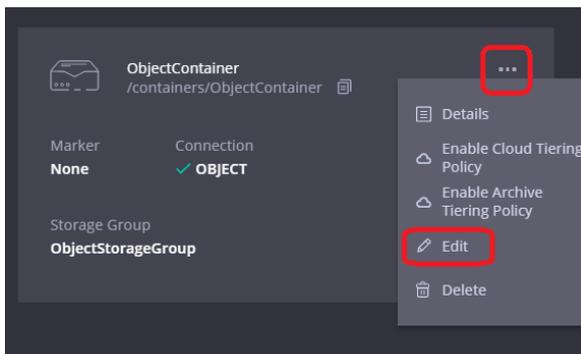
3 Click **Next**.



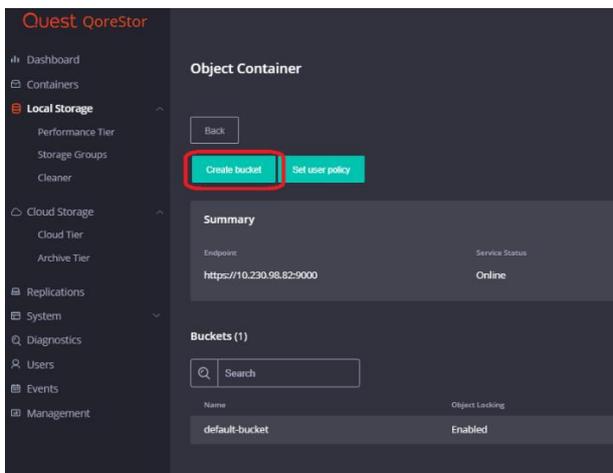
4 Verify the summary is correct and click **Finish**.



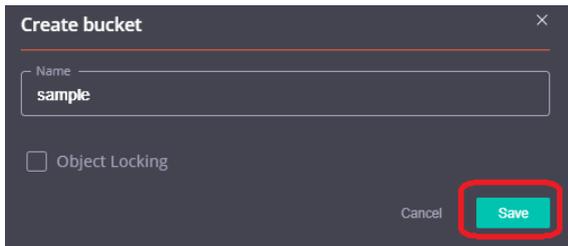
5 The Object Container is now created but we need to create a bucket other than the default. Click the **ellipsis** on the container and click **Edit**.



6 On the Object Container page click **Create bucket**.

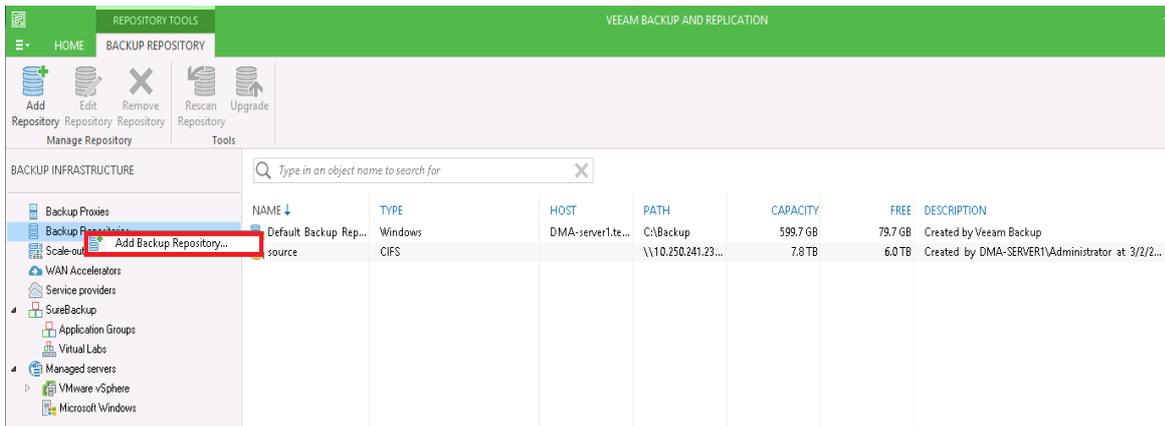


7 Name the bucket then click **Save**.



Adding the QoreStor Object Container(S3) as a repository in Veeam

1 In the Backup Infrastructure section, right-click Backup Repositories, and select Add Backup Repository.



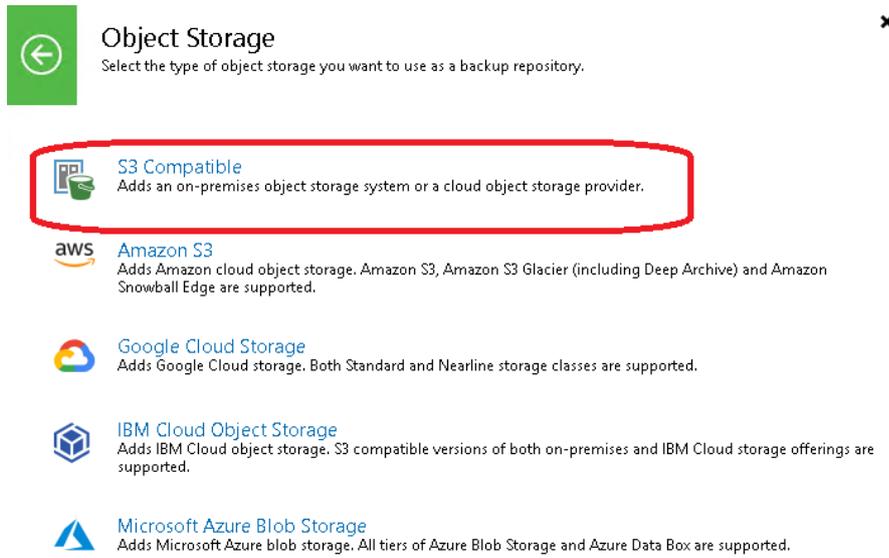
2 Click Object storage.

Add Backup Repository

Select the type of backup repository you want to add.

-  **Direct attached storage**
Microsoft Windows or Linux server with internal or direct attached storage. This configuration enables data movers to run directly on the server, allowing for fastest performance.
-  **Network attached storage**
Network share on a file server or a NAS device. When backing up to a remote share, we recommend that you select a gateway server located in the same site with the share.
-  **Deduplicating storage appliance**
Dell EMC Data Domain, ExaGrid, HPE StoreOnce or Quantum DXi. If you are unable to meet the requirements of advanced integration via native appliance API, use the network attached storage option instead.
-  **Object storage**
On-prem object storage system or a cloud object storage provider. Object storage can only be used as a Capacity Tier of scale-out backup repositories, backing up directly to object storage is not currently supported.

3 Click S3 Compatible.

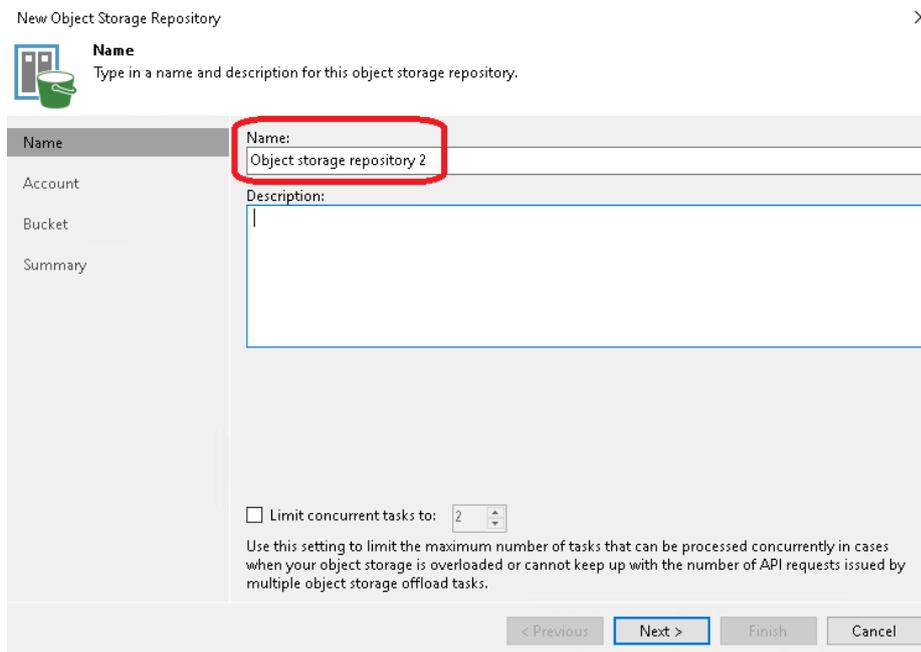


Object Storage ✕

Select the type of object storage you want to use as a backup repository.

- S3 Compatible**
Adds an on-premises object storage system or a cloud object storage provider.
- Amazon S3**
Adds Amazon cloud object storage. Amazon S3, Amazon S3 Glacier (including Deep Archive) and Amazon Snowball Edge are supported.
- Google Cloud Storage**
Adds Google Cloud storage. Both Standard and Nearline storage classes are supported.
- IBM Cloud Object Storage**
Adds IBM Cloud object storage. S3 compatible versions of both on-premises and IBM Cloud storage offerings are supported.
- Microsoft Azure Blob Storage**
Adds Microsoft Azure blob storage. All tiers of Azure Blob Storage and Azure Data Box are supported.

4 Define an object storage repository device name then click **Next**.



New Object Storage Repository ✕

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

Name: Object storage repository 2

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

- 5 Click **Add** on the credentials line.

New Object Storage Repository

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

Name

Account

Bucket

Summary

Service point:
https://hostname:9000

Region:
us-east-1

Credentials:
backup_user (last edited: 320 days ago) Add...

[Manage cloud account:](#)

Use the following gateway server:
R720-40.systest.ocarina.local (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

- 6 Add the username with the object role in QoreStor in the **Access Key** line. Add the password for that user to the **Secret** line. By default this password is St0r@ge! (The "0" in the password is the numeral zero).

Credentials

Access key: backup_user

Secret key: ●●●●●●●●

Description:

OK Cancel

- 7 Add the QoreStor access information to the **Service Point** line. This is usually `https://<hostname>:9000` or `https://<ipAddress>:9000` then click **Next**.

New Object Storage Repository ×

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

Name	Service point: <input type="text" value="https://hostname:9000"/>
Account	Region: <input type="text" value="us-east-1"/>
Bucket	Credentials: <input type="text" value="backup_user (last edited: 320 days ago)"/> Add...
Summary	Manage cloud accounts

Use the following gateway 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.

8 If you get a certificate security alert, click **Continue**.

Certificate Security Alert ×

 Site certificate cannot be verified. Continue anyway?

Remote certificate chain errors:
 UntrustedRoot (A certificate chain processed, but terminated in a root certificate which is not trusted by the trust provider.)

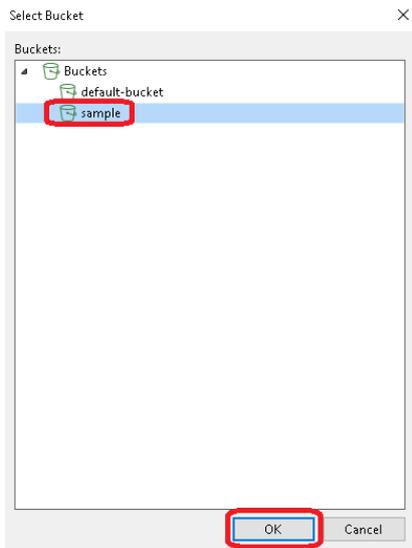
9 On the bucket, page click **Browse...** under the bucket line.

New Object Storage Repository ×

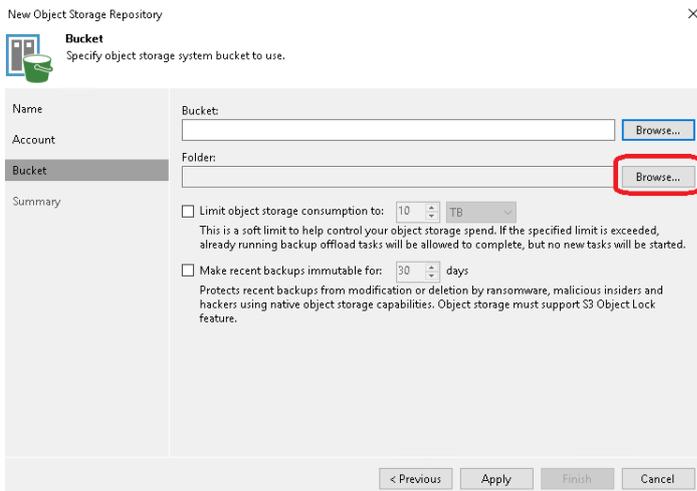
Bucket
Specify object storage system bucket to use.

Name	Bucket: <input type="text"/> Browse...
Account	Folder: <input type="text"/> Browse...
Bucket	<input type="checkbox"/> Limit object storage consumption to: <input type="text" value="10"/> TB <small>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.</small>
Summary	<input type="checkbox"/> Make recent backups immutable for: <input type="text" value="30"/> days <small>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.</small>

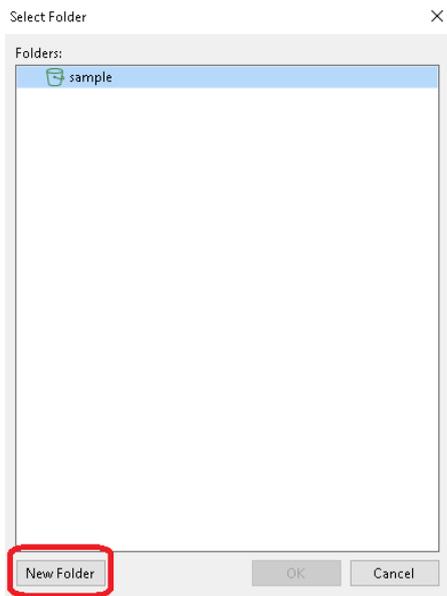
10 Select the bucket name created in the **Creating an Object Container(S3) in QoreStor** section of this guide. Click **OK**.



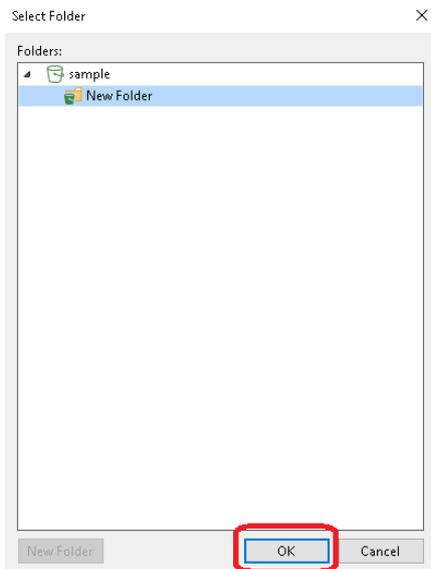
11 Back on the bucket page click **Browse...** under the folder line



12 Click **New Folder** and define a folder name.



13 Select the newly created folder and click **OK**.



14 Back on the bucket page click **OK**.

New Object Storage Repository ×

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

Name	Bucket: <input type="text" value="sample"/> Browse...
Account	
Bucket	Folder: <input type="text" value="New Folder"/> 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.

< Previous **Apply** Finish Cancel

15 Verify the Summary and click **Finish**.

New Object Storage Repository ×

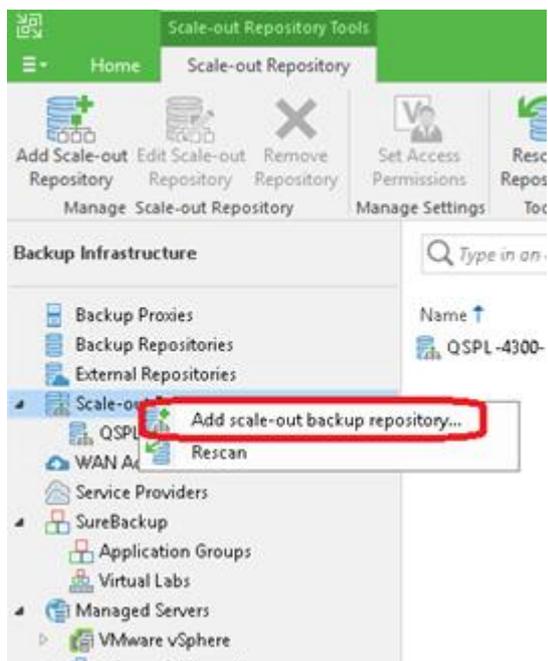
 **Summary**
You can copy the configuration information below for future reference.

Name	Summary:
Account	Object storage repository was successfully created.
Bucket	Name: Object storage repository 2
Summary	Description: Type: S3-compatible Gateway server: not selected Service point: https://hostname:9000 Region: us-east-1 Bucket: sample Concurrent tasks limit: unlimited Storage consumption limit: unlimited Recent backups will not be immutable

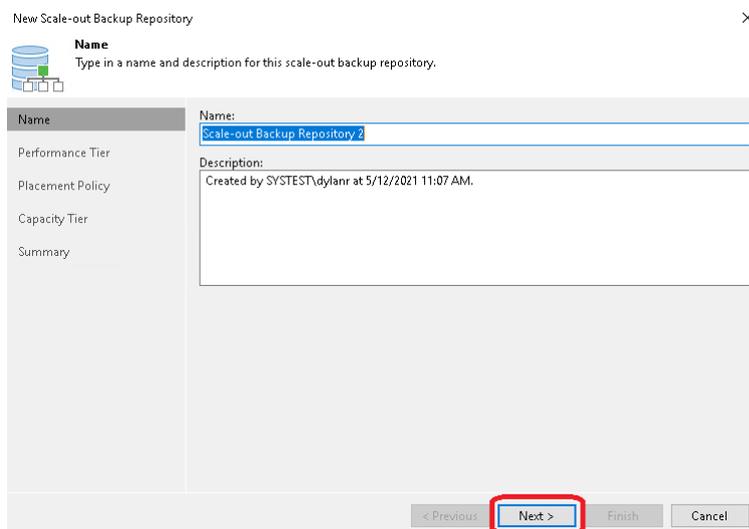
< Previous Next > **Finish** Cancel

Adding the Object Container(S3) as a capacity tier to a Scale-Out repository

- 1 In the Backup Infrastructure section, right-click Scale-out Repositories, and select Add Scale-out backup repository.



- 2 Click **Next**.



- 3 **Add** an existing spindled disk Repository or QoreStor Performance Tier-based Repository to this page. Click **Next**.

in the Veeam section of this guide. Set the retention age for the object repository, keep in mind restores will be quicker from the Performance Tier. Click Apply.

DO NOT USE ENCRYPTION

The screenshot shows the 'Capacity Tier' configuration page in the 'New Scale-out Backup Repository' wizard. The 'Capacity Tier' section is selected in the left-hand navigation pane. The main area contains several options:

- Extend scale-out backup repository capacity with object storage:** This option is checked and highlighted with a red box. Below it, a dropdown menu shows 'Object storage repository 2' and an 'Add...' button.
- Copy backups to object storage as soon as they are created:** This option is unchecked.
- Move backups to object storage as they age out of the operational restore window:** This option is checked. Below it, a spinner control is set to '14' days, also highlighted with a red box.
- Encrypt data uploaded to object storage:** This option is unchecked.

At the bottom of the wizard, the 'Apply' button is highlighted with a red box.

Warning: Do not configure Encryption in Veeam, this will cause QoreStor savings to be extremely low. Instead, configure the Object Container to use encryption in QoreStor.

6 Verify the Summary and click **Finish**.

The screenshot shows the 'Summary' page of the 'New Scale-out Backup Repository' wizard. The 'Summary' section is selected in the left-hand navigation pane. The main area displays a message: 'Scale-out backup repository was created successfully.' At the bottom of the wizard, the 'Finish' button is highlighted with a red box.

Using Instant Recovery with QoreStor

Veeam's Instant VM Recovery immediately restores a virtual machine (VM) into your production environment by running it directly from the backup file.

Instant VM Recovery uses patented vPower® technology to mount a VM image to a production VMware vSphere or Microsoft Hyper-V host directly from a compressed and deduplicated backup file.

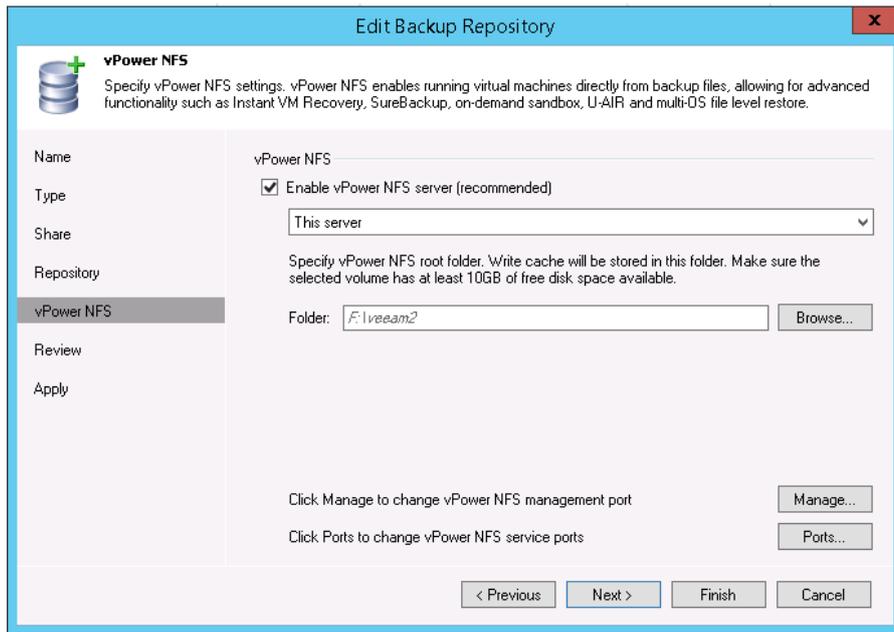
By default, all changes to virtual disks that take place while the VM is running, are logged to auxiliary redo logs residing on the NFS server (Veeam backup server or backup repository). These changes are discarded as soon as a restored VM is removed, or merged with the original VM data when VM recovery is finalized, that is when VM is migrated back to production storage.

Veeam vPower NFS service is a Windows service that runs on a windows backup repository server and enables it to act as an NFS server

Instant Recovery with ESX

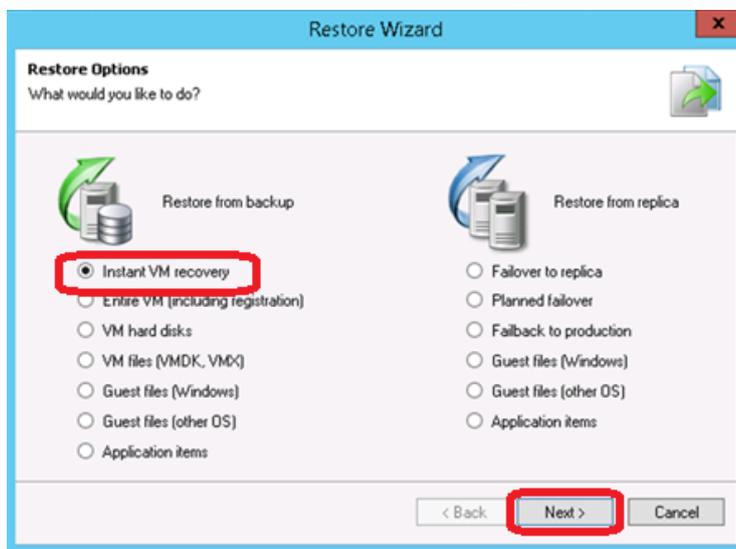
Enabling Instant Recovery with ESX

- 1 Create a backup job for the required VM as described in Section 3, the only difference is to set the **vPower NFS Datastore** in the **vPower NFS** tab.
- 2 Check the checkbox **Enable vPower NFS Server** option on the **vPower NFS** tab and select the appropriate folder as the NFS Datastore.
- 3 NFS Datastore can also be configured on different Windows servers if required and can be done by selecting dropdown and adding the host along with credentials.

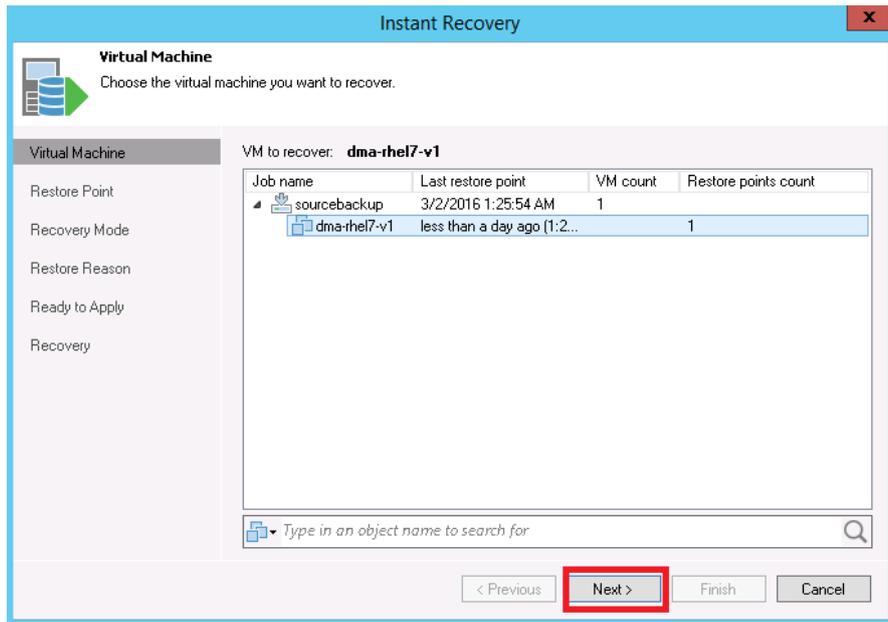


Performing Instant Recovery for ESX

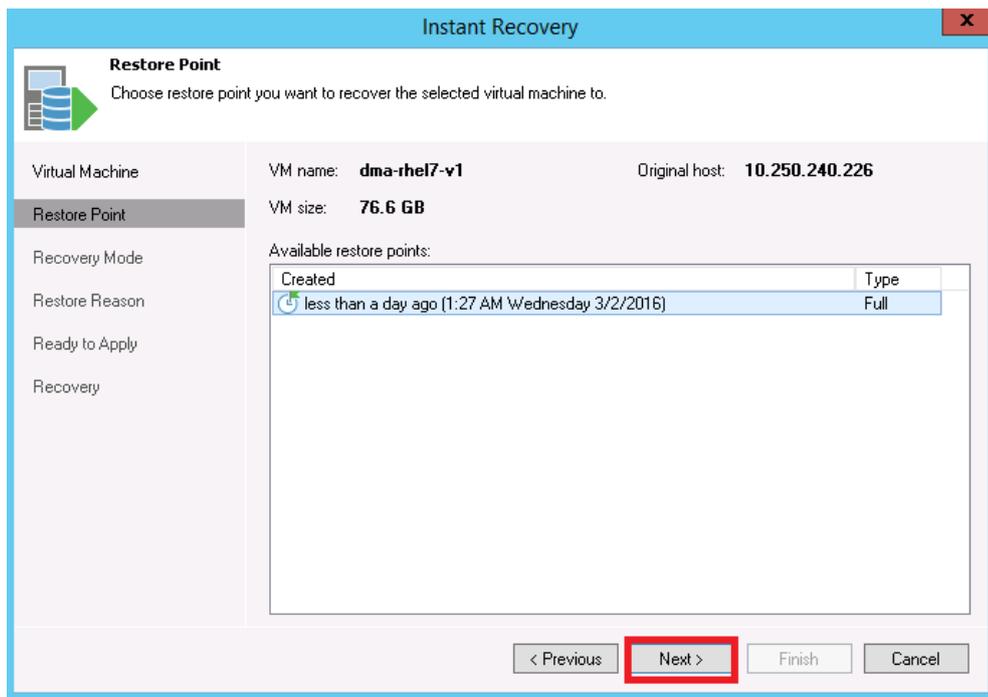
- 1 On Veeam Server's console, click the **Restore Wizard** option, then select the **VMware** option and select **Instant VM recovery**.



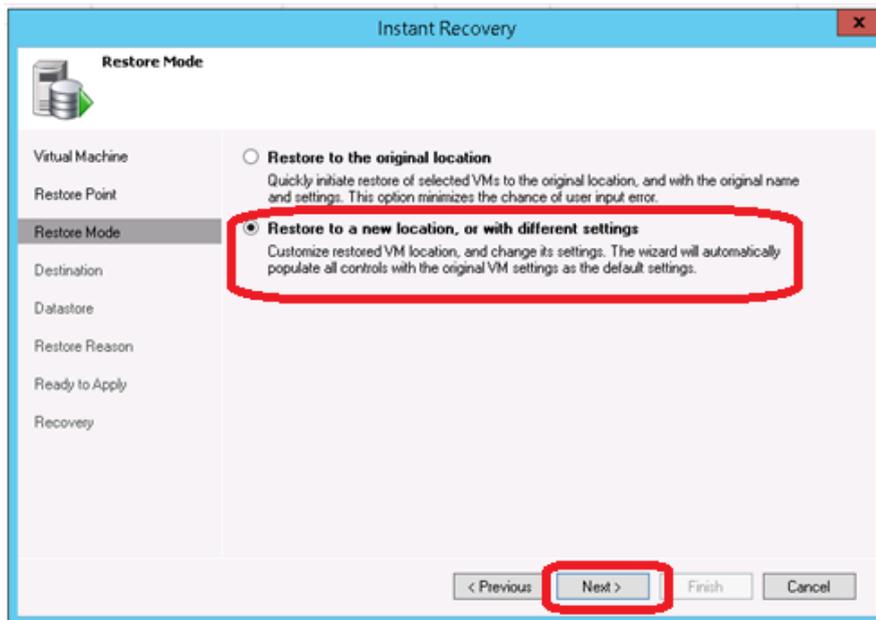
- 2 Select the virtual machine to be recovered and click **Next**.



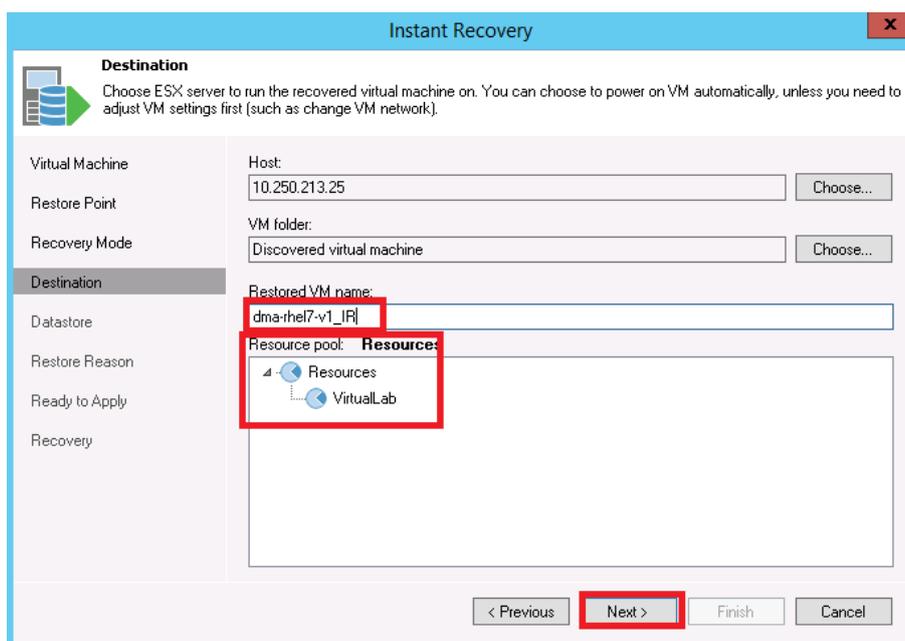
3 At the **Restore Point** step, select the restore point desired.



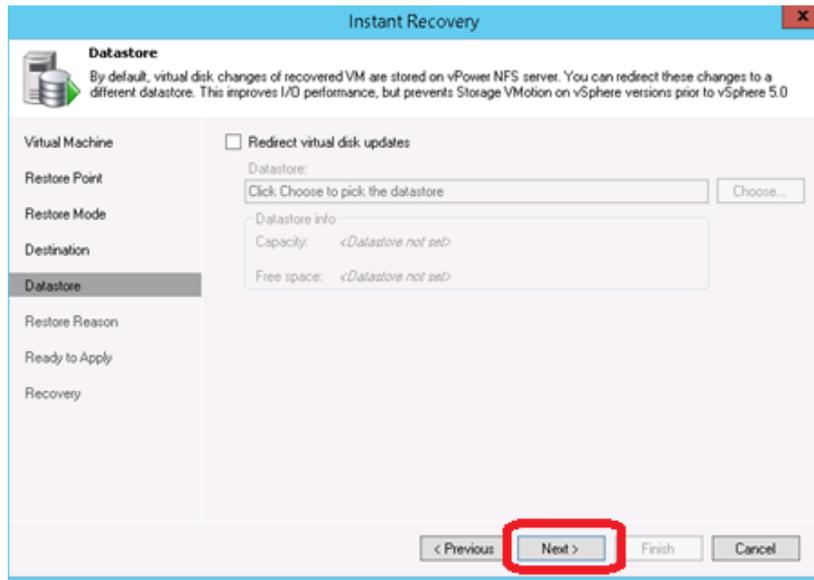
4 At the **Restore Mode** step, select **Restore to a new location, or with different settings**.



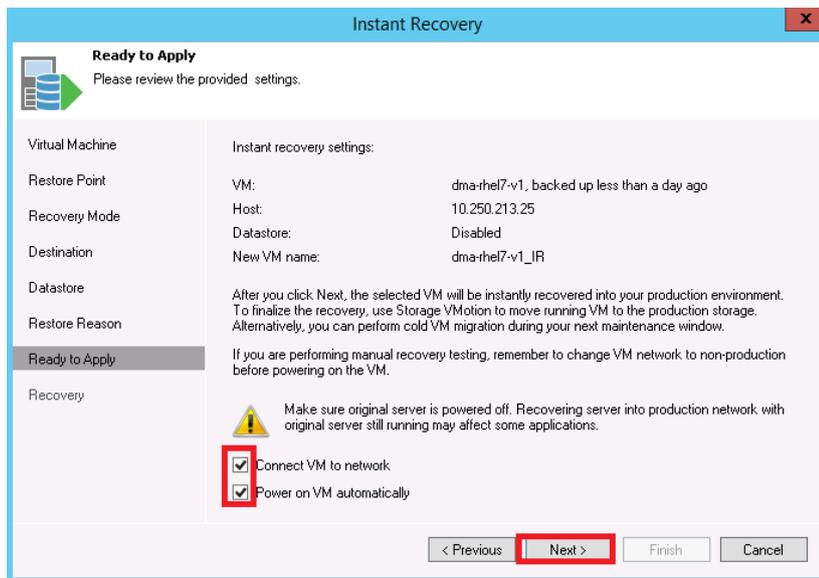
- 5 At the **Destination** step, select the ESX host on which the VM should be restored instantly. In the **Resource pool** box, select the resource pool to which the restored VM should belong.
- 6 In the **Restored VM name** field, set the desired VM name.



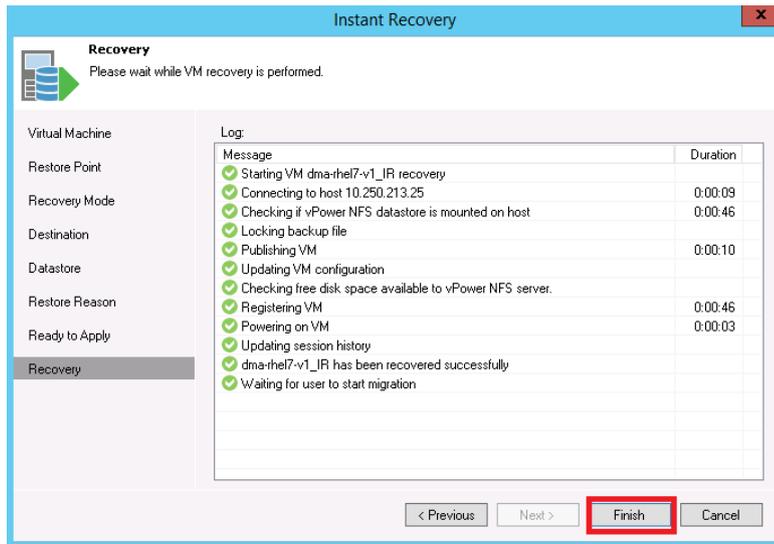
- 7 At the **Datastore** tab, leave the **Redirect virtual disk updates option** unchecked. This will let you use Storage vMotion to migrate the VM to production after the VM is recovered from the backup.



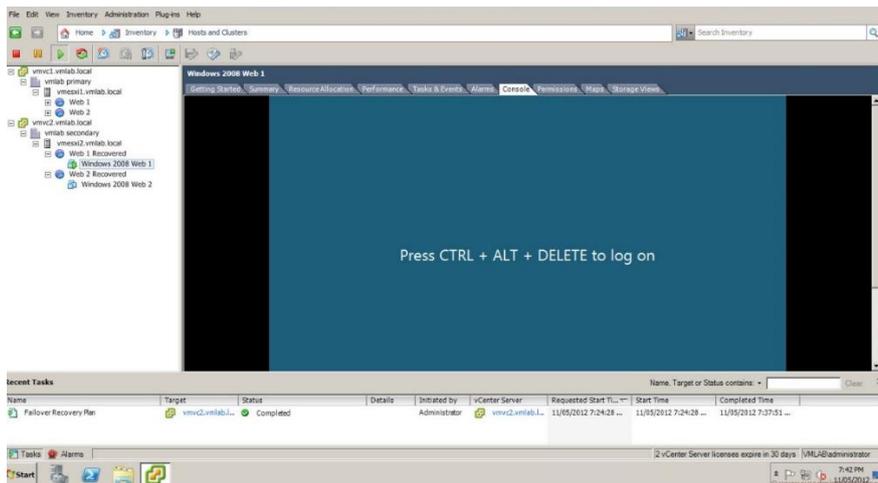
- 8 In the **Ready to Apply** screen, enable **Connect VM to network** and **Power on VM automatically**.



- 9 Click **Finish** to start Instant VM Recovery



10 Open the vSphere client and make sure that the restored VM is started on the ESX host you selected.



11 In Veem Backup & Replication, open the **Backup & Replication** view, select the **Instant Recovery** node in the inventory pane, and make sure that the Instant VM Recovery session is available and mounted.

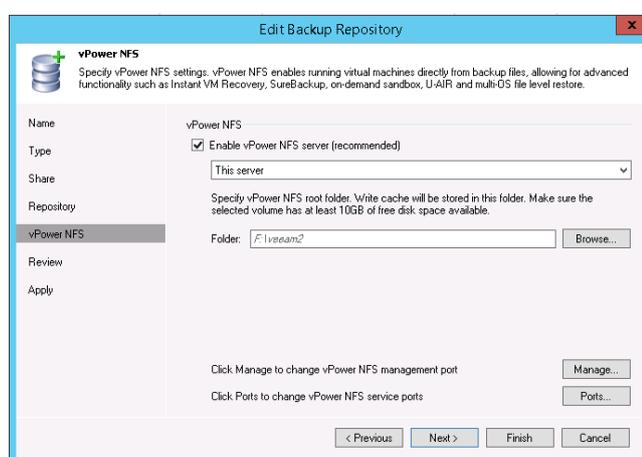


Instant Recovery with Hyper-V Server

Enabling Instant Recovery with Hyper-V

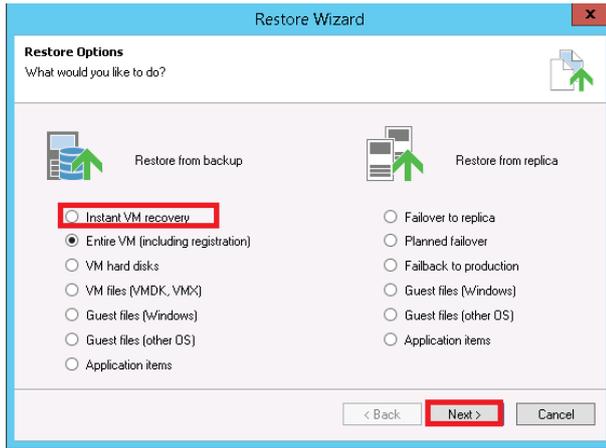
- 1 Create a backup job for the required VM as described in Section 3 and the only difference is to set the **vPower NFS Datastore** in the **vPower NFS** tab as shown in the following screenshot.
- 2 Select **Enable vPower NFS Server** on the **vPower NFS** tab.

i **NOTE:** There is no need to provide a folder as an NFS Datastore. In the case of the Hyper-V, cache data is directly stored at the Hyper-V server datastore location.

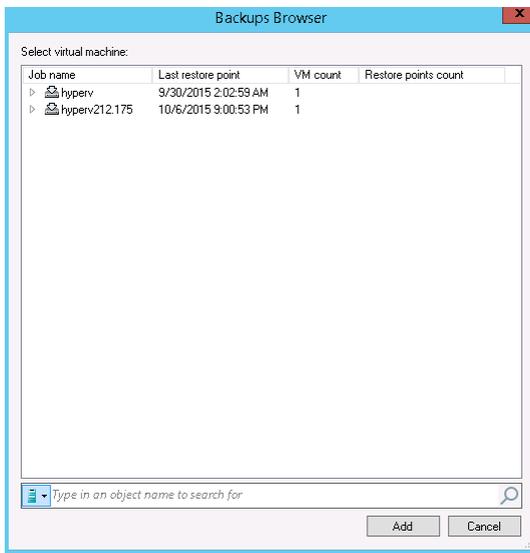


Performing Instant Recovery for Hyper-V

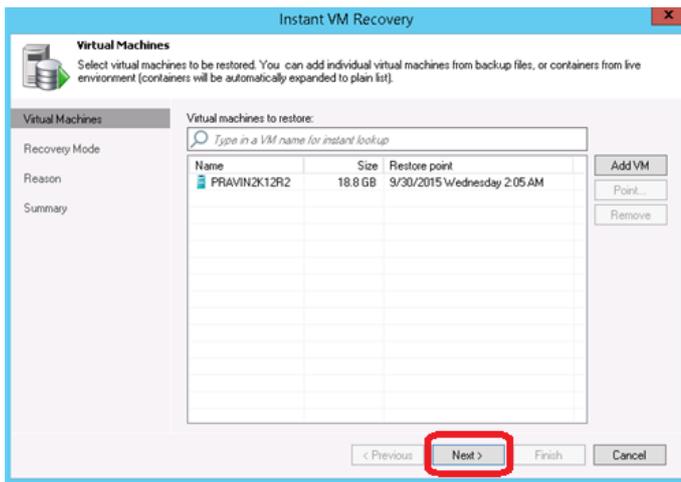
- 1 On the Veeam Backup and Replication console, click the **Restore Wizard**, select **Hyper-V**, and then select the Instant VM recovery.



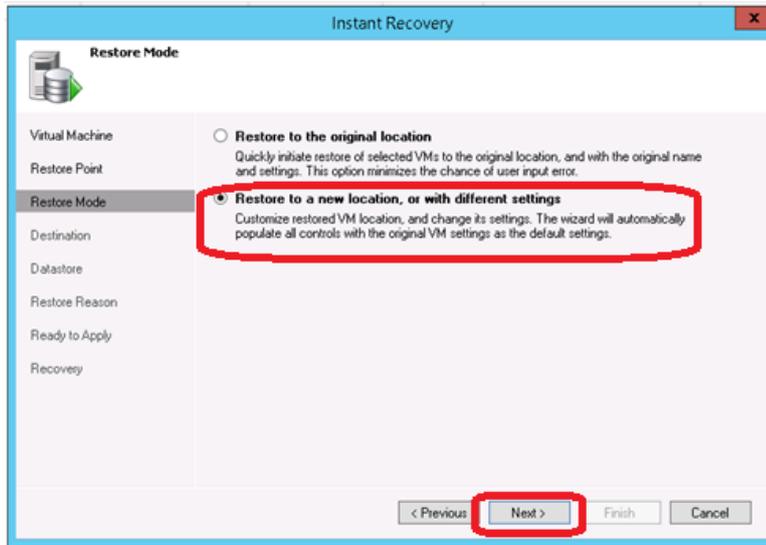
2 Select the virtual machine to be recovered.



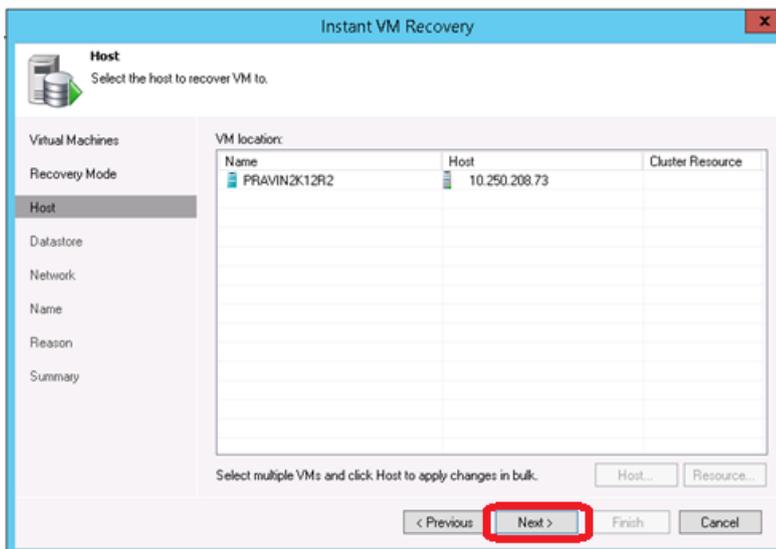
3 Select the desired restore point.



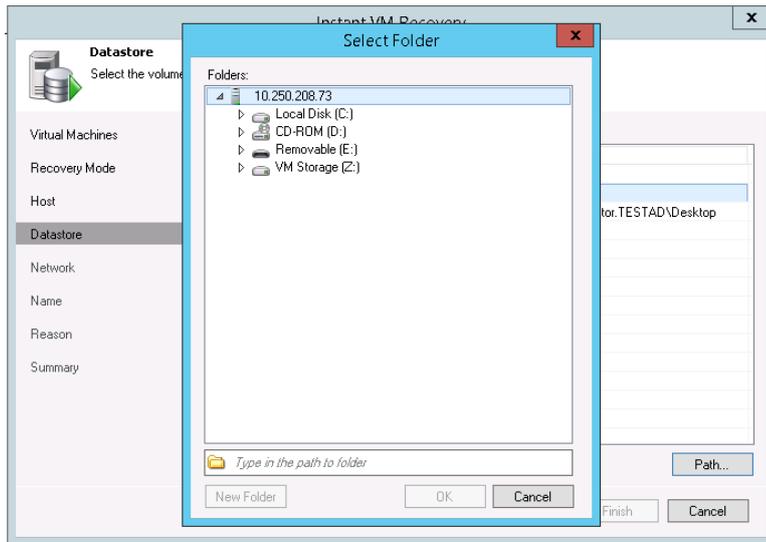
- 4 At the **Restore Mode** step, select **Restore to a new location, or with different settings**.



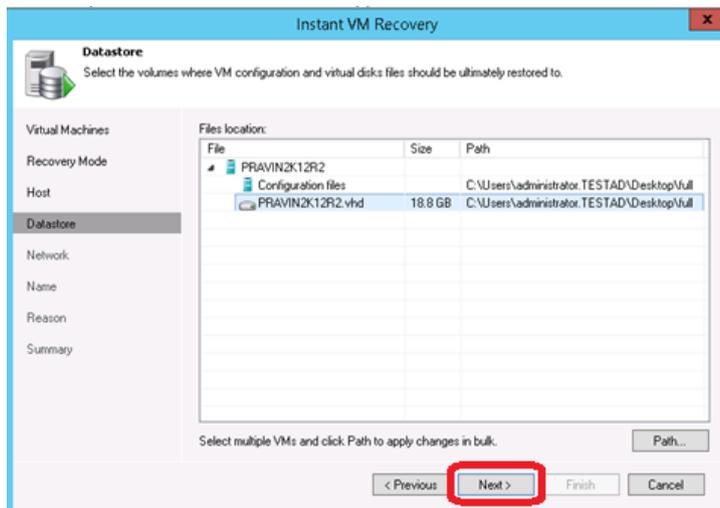
- 5 Select the Host to which your VM should be recovered.



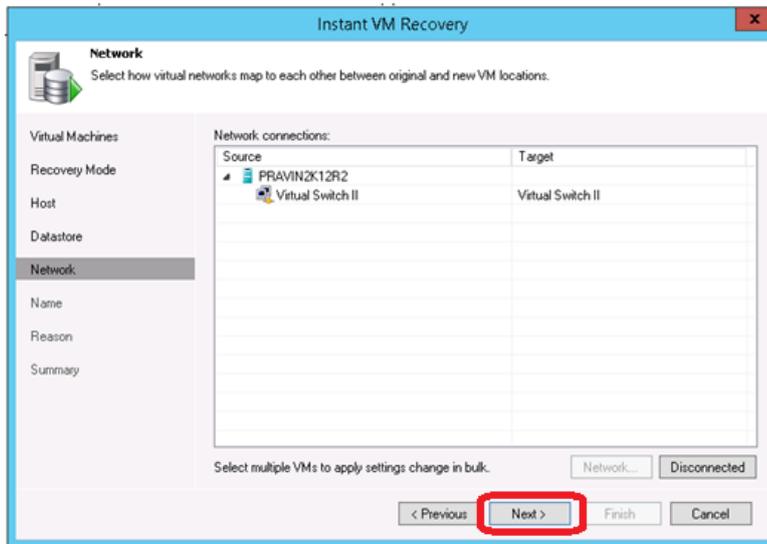
- 6 In the Datastore step, provide the location to temporarily store the cache data.



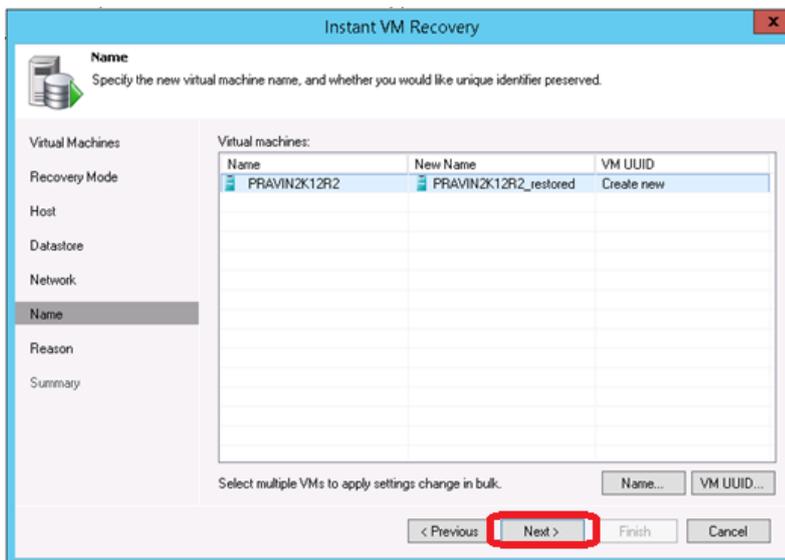
7 After providing the details the screen will look like this:



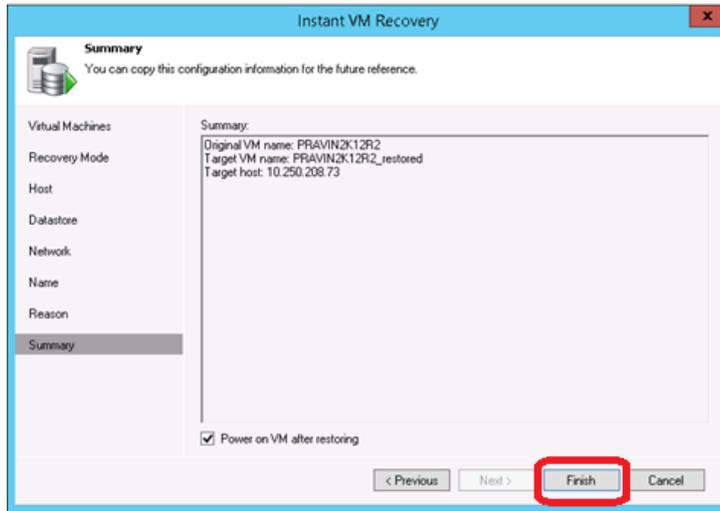
8 In the **Network** section, select the Virtual Networks map to use with the new VM.



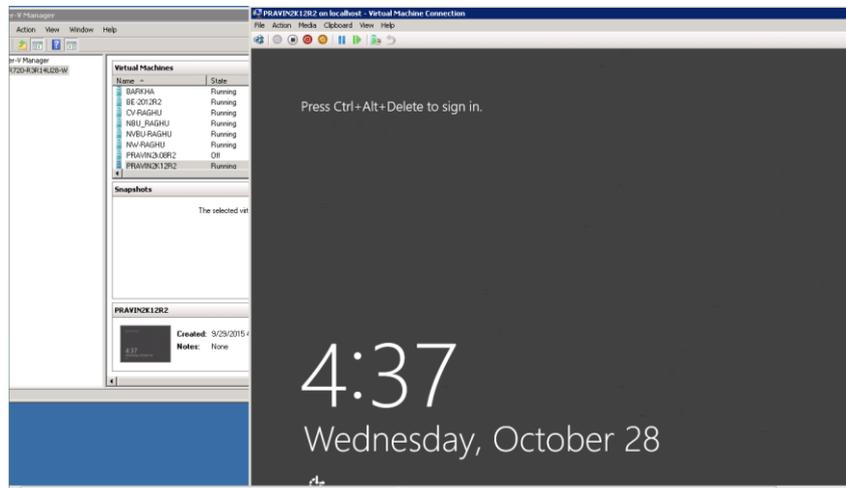
9 In the **Restored VM name** field, set the desired VM name.



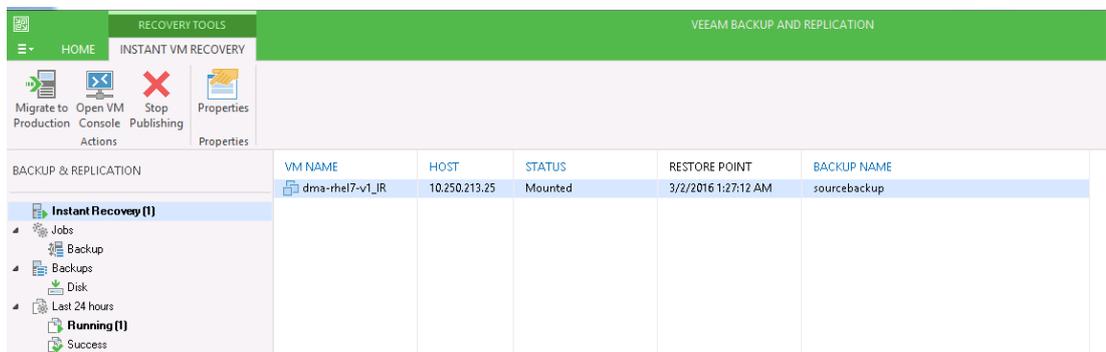
10 Click **Finish** to start the recovery.



11 Open Hyper-v Client and make sure that the restored VM is started on the host you selected.



12 In Veeam Backup & Replication, open the **Backup & Replication** view, select **the Instant Recovery** node in the inventory pane, and make sure that the Instant VM Recovery session is available and mounted.

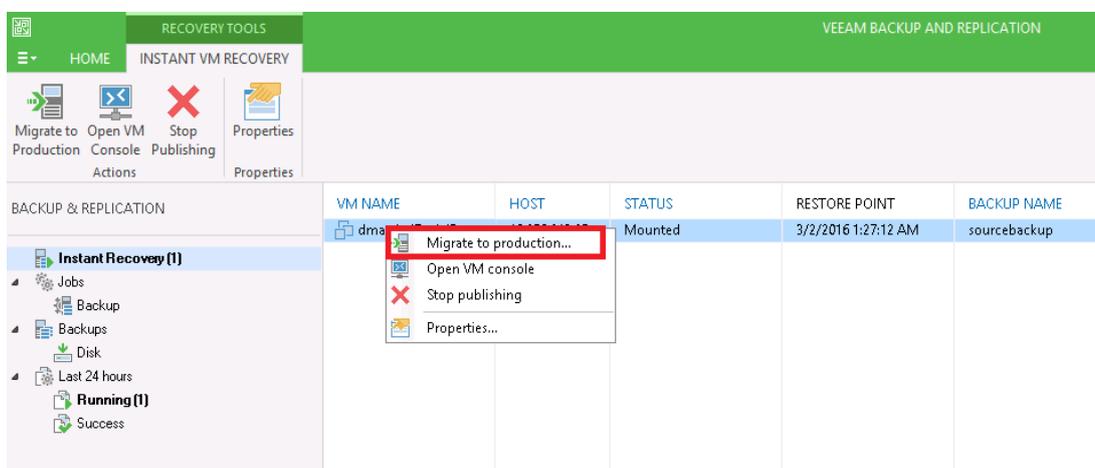


Finalizing Instant Recovery

Migrating VM to a production

For VM migration, you can use VMware Storage vMotion, replicate or copy a VM to production with Veeam Backup & Replication, or use Veeam's Quick Migration. When you migrate the VM to production, the VM data is copied from the backup to production storage. The VM data is pulled from the backup and consolidated with changes made to the VM (redo logs). To migrate the restored VM with Quick Migration:

- 1 Open the **Backup & Replication** view in Veeam Backup & Replication.
- 2 In the inventory pane, select **Instant Recovery**.
- 3 In the working area, right-click the name of the recovered VM and select **Migrate to production**.



Terminating the Instant VM Recovery Session

When you terminate the Instant VM Recovery session, the VM is unpublished from the ESX host, and redo logs are cleared from the vPower NFS datastore. To terminate the current Instant VM recovery session

- 1 Open the **Backup & Replication** view in Veeam Backup & Replication.
- 2 In the **Inventory** pane, select **Instant Recovery**.
- 3 In the working area, right-click the name of the recovered VM and select **Stop publishing**.

QoreStor and Veeam Fast Clone for Hyper-V 2016 backups or Data Copy

Fast clone allows for synthetic full backups of Hyper-V systems or Data Copy jobs with VMs on the ReFS file system with less read performance impact on the QoreStor system. This is achieved through SMB commands and offloading data block copying of existing data to internal operations on the QoreStor instance. It is recommended to configure a new QoreStor repository rather than use a pre-existing one. This is because the existing repository will need to be removed from Veeam to recognize the Fast Clone feature. To do that all Jobs referencing it will need to be moved to other devices or deleted as well. By creating a new container to add as a repository within the same Storage Group, no savings impact will be noticed.

Requirements of Fast Clone

Fast clone is a combination of a Microsoft ReFS filesystem operation, SMB command, Hyper-V backup, and Veeam operation. When considering Fast Clone for QoreStor the following is required:

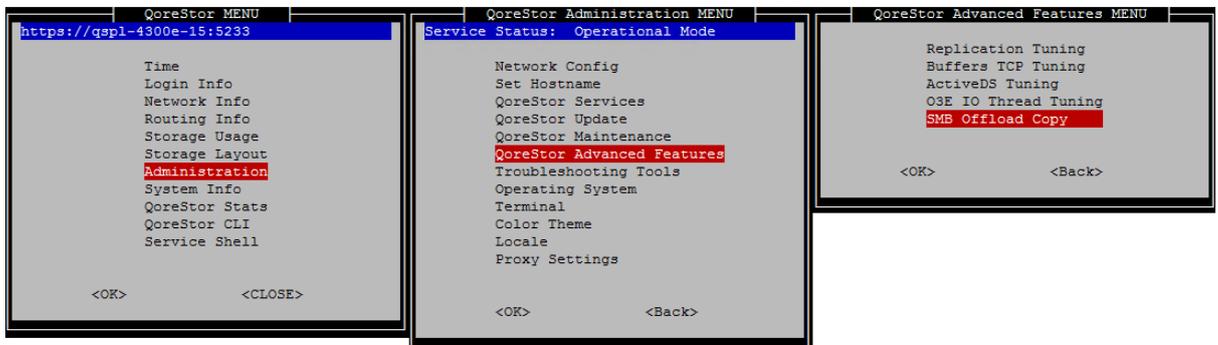
- Veeam 9.5 Update 4 or higher is required.
- The Hyper-V server or Data Copy job proxy [source](#) is running Microsoft Server 2016.
- The VMs for Data Copy job files need to be housed on ReFS File System. NTFS partitions will not work for Fast Clone operations.
- SMB 3.1.1 is required (This is taken care of by the QoreStor version requirements).
- The Veeam backup repository requires the use of the “Align backup file data blocks” option
 - This option will become automatically selected and greyed out making unchecking the option impossible.
- The QoreStor instance is running 6.0 HF2.
- The QoreStor instance has Fast Clone/SMB offload enabled. This setting is off by default.
- The Veeam Proxy moving the data or the Hyper-V server will need to have the Quest Rapid CIFS driver installed and at version 4.0.3220.1 or newer.

- Any Veeam repositories added before enabling Fast Clone/SMB offload will need to be removed and re-added within Veeam to recognize the newly supported option. This is not required if they are not used with Fast Clone jobs.
- Synthetic full operations will need to be configured for all preexisting Veeam backup or Data Copy jobs.

Configuring a new Fast Clone Repository

In this section, we are going to assume the QoreStor being added is new to Veeam. In the following section, we'll cover additional steps to reconfigure existing QoreStor repositories in Veeam.

- 1 First, we'll need to enable Fast Clone support in the QoreStor instance. This is easily done by logging into the qsservice user via ssh. If this is your first time logging in as the qsservice user please reference the QoreStor Deployment Guide for information. Please note enabling the SMB Offload/Fast Clone feature does restart the QoreStor services which could result in failed backup or data copy jobs.
- 2 Once you've logged in via SSH you'll be greeted by the QoreStor Menu. Select **Administration**, followed by **QoreStor Advanced Features**, then the **SMB Offload Copy** option.

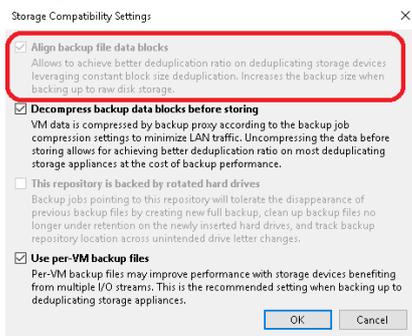


- 3 Select the **Enable SMB Server Offload Copy Support** option, then select **Yes** followed by **Ok**.

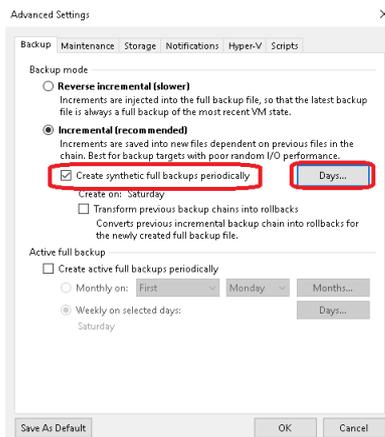


- 4 Wait for the QoreStor services to restart and for the system to become operational again.
- 5 Install the 4.0.3220.1 or newer Quest Rapid CIFS driver on the Veeam Backup and Replication server as well as any Hyper-V server or Veeam proxy that will be used. Please follow the **Installing Rapid CIFS on a Veeam Windows Proxy** section for steps to do this.

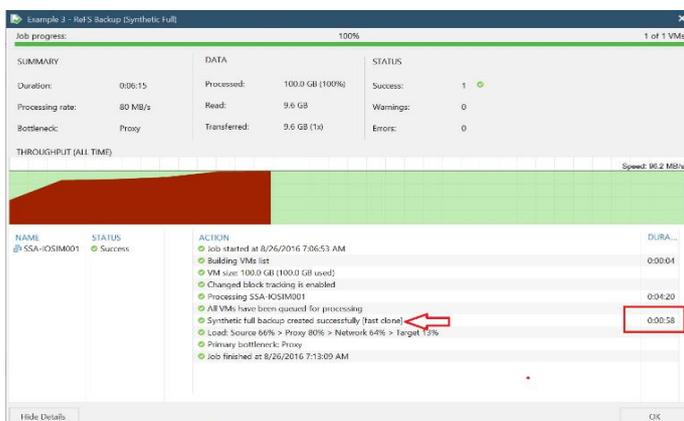
- 6 Create a new CIFS container and add it to Veeam as a repository by following the **Creating a CIFS container for use with Veeam** and **Adding the QoreStor CIFS container as a repository in Veeam** sections of this guide. Ensure the **Align backup file data blocks** option is checked when adding the repository to Veeam. This will likely be automatically checked and greyed out if Fast Clone support is recognized by Veeam.



- 7 Create a new Hyper-V backup or Data Copy job following the **Creating a backup job with the QoreStor system as the target** section of this guide ensuring to use of the Synthetic full option in the job settings.



- 8 For the next Synthetic Full, you should see Fast Clone referenced in the job details.

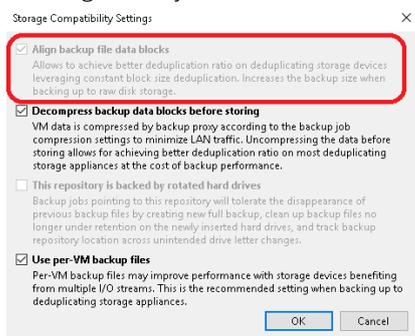


Reconfiguring an Existing QoreStor Repository for Fast Clone

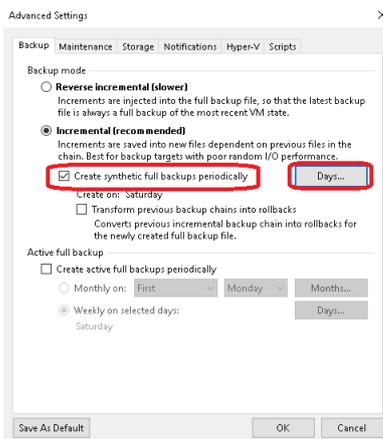
In this section, we'll cover additional steps needed to get an existing QoreStor repository recognized as supporting Fast Clone by Veeam. To achieve this the existing repository will need to be removed and re-added to Veeam. This will involve pointing existing jobs to other repositories or deleting them outright.

Warning: This is an advanced operation and should only be attempted by a customer comfortable with the Veeam product. Quest recommends creating a new Repository in the same Storage Group and leaving your existing repository in place rather than following these steps.

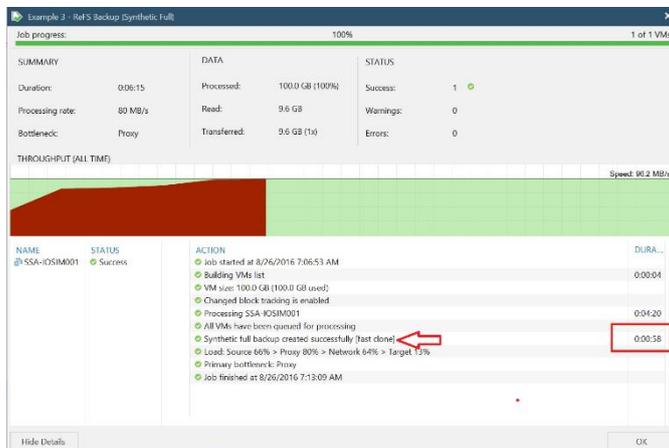
- 1 Follow steps 1 – 5 in the **Configuring a new Fast Clone Repository** section.
- 2 Perform a manual Veeam configuration DB backup and take a copy of that backup file from the repository
- 3 Clone all existing jobs going to the original repository. Do not edit these jobs to configure them with a backup repository yet.
- 4 Remove all existing Veeam Jobs going to the original repository, in 9.5 U4 this should leave the backup files in place and only remove the job and backup file references from the Veeam configuration database.
- 5 Remove the original repository from Veeam, again in 9.5 U4 this should leave the backup files in place and only remove the job and backup file references from the Veeam configuration database.
- 6 Add the original repository back to Veeam, ensuring to select all advanced storage options suggested in the **Adding the QoreStor CIFS container as a repository in Veeam** section of this guide. The **Align backup file data blocks** should be automatically checked and greyed out if Fast Clone support is recognized by Veeam. If not double check all previous steps.



- 7 Run a rescan of the repository once added to Veeam, this may take some time depending on the number of save sets existing in the repository. This will import the existing files into the configuration database and make sure they are restorable.
 - a If the backups are still not restorable run a Veeam configuration backup restore using the backup you manually created. This will put your Veeam server back into the state it was before any jobs were cloned or removed.
- 8 Edit your cloned jobs to use the newly re-added repository. Ensure the Synthetic feature is selected in the job advanced options for every cloned job.



- 9 For the next Synthetic Full, you should see Fast Clone referenced in the job details.



Performance Tier

A Performance Tier allows you to define a set of faster disks as a Storage Group and create a container within that group. This Performance container will always read/write to these faster disks which will allow operations like restores and standard (non-fast clone) synthetic backups to occur quickly. This tier does not stage data off to the standard disks, this is because a restore of synthetic operation reading from the standard disks would still hamper the operation. All data written to the Performance Tier stays within the performance Tier. Because of this, it is recommended to write only specific jobs, which are required to be highly available and are sized to fit within the performance tier size. Please read the QoreStor User Guide for more details about the Performance Tier.

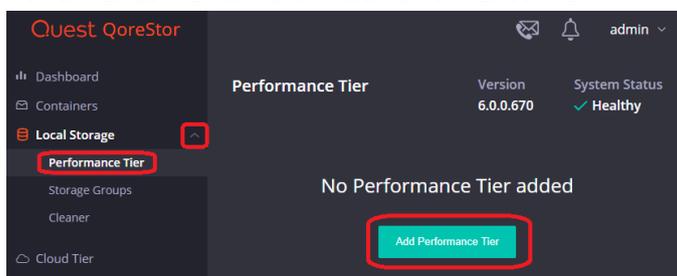


Warning: Please note that once a Performance Tier is added to a system it cannot be easily removed and attempting to do so will most likely result in the destruction of data. Please disable any backup or data copy jobs to the QoreStor system and contact support before attempting removal to find out if this is possible.

Setting up Performance Tier with QoreStor

In this section, we are not going to cover adding a device, creating a partition, creating an XFS filesystem, or defining a mount point in detail. Please reference the QoreStor Installer Guide for this information.

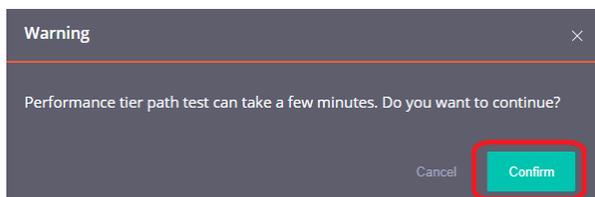
- 1 We first need to cable and add the disks to the OS level. Once seen as a device in the OS an aligned partition will need to be created, an XFS file system created, and a mount point defined in fstab that includes mount option requirements defined in the QoreStor Installer guide.
- 2 Once a file system path to the high-performance storage is added the next step is to add that path as a performance tier in QoreStor. In the QoreStor UI expand **Local Storage** and select the **Performance Tier** tab. Click **Add Performance Tier**.



- 3 Enter the performance tier mount path and click the **Test** button.



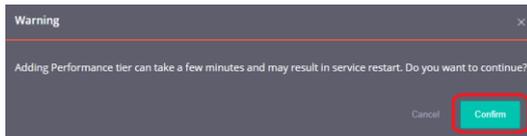
- 4 Click the **Confirm** button.



- 5 If the path gets the expected performance click **Add**.

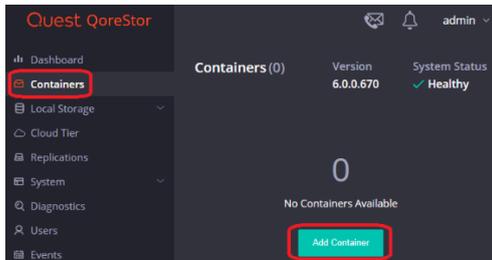


- 6 Click **Confirm** to finish adding the performance Tier, QoreStor services will be restarted

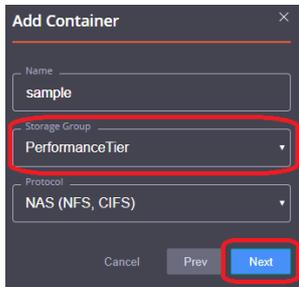


- 7 Once the performance Tier is added you will be logged out. Once logged back in the Performance Tier tab will now list a dashboard for the performance Tier.

- 8 Navigate to the Containers tab and click Add Container.



- 9 In the **Storage Group** dropdown, select **Performance Tier**. Input the container **Name** and set the **Protocol** to **NAS (NFS, CIFS)**. Click **Next**.



- 10 Follow the rest of the steps listed in the **Creating a CIFS container for use with Veeam** and **Adding the QoreStor CIFS container as a repository in Veeam** sections of this guild to finish configuring your Performance Tier container.

Optimizing Performance Tier via Sync Always option

Veeam suggests enabling sync always on CIFS shares. This share-level option decides whether every write to disk should be followed by a disk synchronization before the write call returns control to the client. Setting this to yes can decrease performance but does add some more resiliency to writes in case of interruption of the QoreStor system before writes were synced to disk. We do not recommend this option in cases where performance is a key factor

1. On the QS system run the following command:

```
/opt/qorestor/bin/connection --update --name <container name> --type CIFS --options "sync always"=yes
```

Cloud/Archive Tier

Cloud Tier

Cloud Tier allows per container tiering of deduplicated data to low-cost cloud storage. This enables several potential workflows. Namely the ability to keep longer retention while using less physical space on-site or duplicate archival to the cloud. This is done by establishing a Cloud Tier connection and defining per container policies by which to tier data to the cloud. The policy manager allows for tiering based on time limitations and optionally filtering included and excluded files. It is important to note that individual data blocks will be tiered off not whole backup files. This means if a data block is found frequently over multiple backups it will not necessarily be tiered to the cloud.



Warning: Once a container is configured as Cloud Tier the only way to remove it would be to delete the container or contact Support to fully restore all data blocks from the Cloud. This might involve a read cost from the cloud provider



Warning: It is important to fully consider your Veeam Job configuration and policy configuration when deploying Cloud Tier. Failure to do so could result in unexpected charges from the cloud provider or even failing backup jobs. Please read this section in its entirety as well as check the Cloud Tier section of the QoreStor User Guide.

Important Considerations for Cloud Tier with Veeam

Cloud tiering is achieved by sending deduplicated data blocks to low-cost cloud storage on a cloud provider. These data blocks are identified via a per-container policy manager. The Policy manager options are Idle Time, On-Prem Retention, Include/Exclude Directory paths, and Include/Exclude file types.

- **Idle Time before cloud migration** – Replicates stable data blocks idle for more than the selected number of days/hours to the cloud. After this completes data blocks will be located both On-Premises and on the cloud. All restores will come from the On-Premises data block and not induce any cost. Any attempted modification of files after this idle time will result in access-denied errors. This is why the job type should be considered in Veeam, more on this later in this section.
- **On-Prem Retention Age** – After the selected number of days/hours data blocks that have replicated to the cloud will be removed from On-Premises storage. After this, any data reads, such as restore

or synthetic full backups, will be from the Cloud Provider. This can be slower and induce costs from the provider.

- **Folder Paths** – Allows for including or excluding specific paths from cloud tiering replication. Usually, this feature shouldn't be needed with Veeam.
- **File Extensions** – Allows for including or excluding specific file types from cloud tiering replication. Usually, this feature shouldn't be needed with Veeam.

In most cases, with Veeam, Only Idle time and On-Prem Retention need to be considered.



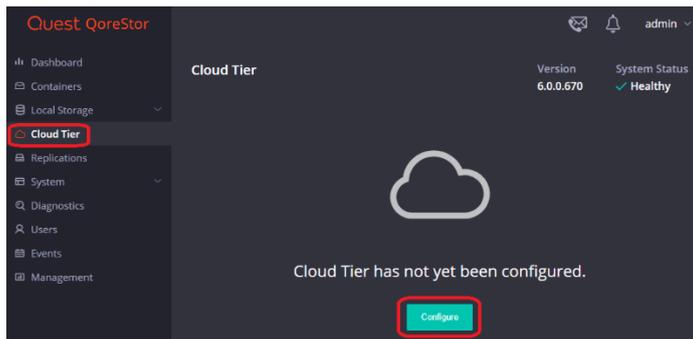
Warning: **Idle time** is especially important to consider with two workflows. Forever Forward Incremental and forward incremental with Synthetic Full Backups.

- **Forever Forward Incremental** – Quest recommends against using Forever Forward Incremental jobs with Cloud Tiering at all. In this workflow, a full backup is taken initially and kept, every backup after this will be incremental. Importantly once retention is met the **Original** Full Backup file has the older incremental injected into it. This means the oldest file in a backup chain is modified by Veeam. If this first full is determined idle by the policy manager "**Idle Time before cloud migration**" setting, any attempts at modifying it will fail with access denied errors. Even if the Full backup is excluded from cloud tiering the oldest incremental will be read from the cloud resulting in a charge from the cloud provider.
- **Forward Incremental with Synthetic Full Backups** – Quest recommends considering your Synthetic Full schedule when using this workflow with Cloud Tiering. In this workflow, you schedule a periodic Synthetic operation in your backup job. This can be daily, weekly, or monthly. In this workflow, the initial backup will be a full backup. The following days will be incremental backups until your next scheduled synthetic full backup. During the synthetic full Veeam will read from the most recent Full as well as every incremental after it. All of this data will be written into a new Full backup file. It's important that your "**On-Prem Retention Age**" setting is longer than your synthetic schedule. If this isn't done the Synthetic operations will result in cloud reads which will result in performance impact and induce cost from the cloud provider.
- **Forward Incremental with Active Full backups** – All new backups will be written into new full or incremental backup files. There is no consideration for this backup time and it will work without issue with Cloud Tiering.
- **Reverse Incremental** – In this workflow, a full backup is taken initially. Each additional backup will be incremental which is then injected directly into the full. After the inject an incremental file is left with all the data removed from the full. These files are okay to tier to the cloud without issue. The injection means the full backup will be modified every backup instead of a new file created. The "**Idle Time before cloud migration**" setting needs to be longer than your scheduled incremental backup frequency. This will likely be easy to achieve since incremental backups typically happen frequently. Failure to do so will result in access denied errors.

Setting up Cloud Tier

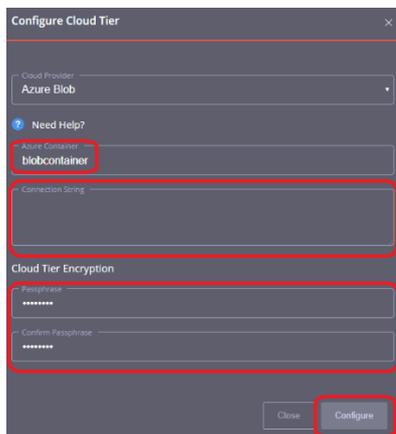
Before setting up Cloud Tier it's important to gather some information from your cloud provider. If using Azure, you will need your Connection String, this can be found on your Azure portal under your blob storage account. If using AWS, Wasabi, or an S3 Compatible cloud provider you will need your Access Key, Secret Key, Region, and Endpoint setting (if using a cloud emulator). These can be found on your AWS console or from your cloud provider.

- 1 In the QoreStor UI select the **Cloud Tier** tab then click the **Configure** button.



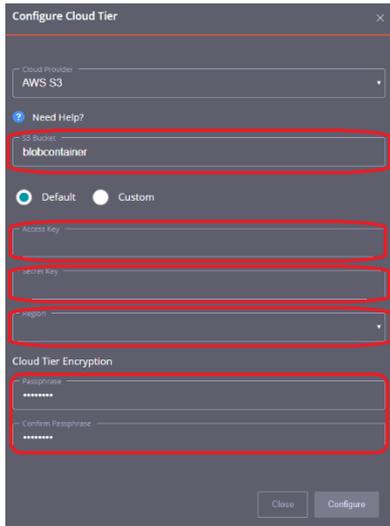
- 2 For Azure enter your Azure Container name, this will be created automatically in the cloud. Enter your Connection string from the Azure portal and your passphrase. This passphrase is user-defined and used to securely encrypt all files written to the cloud provider. Finally, click Configure.

i **NOTE:** Please note the Azure Container name need to be lower case and some symbols are not allowed. This is a limitation of Azure



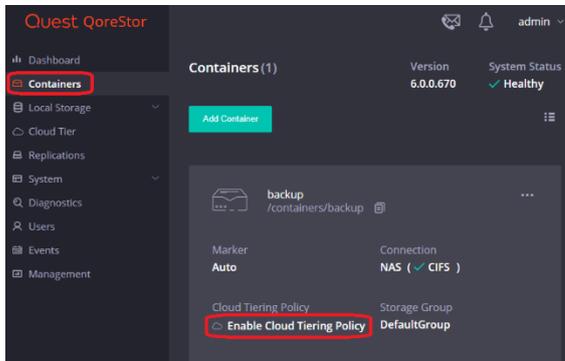
- 3 For AWS, Wasabi, or S3 compatible enter your S3 bucket name, and this will be created. Enter your Access Key, Secret Key, Region, and passphrase used to encrypt all data written to the cloud provider.

i **NOTE:** Please note the S3 Bucket name need to be lower case and some symbols are not allowed. This is a limitation of S3



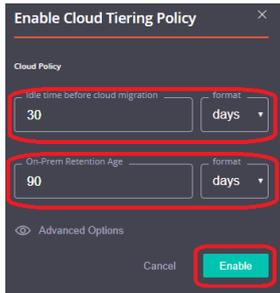
- 4 At this point, Cloud Tier should show as configured and the **Cloud Tier** tab will be populated with statistics. The next step will be to Enable the Cloud Tiering Policy on individual containers.
- 5 Select the **Containers** tab and find or create a container. Click the “Enable Cloud Tiering Policy” hyperlink on this container.

Warning: Once a container is configured as Cloud Tier the only way to remove it would be to delete the container or contact Support to fully restore all data blocks from the Cloud. This might involve a read cost from the cloud provider

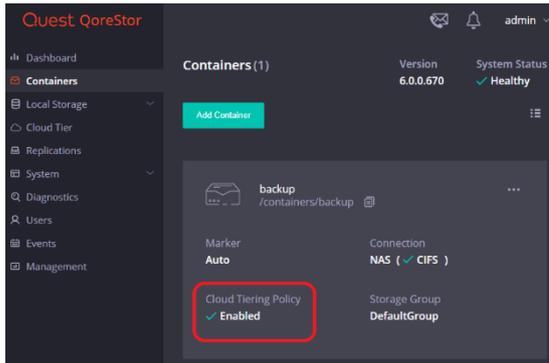


- 6 Define the **Idle tie before cloud migration** and **On-Prem Retention Age**, and click **Enable**.

Warning: Please reference the [Important Considerations for Cloud Tier with Veeam](#) section of this guide before defining idle time and retention age.



- 7 The container will not show as having Cloud Tiering Policy enabled. Idle data will now automatically tier to the cloud provider.



Archive Tier

Important Considerations for Archive Tier with Veeam

QoreStor's archive tier feature enables QoreStor data to be quickly and easily archived to long-term Amazon S3 Glacier or Amazon S3 Glacier Deep Archive storage. Using Veeam and a supported protocol (Object container(S3)), files can be written to a QoreStor container and migrated to your archive tier according to easily defined policies. QoreStor provides a policy engine that allows you to set file age and on-premises retention criteria to be used in identifying which files are most suited for replication to the cloud. Policies are defined at the container level and apply to all files within that container. Using the QoreStor Cloud Policy, you can replicate files based on:

- **Idle time** - replicate stable files idle for more than the selected number of hours.
- **File extensions** - replicate files that match or do not match names in a list of extensions.
- **Regular expressions** - include or exclude files based on their match to configured regular expressions.
- **File locations** - replicated files in a list of directories, or all files except those in a list of directories.

Any data that is archived from the QoreStor instance by the archive tier is encrypted with zero knowledge encryption. The encryption keys are solely owned by you. If the encryption keys are placed in the archive tier, a passphrase is used to encrypt those keys and that passphrase is only known to you. For added security, QoreStor obfuscates metadata names such as blockmap and data store objects that are stored in the archive tier.

Data stored in the archive tier is not available for immediate recovery. When recovery is initiated, the data stays in the archive tier while a copy is made in S3 standard storage and kept for an amount of time specified by the **archive_retention_in_warm** parameter. Although recovery times may vary, the general expectations for recovery times are:

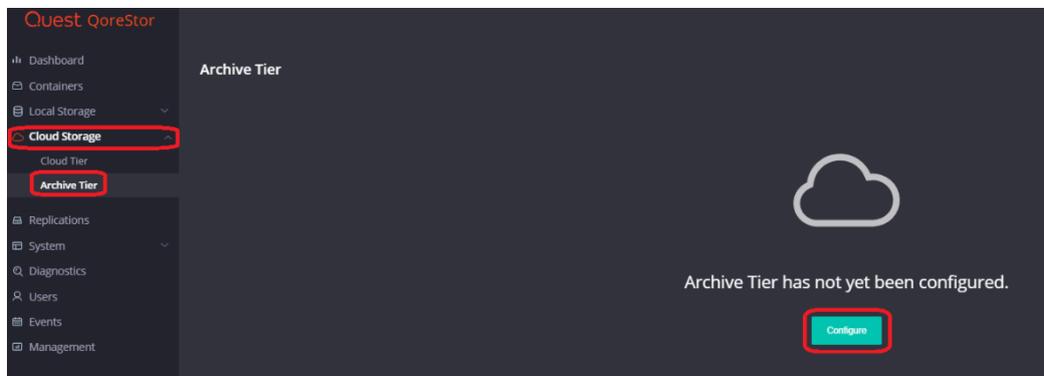
- Amazon S3 Glacier storage: 3-5 hours
- Amazon S3 Glacier Deep Archive: within 12 hours

Setting up Archive Tier

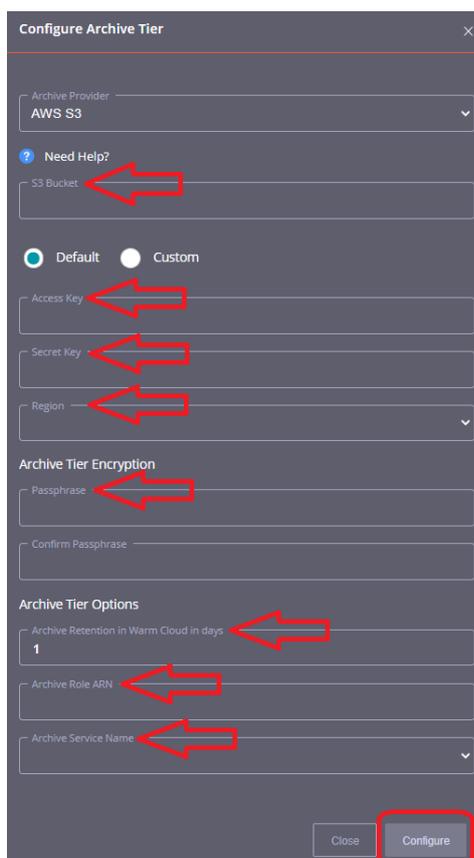
Archive Tier is a feature that allows a QoreStor system to tier deduplicated blocks of files to an AWS glacier/deep archive via S3 protocol. Once added one or more containers can be added to a policy. How that policy is configured can determine how long the data is available on-prem in QoreStor, how long it's available both on-prem and in the archive simultaneously, and finally at what point is it only available in the

cloud. Archive Tier restores are more difficult, careful consideration should be given to how long the data should be available on-prem before configuring the archive tier.

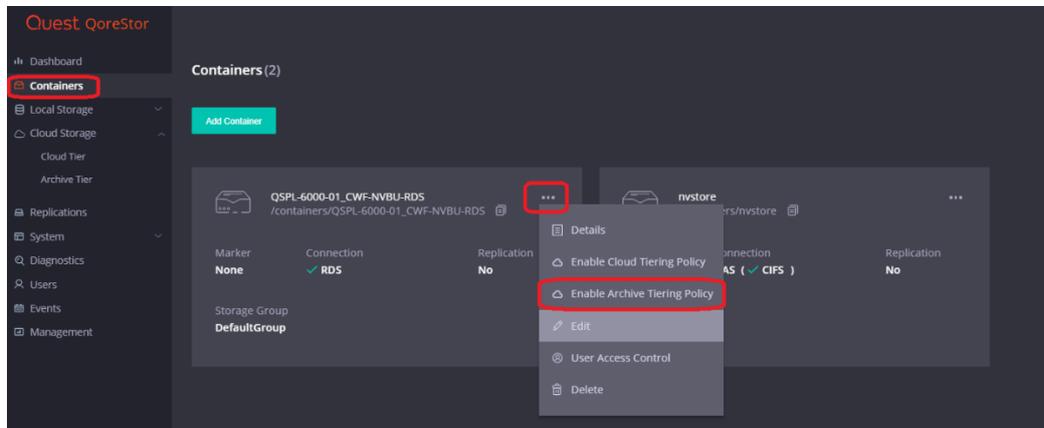
1. Open the QoreStor UI, expand the **Cloud Storage** section, and select the **Archive Tier** page. Click the **Configure** button.



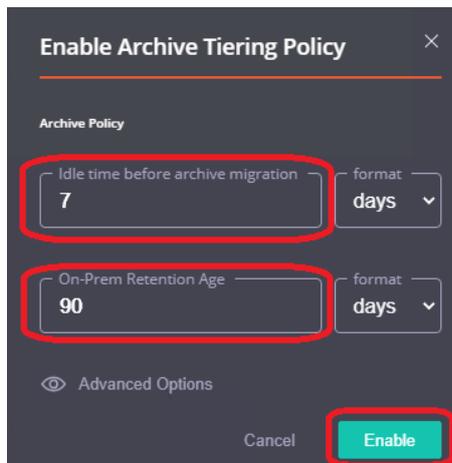
2. You will have to provide several bits of information from your AWS account including the **access key, secret, correct region, ARN role,** and select an **Archive Service Name**. The **S3 bucket name** will be created and is character limited by the provider. Also please make sure to keep your **passphrase**, without this the data is not recoverable in a Disaster Recovery scenario. Finally, click **Configure**.



3. We need to add an Archive tiering policy to a specific container. Do this by navigating to the **Containers** page, selecting the **ellipsis** in the top right corner of the specific container, and clicking **Enabled Cloud Tiering Policy**.



4. In the next window, we need to define the policy. **Idle time before archive migration** specifies the number of hours/days datablocks must be kept idle before being sent to the cloud. The **On-Prem Retention age** specifies the number of hours/days files will be kept locally after they are sent to the archive. Finally, click **Enable**.

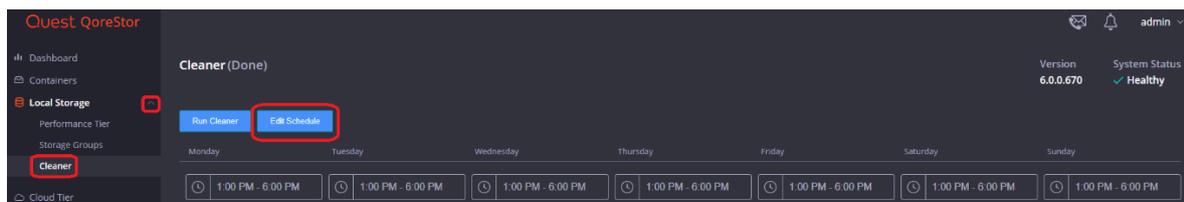


Setting up the QoreStor system cleaner

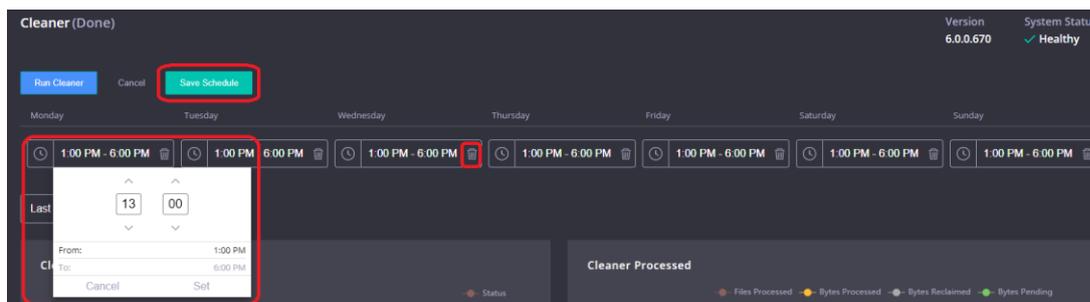
Performing scheduled disk space reclamation operations are recommended as a method for recovering disk space from system containers in which files were deleted as a result of deduplication.

The system cleaner runs during idle time. If your workflow does not have a sufficient amount of idle time on a daily basis, then you should consider scheduling the cleaner to force it to run during a scheduled time. If necessary, you can perform the procedure shown in the following example screenshot to force the cleaner to run. After all of the backup jobs are set up, the QoreStor system cleaner can be scheduled. The QoreStor system cleaner should run at least 40 hours per week when backups are not taking place, and generally after a backup job has been completed. Refer to the *QoreStor Series Cleaner Best Practices* white paper for guidance on setting up the cleaner.

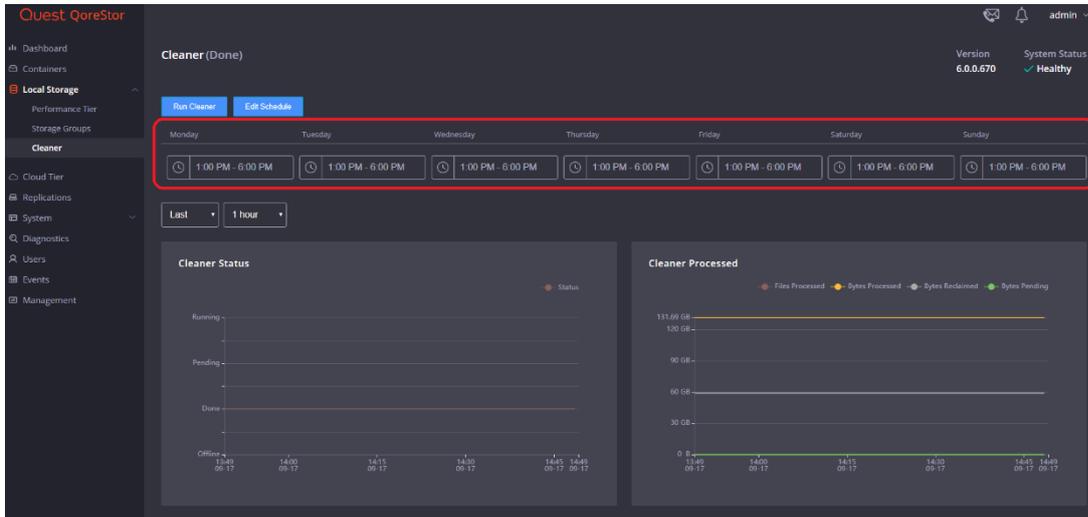
- 1 In the QoreStor system GUI, expand the **Local Storage** tab then click **Cleaner**, and finally **Edit Schedule**.



- 2 Define the schedule and click **Save Schedule**.



3 The new cleaner event is displayed on the **Cleaner** Tab.



Monitoring deduplication, compression, and performance

After backup jobs have run, the QoreStor system tracks capacity, storage savings, and throughput in the QoreStor dashboard. This information is valuable in understanding the benefits of the QoreStor software.

NOTE: Deduplication ratios increase over time. It is not uncommon to see a 2-4x reduction (25-50% total savings) on the initial backup. As additional full backup jobs are completed, the ratios will increase. Backup jobs with a 12-week retention will average a 15x ratio, in most cases.

