



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

Technical White Paper

Quest Engineering
February 2022



© 2022 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 1, 2022

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	9
Creating an NFS container for use with Veeam.....	16
Adding the QoreStor NFS container as a repository in Veeam	19
Configuring Rapid CIFS for Veeam.....	27
Windows prerequisite.....	27
Installing Rapid CIFS on a Veeam Windows Proxy.....	27
Creating a backup job with the QoreStor system as a target.....	31
Setting up QoreStor system replication.....	37
Creating a CIFS/NFS replication session	37
Restoring from the replication target.....	41
QoreStor for Scale-Out Capacity.....	47
Using QoreStor as a Veeam Scale-Out Capacity Tier via Object Container(S3)	47
Creating an Object Container(S3) in QoreStor.....	47
Adding the QoreStor Object Container(S3) as a repository in Veeam	50
Adding the Object Container(S3) as a capacity tier to a Scale-Out repository	58
Instant Recovery	63
Using Instant Recovery with QoreStor.....	63
Instant Recovery with ESX.....	63
Enabling Instant Recovery with ESX.....	64
Performing Instant Recovery for ESX	64
Instant Recovery with Hyper-V Server	69
Enabling Instant Recovery with Hyper-V.....	70
Performing Instant Recovery for Hyper-V.....	70
Finalizing Instant Recovery.....	77
Migrating the restored VM to production.....	77
Terminating the Instant VM Recovery session	78

QoreStor and Veeam Fast Clone for Hyper-V 2016 backups or Data Copy.....	79
Understanding Fast Clone	79
Requirements of Fast Clone.....	79
Configuring a new Fast Clone repository	80
Reconfiguring an existing QoreStor repository for Fast Clone	82
Performance Tier	85
Understanding Performance Tier	85
Setting up Performance Tier with QoreStor	85
Optimizing Performance Tier via Sync Always option.....	88
Cloud/Archive Tier.....	89
Cloud Tier	89
Important Considerations for Cloud Tier with Veeam.....	89
Setting up Cloud Tier.....	91
Archive Tier	95
Important Considerations for Archive Tier with Veeam.....	95
Setting up Archive Tier.....	96
Setting up the QoreStor system cleaner	99
Monitoring deduplication, compression and performance	101

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/>



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

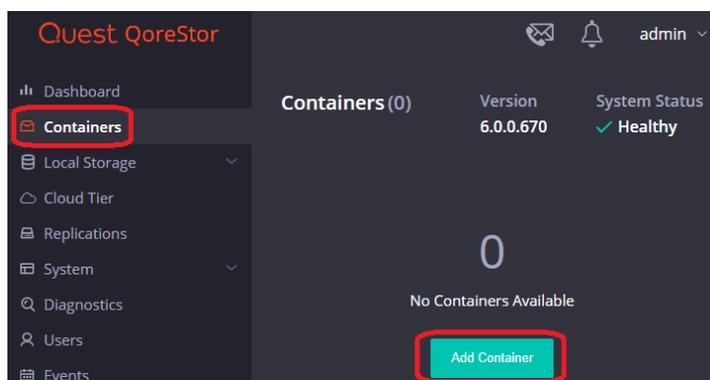
- Creating a CIFS container for use with Veeam
- Creating an NFS container for use with Veeam
- Configuring Rapid CIFS for Veeam
- Creating a backup job with the QoreStor system as a target
- Setting up QoreStor system replication

Creating a CIFS container for use with Veeam

To create a CIFS container for use with Veeam

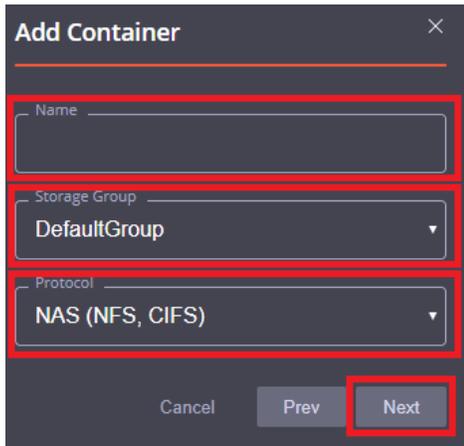
- 1 In the left navigation, click **Containers**, and then click **Add container**.

Figure 1: QoreStor Containers page



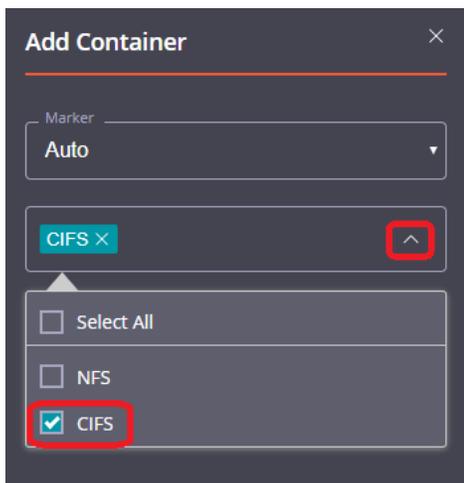
- 2 In the Add Container dialog, enter a container **Name**, select a **Storage Group** or leave the **DefaultGroup**, and then from the **Protocol** drop-down menu select **NAS (NFS, CIFS)**. Click **Next**.

Figure 2: Add Container window



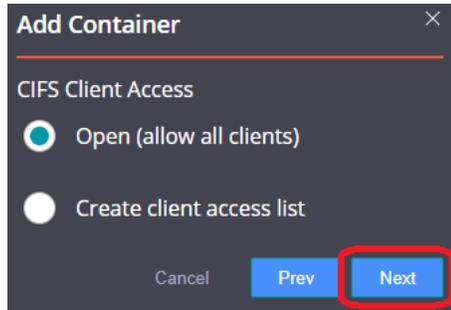
- 3 In the **Protocols** drop-down menu, select **CIFS**. Leave **Marker Type** on **Auto**, then click **Next**.

Figure 3: Add Container window - Protocols



- 4 If needed, fill in the **CIFS Client Access** options, and then click **Next**.

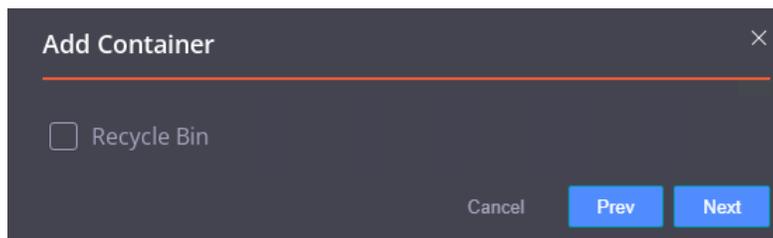
Figure 4: Add Container window – Client Access



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

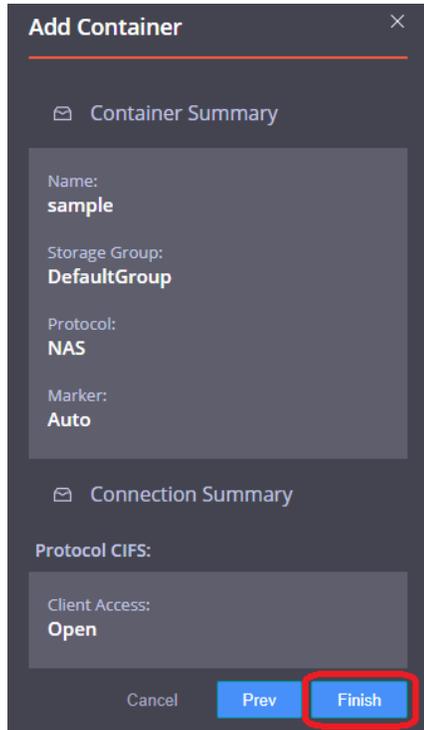
- 5 Optionally, enable the Recycle Bin feature, and then click **Next**.
For more information, see the *QoreStor User Guide*.

Figure 5: Add Container window – Recycle Bin



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

Figure 6: Add Container window - Summary



Adding the QoreStor CIFS container as a repository in Veeam

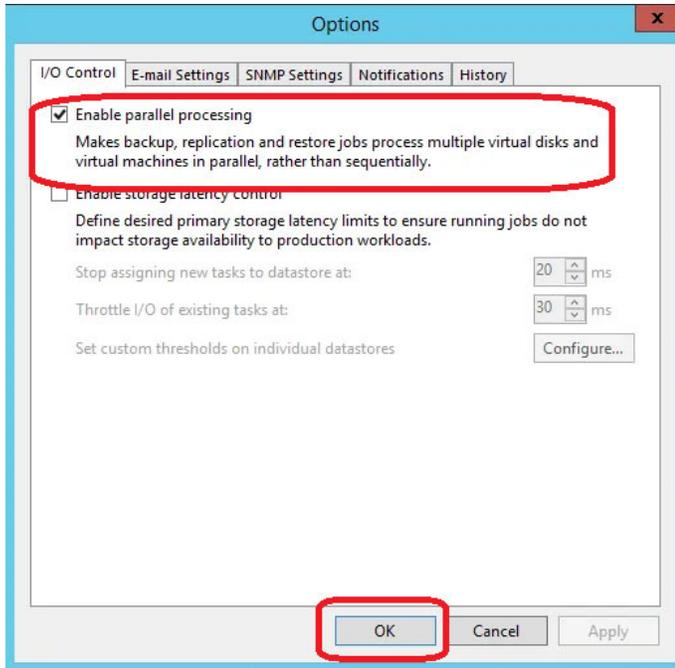
! **CAUTION:** To maximize the QoreStor and Veeam deduplication savings and performance, Quest recommends using the settings provided in this guide for all the data being backed up.

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 same settings.

To add the QoreStor CIFS container as a repository in Veeam

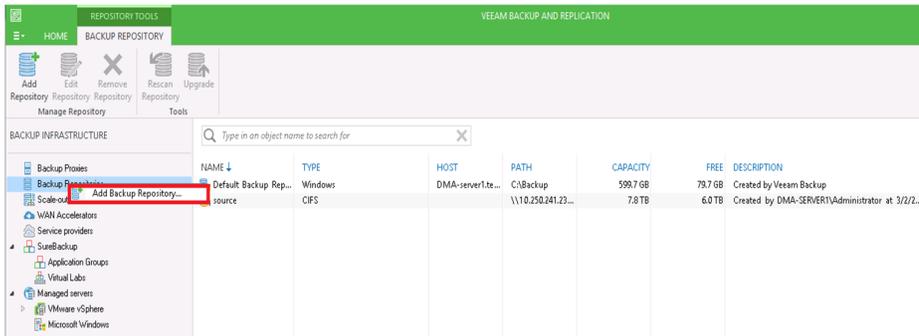
- 1 Open the Veeam Backup & Replication console.
- 2 If using Veeam 9.5 U3 or earlier, select the dropdown **Menu** and click **General Options**
- 3 On the I/O Control tab, select **Enable parallel processing** and click **OK**. This option is missing in Veeam 9.5 U4 and later as it is automatically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

Figure 7: Veeam I/O control options



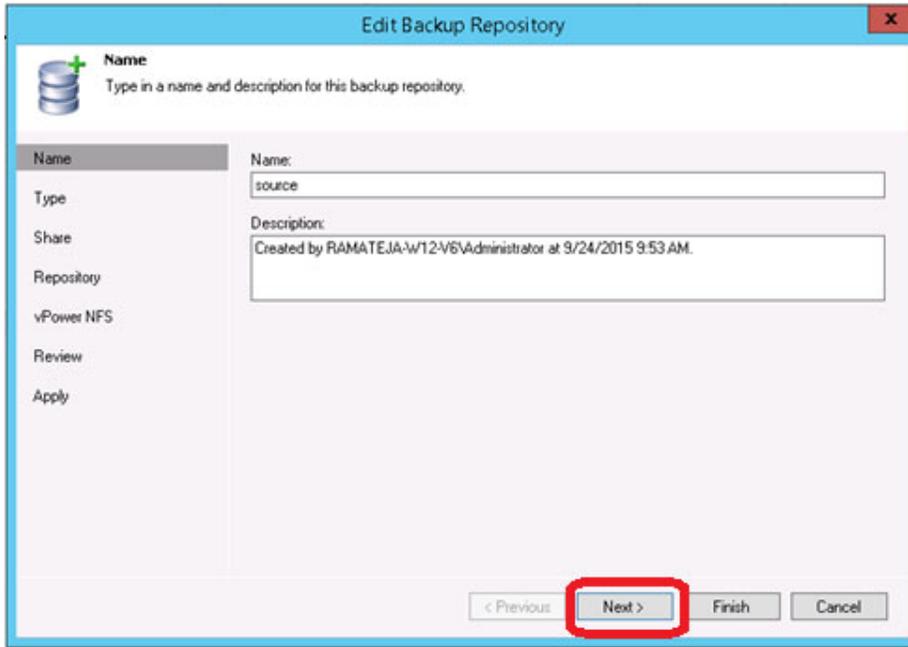
- 4 In the **Backup Infrastructure** section, right-click **Backup Repositories**, and select **Add Backup Repository**.

Figure 8: Veeam Backup Repository page



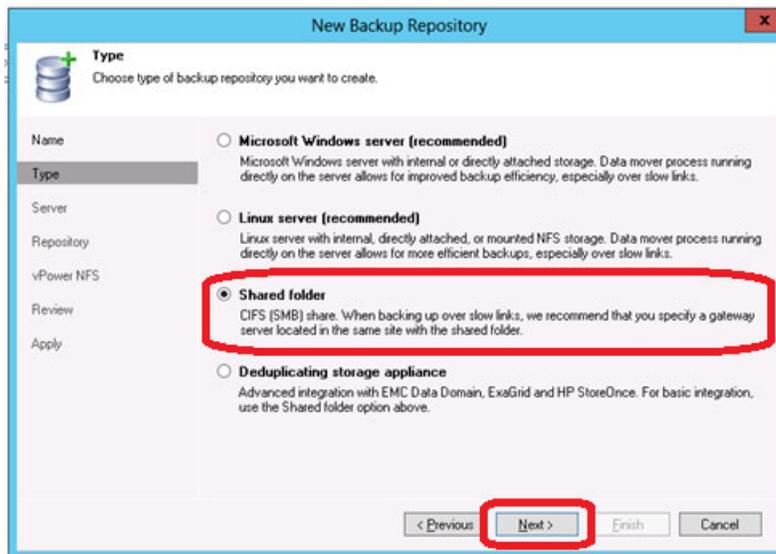
- 5 Enter a name for the QoreStor container repository and click **Next**.

Figure 9: Veeam Edit Backup Repository window - Name



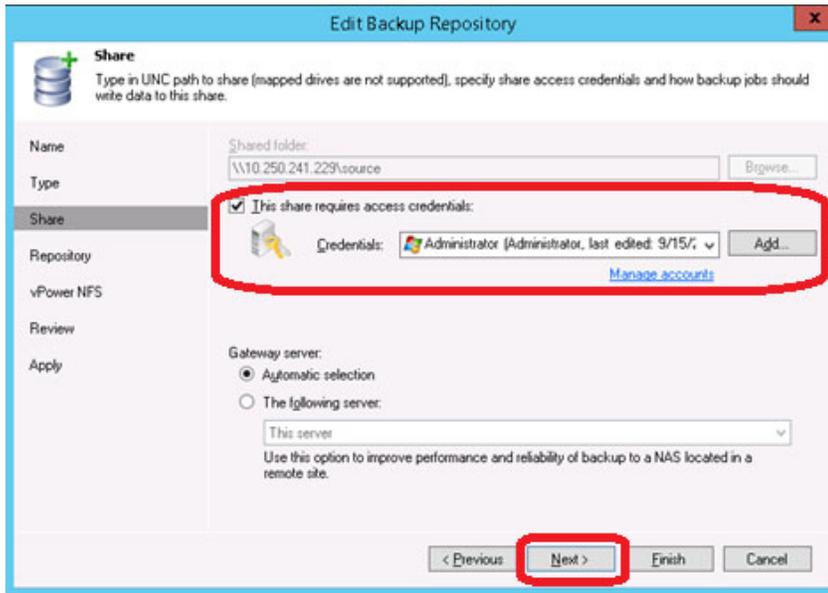
- 6 Select **Shared folder** as the type of backup repository and click **Next**.

Figure 10: Veeam New Backup Repository window - Type



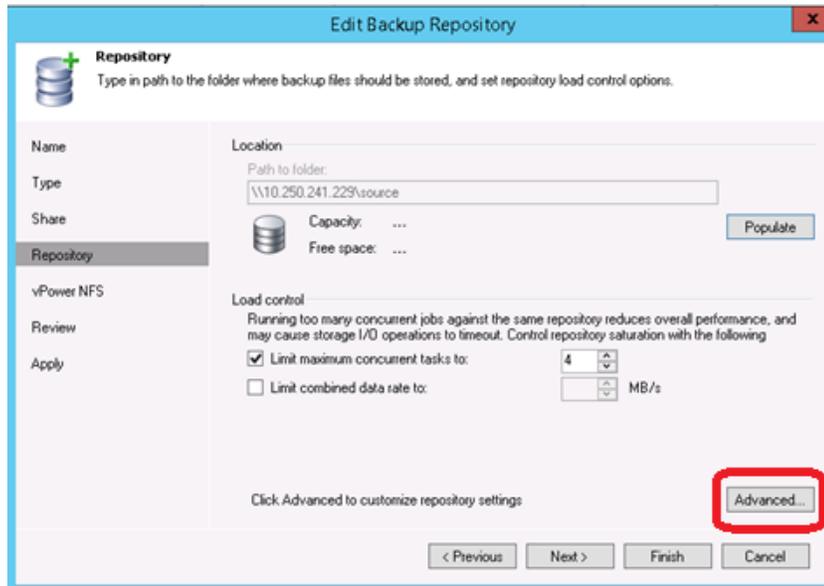
- 7 In the **Shared folder** field, enter the QoreStor container share UNC path (or TCP/IP address to replace hostname), select the **Gateway Server**, and click **Next**.

Figure 11: Veeam Edit Backup Repository window - Share



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

Figure 12: Veeam Edit Backup Repository window - Repository





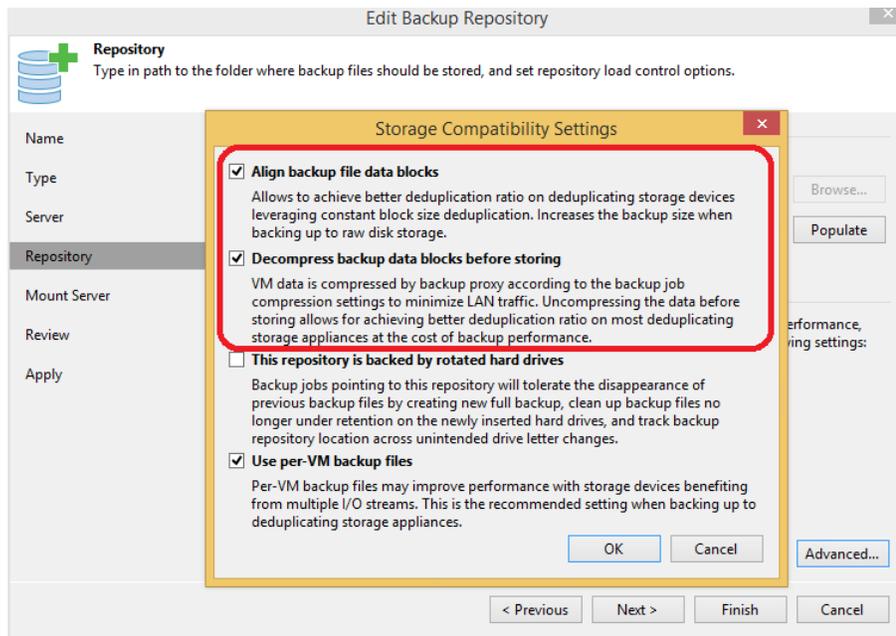
NOTE: For the maximum concurrent jobs supported for CIFS/NFS, see the *QoreStor Interoperability Guide*. The maximum concurrent tasks also depend upon the number of CPU cores of Veeam Servers or proxies.

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



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. Quest does not recommend changing these settings after data has been written to QoreStor.

Figure 13: Veeam storage compatibility settings for data blocks

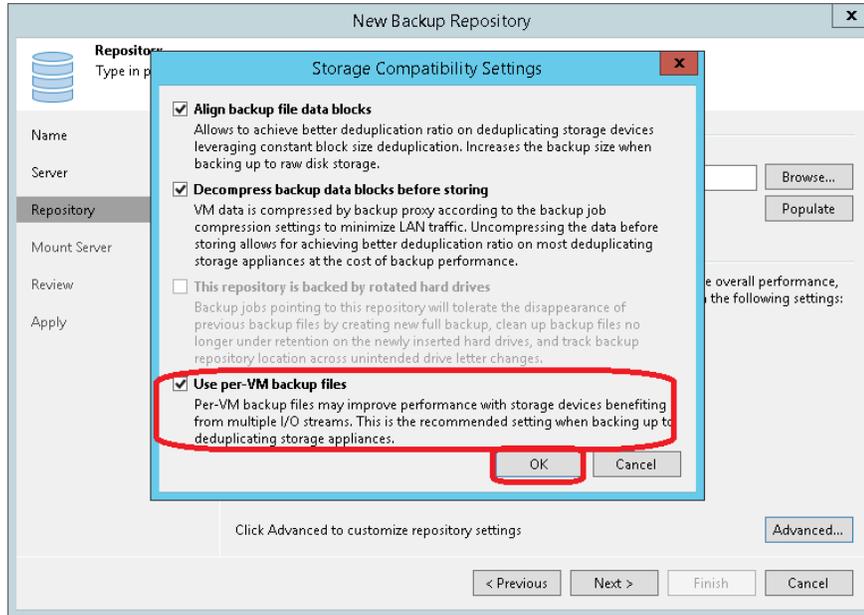


CAUTION: If you change the setting for the **Align backup file data blocks** option after backups are taken, it will impact the deduplication savings for future backups.

- 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.

Figure 14: Veeam storage compatibility settings for VM backup files



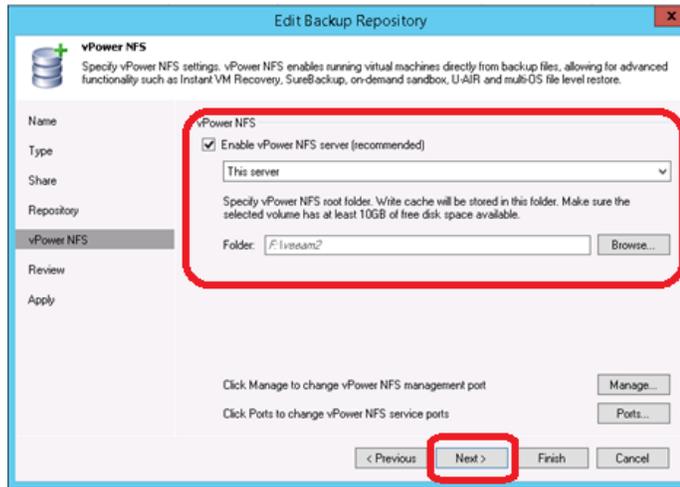
i **NOTE:** If using Veeam 9.5 U3 or earlier, confirmed that you selected **Enable parallel data processing** in Step 3.

i **NOTE:** The Use Per-VM backup files option allows multiple write streams within a single job with parallel processing enabled. Quest recommends selecting this option as it dramatically improves overall job backup performance.

11 Click **Next**.

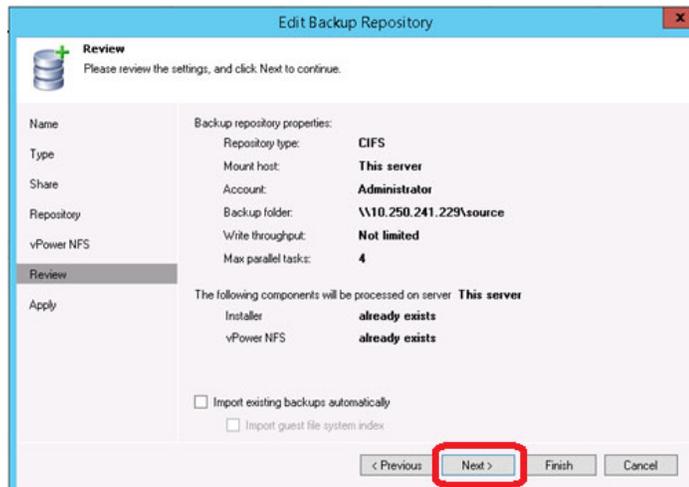
12 To use the Instant Recovery feature, select **Enable vPower NFS server (recommended)**.

Figure 15: Veeam Edit Backup Repository window – vPower NFS



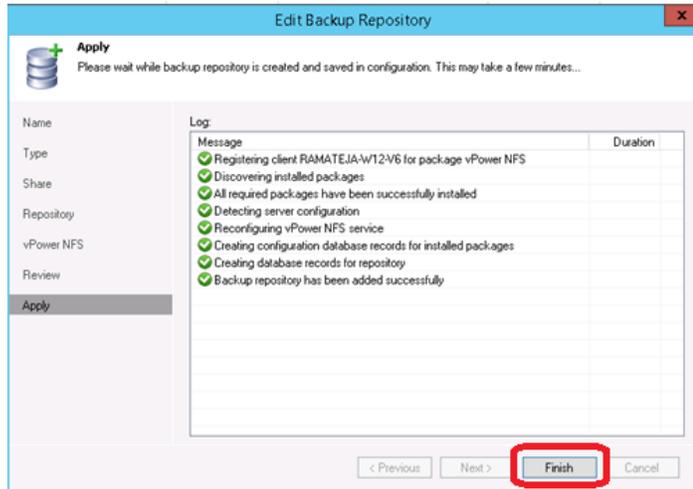
13 On the review page, verify the settings, and click **Next** to apply changes.

Figure 16: Veeam Edit Backup Repository window - Review



14 Click **Finish**.

Figure 17: Veeam Edit Backup Repository window - Apply

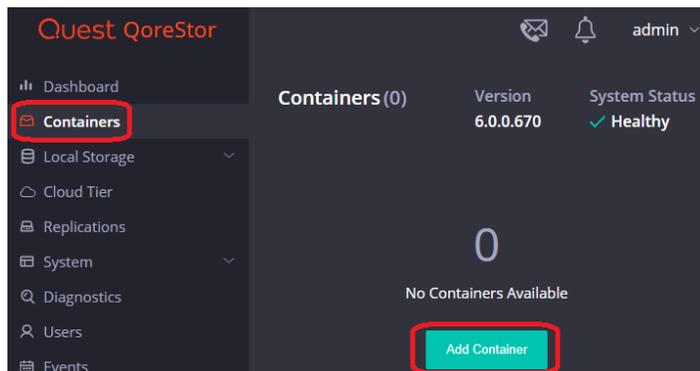


Creating an NFS container for use with Veeam

To create an NFS container for use with Veeam

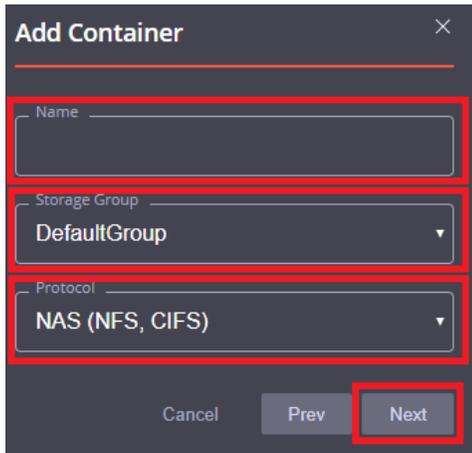
- 1 In the left navigation, click **Containers**, and then click **Add container**.

Figure 18: QoreStor Containers page



- 2 In the Add Container dialog, enter a container **Name**, select a **Storage Group** or leave the **DefaultGroup**, and then from the **Protocol** drop down menu select **NAS (NFS, CIFS)**. Click **Next**.

Figure 19: Add Container window

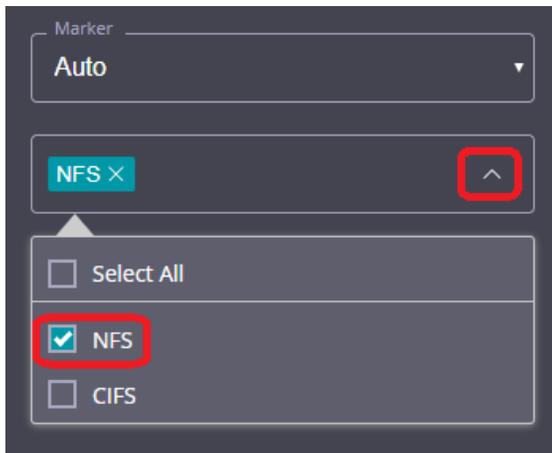


The screenshot shows the 'Add Container' dialog box with the following fields and options:

- Name:** An empty text input field.
- Storage Group:** A dropdown menu currently showing 'DefaultGroup'.
- Protocol:** A dropdown menu currently showing 'NAS (NFS, CIFS)'.
- Buttons:** 'Cancel', 'Prev', and 'Next'. The 'Next' button is highlighted with a red box.

- 3 In the **Protocols** drop-down menu, select **NFS**. Leave **Marker Type** on **Auto**, then click **Next**.

Figure 20: Add Container window - Protocols

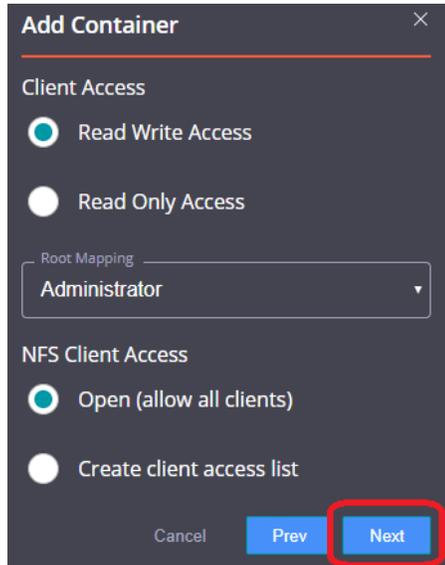


The screenshot shows the 'Protocols' section of the dialog with the following details:

- Marker:** A dropdown menu set to 'Auto'.
- Protocol List:** A list of protocols with checkboxes. 'NFS' is checked and highlighted with a red box. 'CIFS' is unchecked.
- Buttons:** 'Select All' and 'Next'. The 'Next' button is highlighted with a red box.

- 4 If needed, fill in the **NFS Client Access** options and click **Next**.

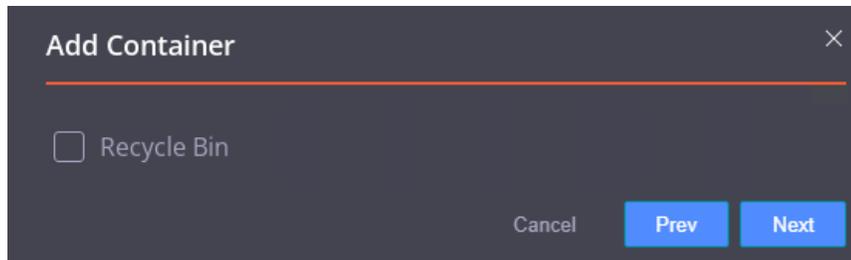
Figure 21: Add Container window – Client Access



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

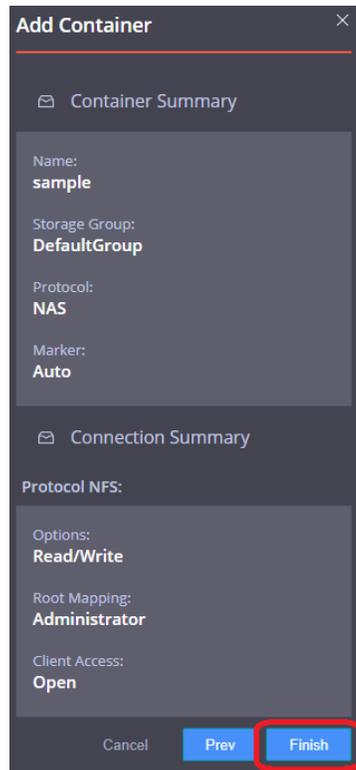
- 5 Optionally, enable the Recycle Bin feature, and then click **Next**. For more information, see the *QoreStor User Guide*.

Figure 22: Add Container – Recycle Bin



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

Figure 23: Add Container - Summary



Adding the QoreStor NFS container as a repository in Veeam

i **NOTE:** The Veeam Server is supported on Windows only. To configure an NFS container from QoreStor as a backup repository, use a Linux server for mounting the NFS container.

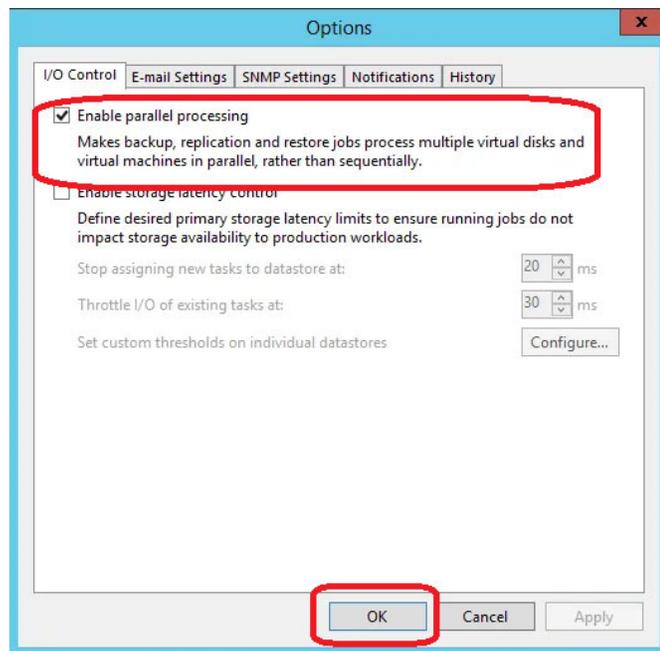
! **CAUTION:** To maximize the QoreStor and Veeam deduplication savings and performance, Quest recommends using the settings provided in this guide for all the data being backed up.

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 same settings.

To add the QoreStor NFS container as a repository in Veeam

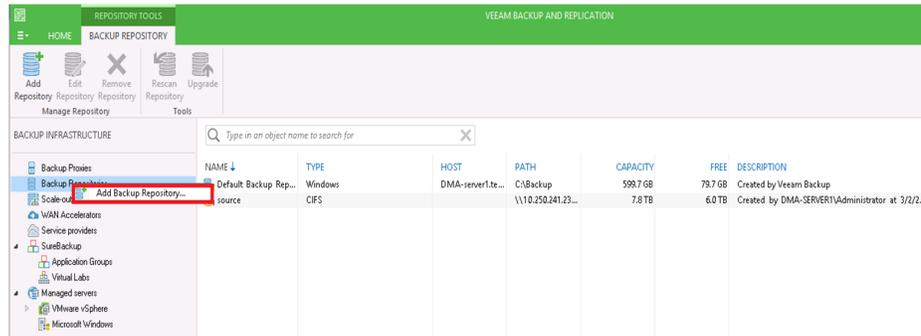
- 1 Open the Veeam Backup & Replication console.
- 2 If using Veeam 9.5 U3 or earlier, select the dropdown **Menu** and click **General Options**
- 3 On the I/O Control tab, select **Enable parallel processing** and click **OK**. This option is missing in Veeam 9.5 U4 and later as it is automatically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

Figure 24: Veeam parallel processing options



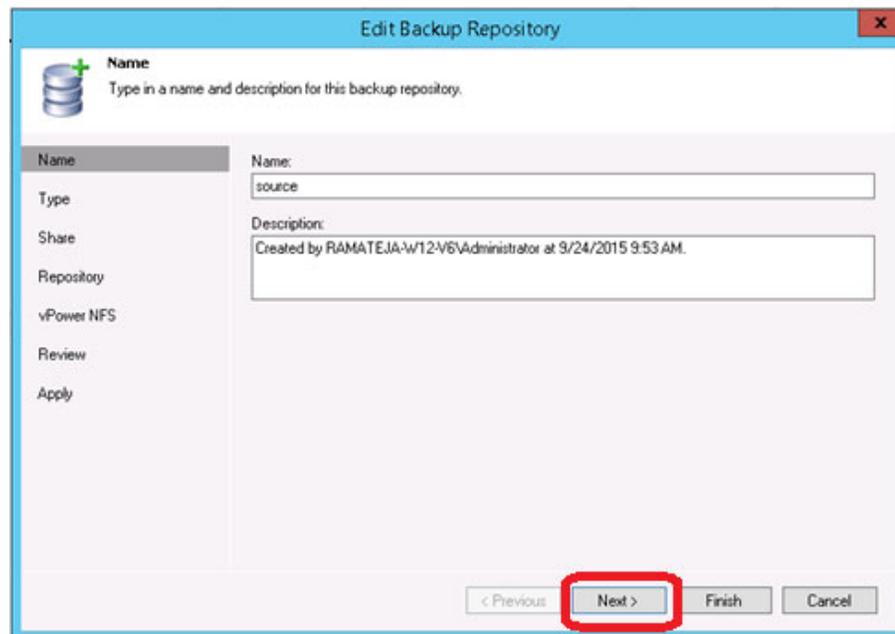
- 4 In the **Backup Infrastructure** section, right-click **Backup Repositories**, and select **Add Backup Repository**.

Figure 25: Veeam Backup Repository page



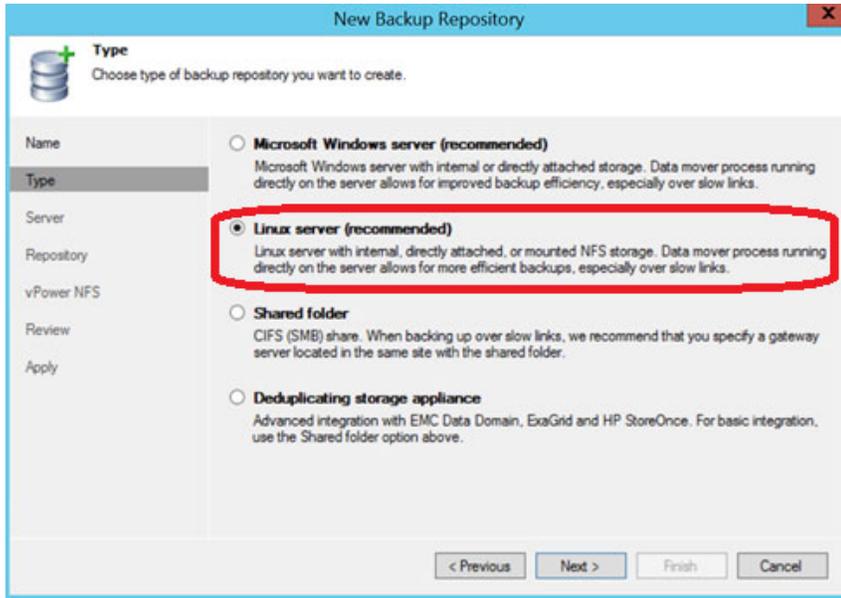
- 5 Enter a name for the QoreStor container repository and click **Next**.

Figure 26: Veeam Edit Backup Repository window - Name



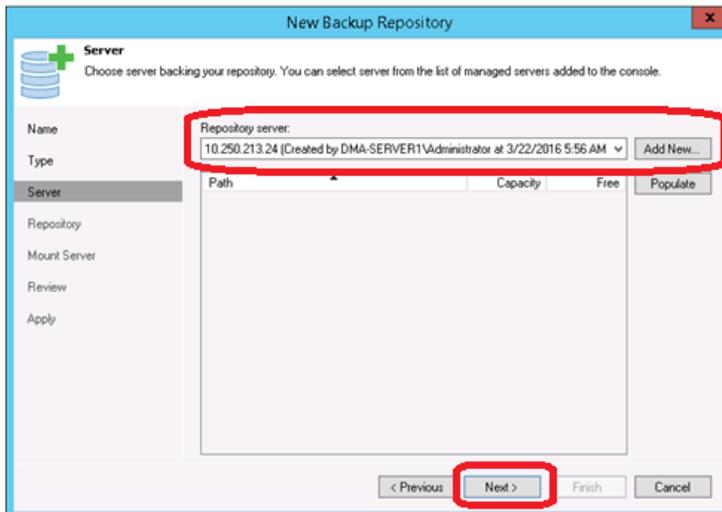
- 6 Select **Linux Server (recommended)** as the type of backup repository, and then click **Next**.

Figure 27: Veeam New Backup Repository window - Type



- 7 Complete one of the following options, and then click **Next**:
 - Under **Repository server**, enter the name of the repository on the Linux server and click **Add New**.
 - If you previously added it, select the server from the list.

Figure 28: Veeam New Backup Repository window - Server



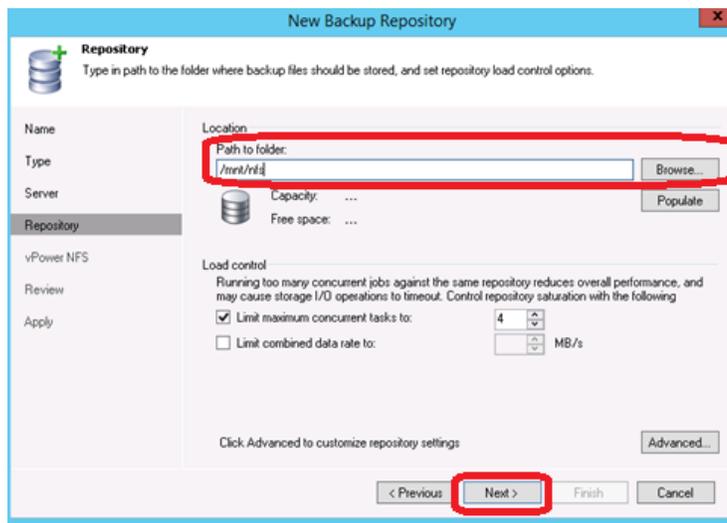
- 8 Mount the QoreStor NFS Container onto a Linux server.

Figure 29: Mounting the container in the Linux command line

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

- 9 Under Location, enter the container mount path. To customize the repository settings, click **Advanced**.

Figure 30: Veeam New Backup Repository window - Repository

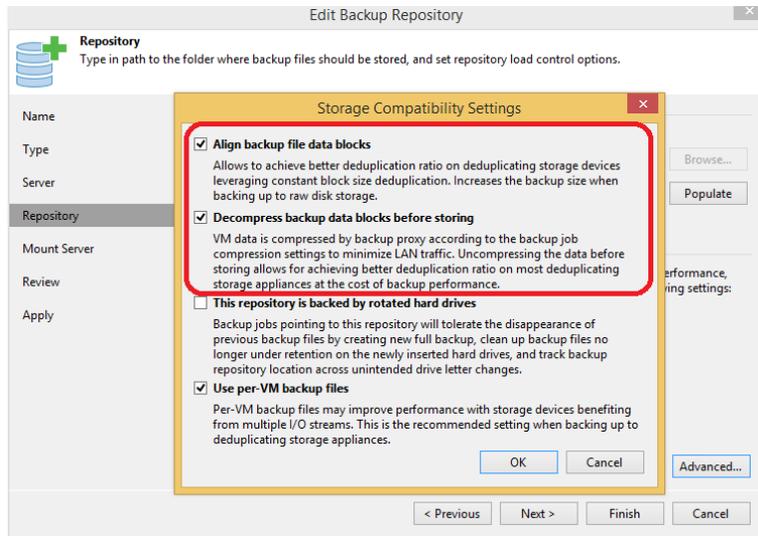


i **NOTE:** For the maximum concurrent jobs supported for CIFS/NFS, see the *QoreStor Interoperability Guide*. The maximum concurrent tasks also depend upon the number of CPU cores of Veeam Servers or proxies.

- 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. Quest does not recommend changing these settings after data has been written to QoreStor.

Figure 31: Veeam storage compatibility settings for data blocks

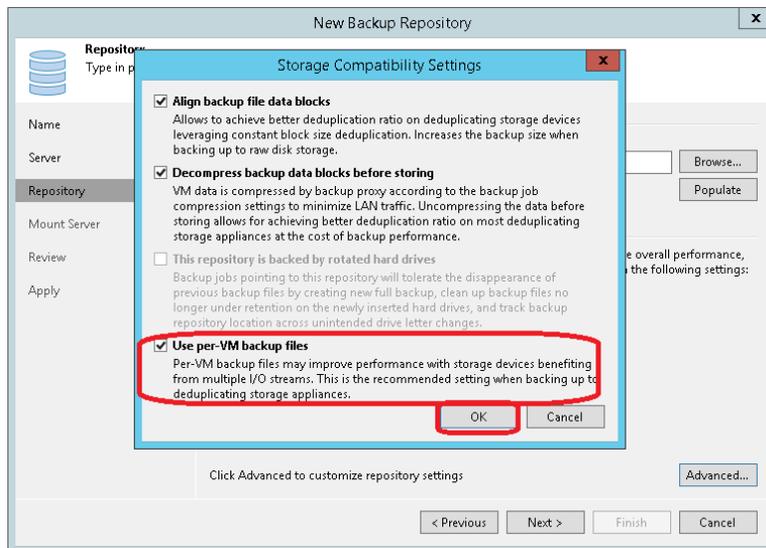


CAUTION: If you change the setting for the Align backup file data blocks option after backups are taken, it will impact the deduplication savings for future backups.

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.

Figure 32: Veeam storage compatibility settings for VM backup files



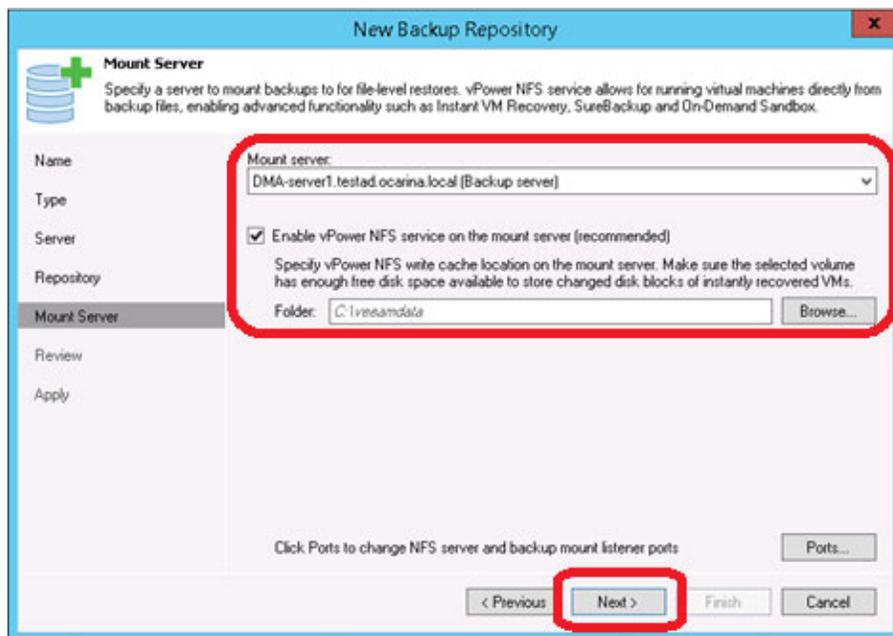
i **NOTE:** If using Veeam 9.5 U3 or earlier, confirmed that you selected **Enable parallel data processing** in Step 3.

i **NOTE:** The Use Per-VM backup files option allows multiple write streams within a single job with parallel processing enabled. Quest recommends selecting this option as it dramatically improves overall job backup performance.

12 Click **Next**.

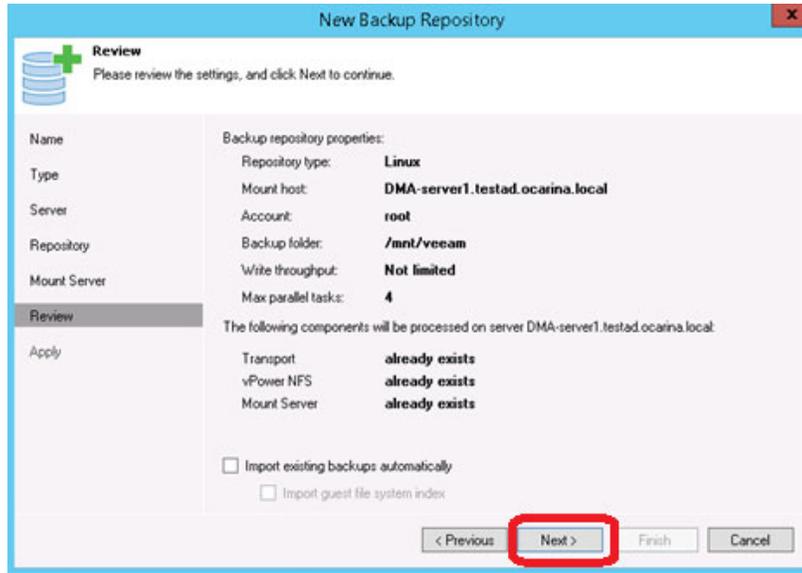
13 Optionally, to use the Instant Recovery feature, select **Enable vPower NFS service on the mount server (recommended)**, and then click **Next**.

Figure 33: Veeam New Backup Repository window – Mount Server



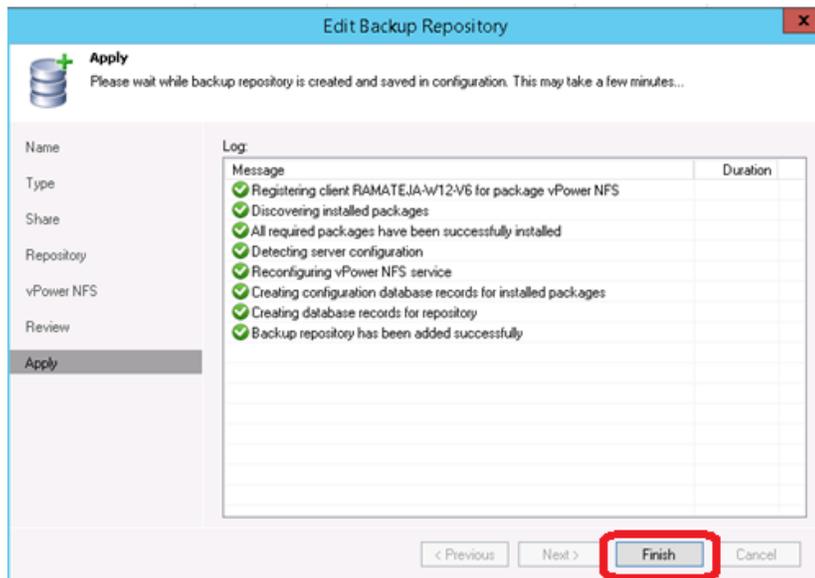
14 On the review page, verify the settings, and click **Next** to apply changes.

Figure 34: Veeam New Backup Repository window – Mount Server



15 Click **Finish**.

Figure 35: Veeam Editor Backup Repository window - Apply



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 prerequisite

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

i **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 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.

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

To install Rapid CIFS on a Veeam Windows proxy

- 1 Download the MSI to the Server/Proxy from 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.

Figure 36: Rapid CIFS Filter Driver Setup Wizard Welcome screen



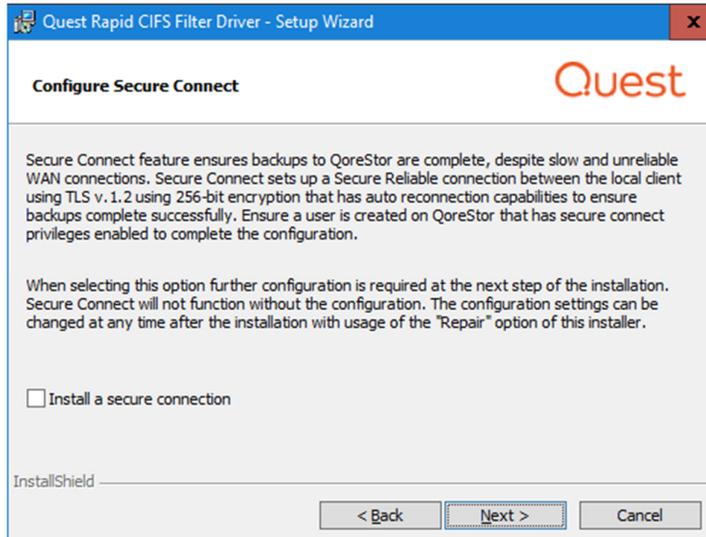
- 5 Read and accept the license agreement to proceed. Click **Next** when ready.

Figure 37: Rapid CIFS Filter Driver license agreement



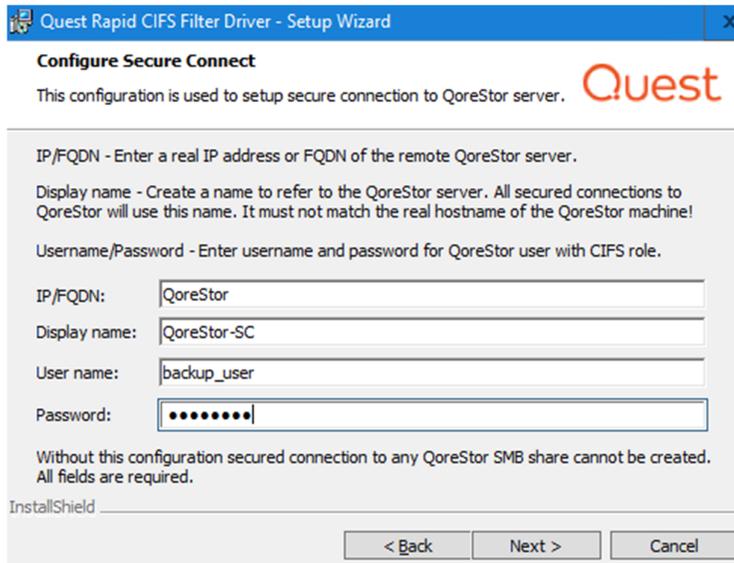
- 6 If installing with Secure Connect for WAN, select **Install a secure connection** and click **Next**.

Figure 38: Rapid CIFS Filter Driver Setup Wizard – Configure Secure Connect



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

Figure 39: Rapid CIFS Filter Driver Setup Wizard – IP/FQDN credentials



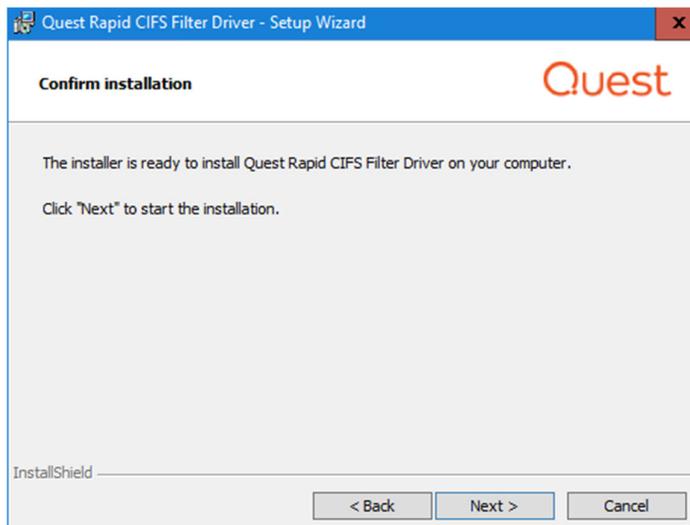


NOTE: When accessing the share from this server use the **Display Name** when accessing the share to leverage Secure Connect. I.E //QoreStor-SC/share

Use the normal IP/FQDN to access WITHOUT secure connect.

- 8 Click **Next**.

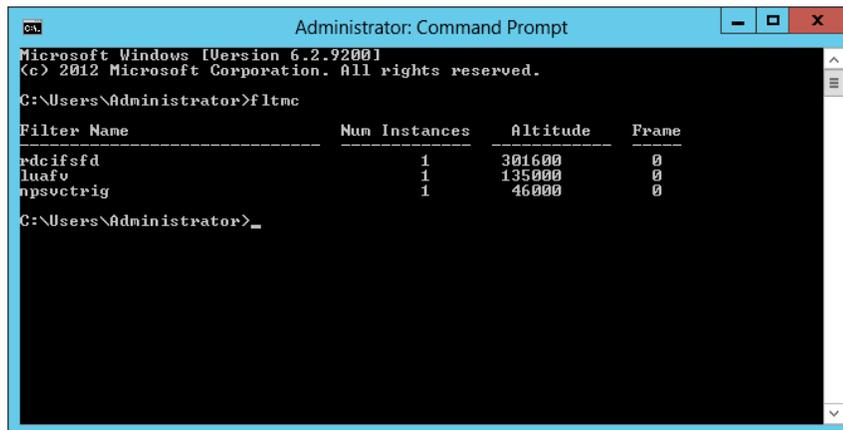
Figure 40: Rapid CIFS Filter Driver Setup Wizard – Confirm Installation



- 9 After the install finishes click **Finish**.

You can optionally verify that the “**rdcifsfd**” driver is loaded automatically; this can be checked by using the command **fltmc**.

Figure 41: Administrator Command Prompt verification

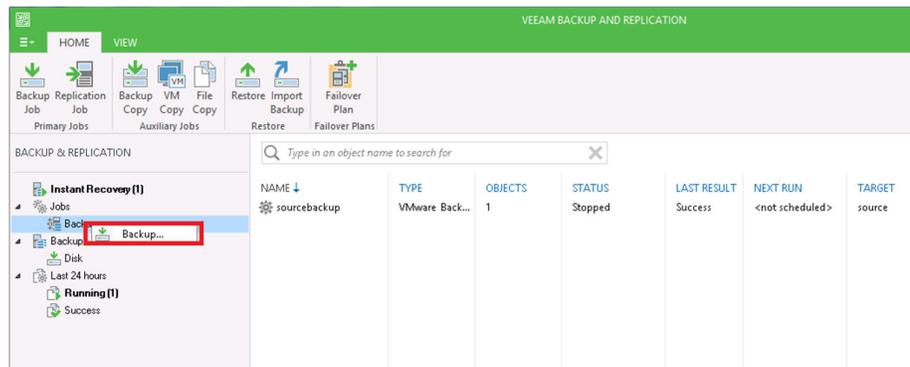


Creating a backup job with the QoreStor system as a target

To create a backup job with the QoreStor system as a target

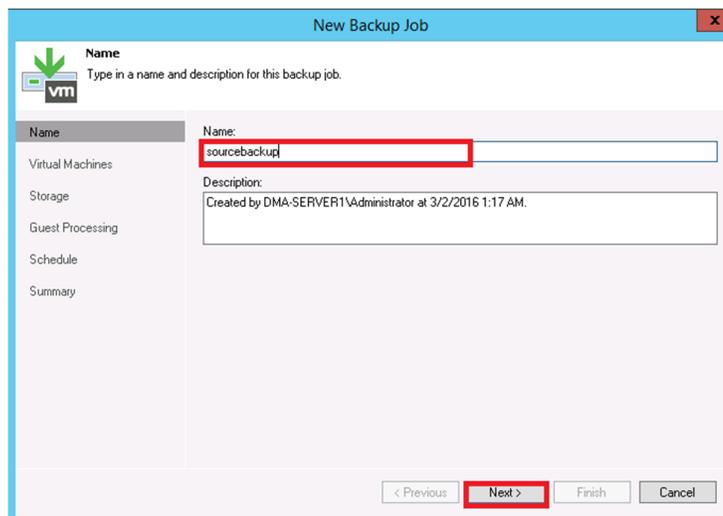
- 1 On the **Backup & Replication** menu, go to **Jobs > Backup**, and right-click **Backup** to create a new backup job.

Figure 42: Veeam Backup page



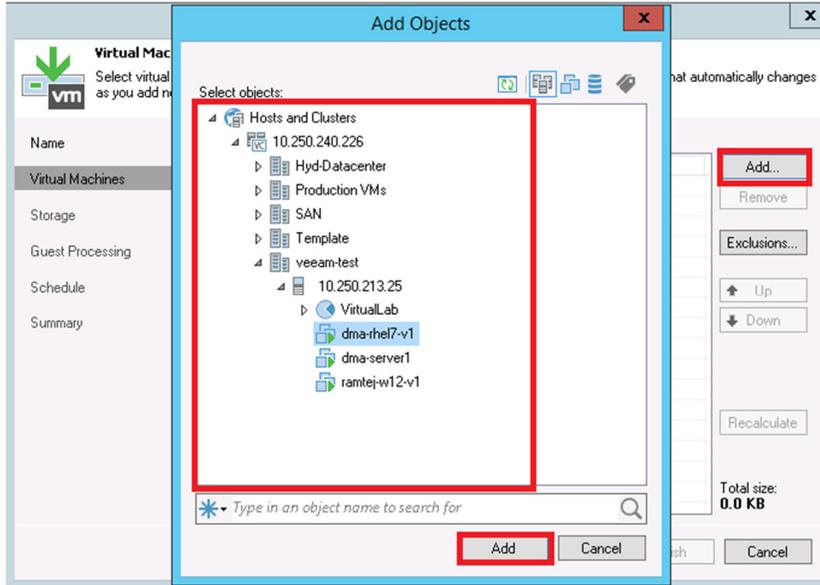
- 2 Provide the backup job name and click **Next**.

Figure 43: Veeam New Backup Job window



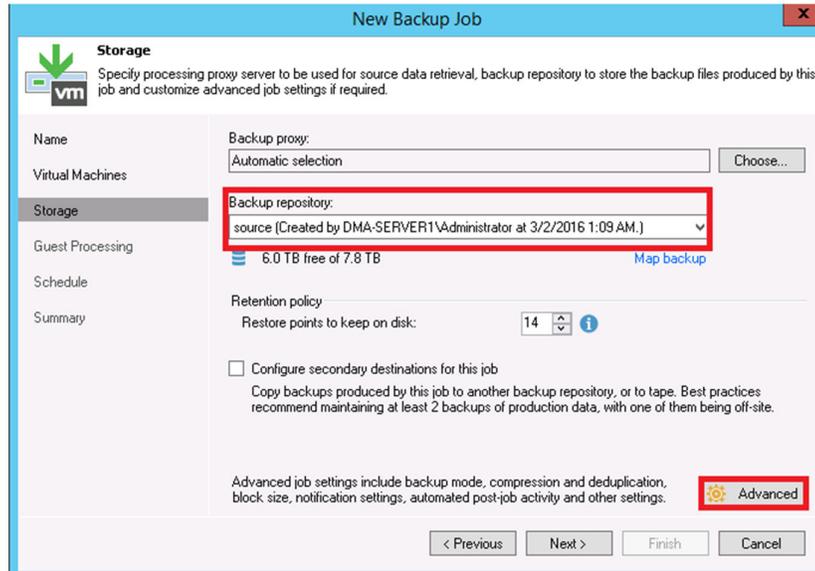
- 3 Select one or more virtual machines, data stores, resource pools, vApps, SCVMM clusters, etc. for backup.

Figure 44: Veeam Add Objects window



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

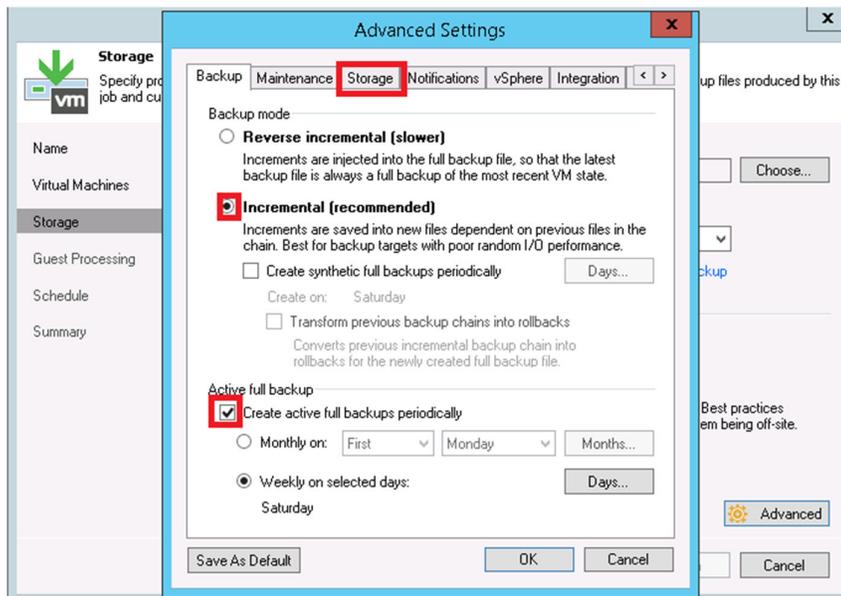
Figure 45: Veeam backup repository



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

i **NOTE:** 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 as if it was running for the first time. Due to read performance requirements during the synthetic operation, only use the Synthetic Full backup option with a Veeam Ready Repository QoreStor instance.

Figure 46: Veeam backup job advanced settings

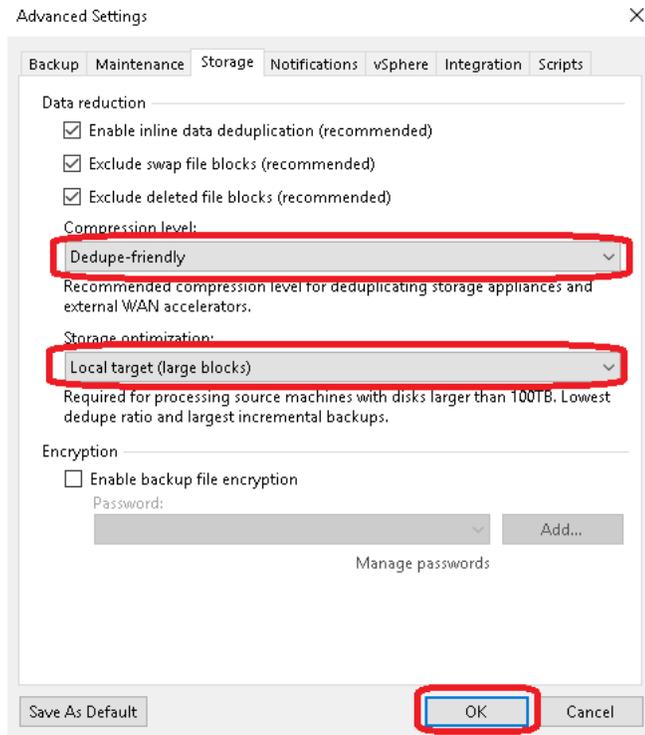


For information on configuring Fast Clone options for Hyper-V 2016 ReFS VM's see QoreStor and Veeam Fast Clone for Hyper-V 2016 backups or Data Copy.

i **NOTE:** Veeam recommends against very long retention combined with infrequent active or synthetic full backups. Run a full backup at least once a month. For specific recommendations, contact Veeam.

- 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).

Figure 47: Veeam advanced storage settings

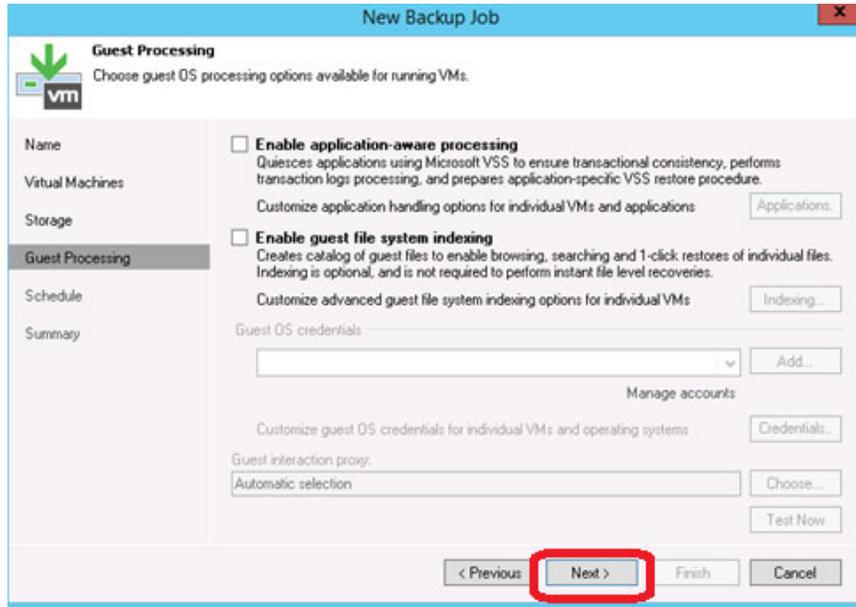


NOTE: For best balance between backup performance and deduplication savings it is recommended to choose these options for all of the backup jobs written to QoreStor.

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.

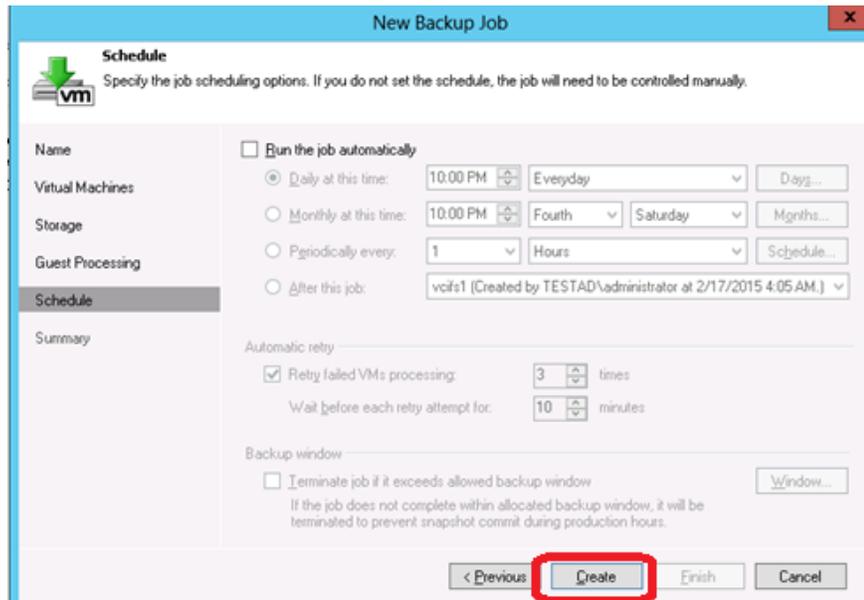
- 7 Enable any optional settings required by your workflow and click **Next**.

Figure 48: Veeam New Backup Job window – Guest Processing



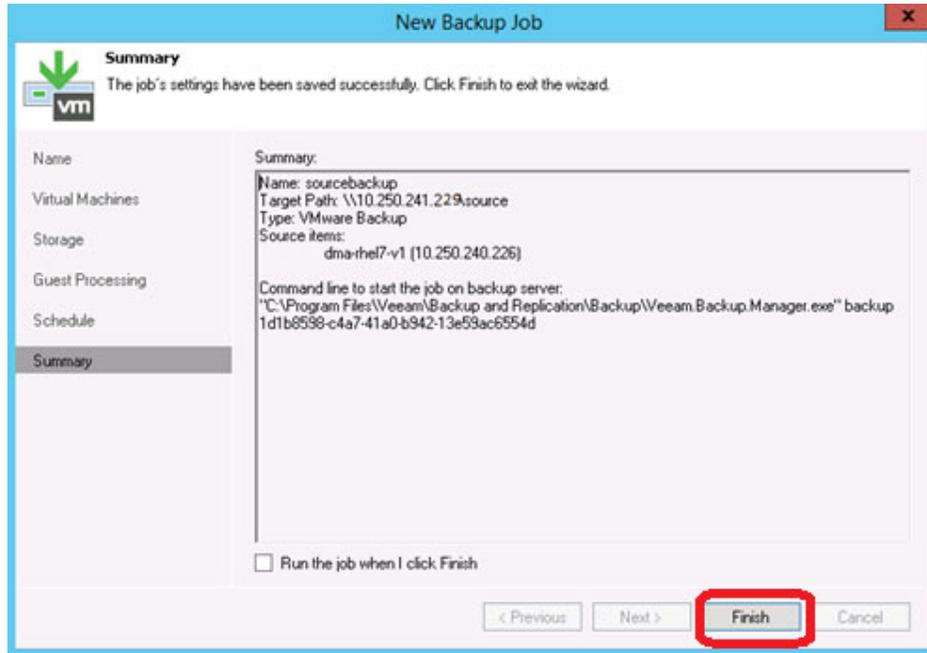
- 8 Schedule the backup and click **Create**.

Figure 49: Veeam New Backup Job window - Schedule



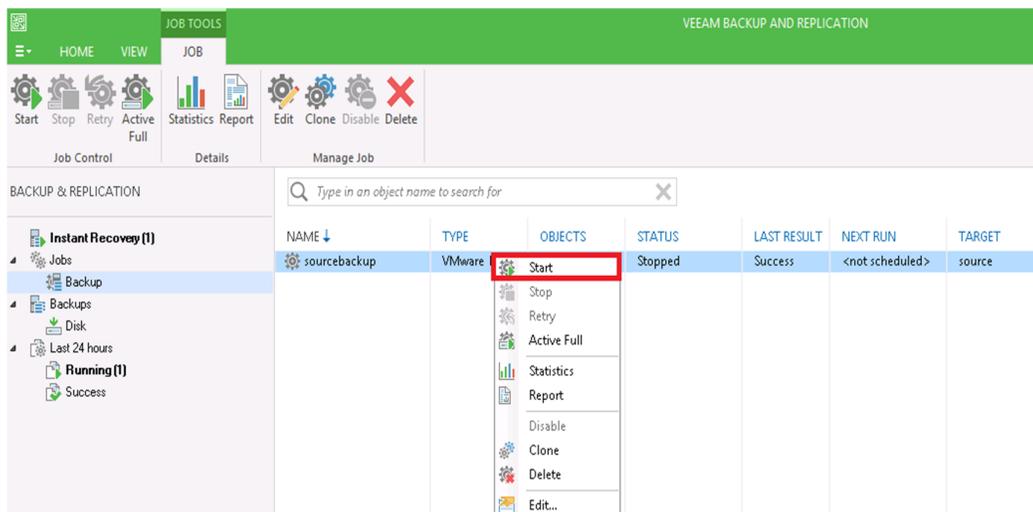
9 Click **Finish**.

Figure 50: Veeam New Backup Job - Summary



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

Figure 51: Veeam backup job start



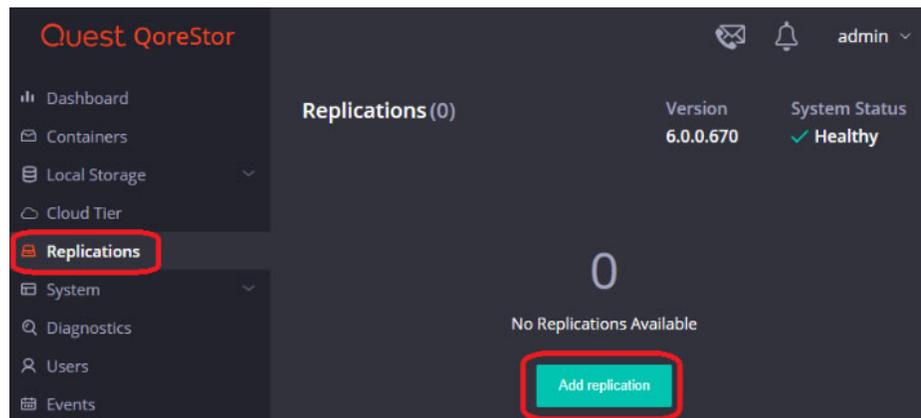
Setting up QoreStor system replication

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

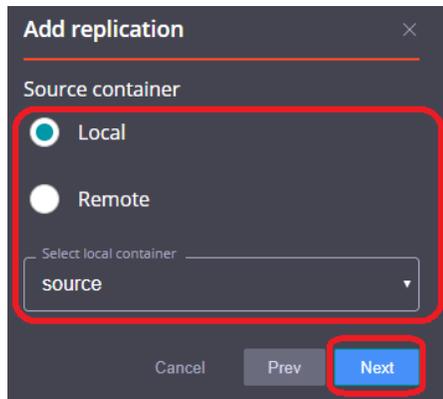
- 1 Create a source container on the source QoreStor system.
- 2 Create a target container on the target QoreStor system.
- 3 On the source QoreStor system, go to the **Replications** Tab. Click the **Add replication** button.

Figure 52: QoreStor Replications page



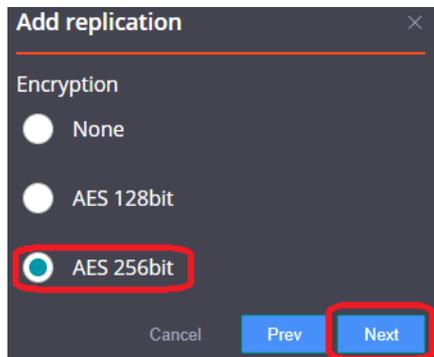
- 4 Select the source Container for Replication and click **Next**.

Figure 53: Add replication – Source container



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

Figure 54: Add replication - Encryption



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

Figure 55: Add replication – Target container

Add replication [X]

Target container

Local

Remote

Username: admin

Password:

Remote Machine: qspl-6300-47.systemst.ocarina.local

Retrieve Containers

select remote container: target

Cancel Prev Next

- 8 Specify any **Bandwidth Limitations** needed in MBps, leaving 0 for unlimited bandwidth. Click **Next**.

Figure 56: Add replication - Limits

Add replication [X]

Limits

Bandwidth Limit for the peer machine [MBps]: 0

Cancel Prev Next

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

Figure 57: Add replication – Summary

Add replication [X]

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**

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

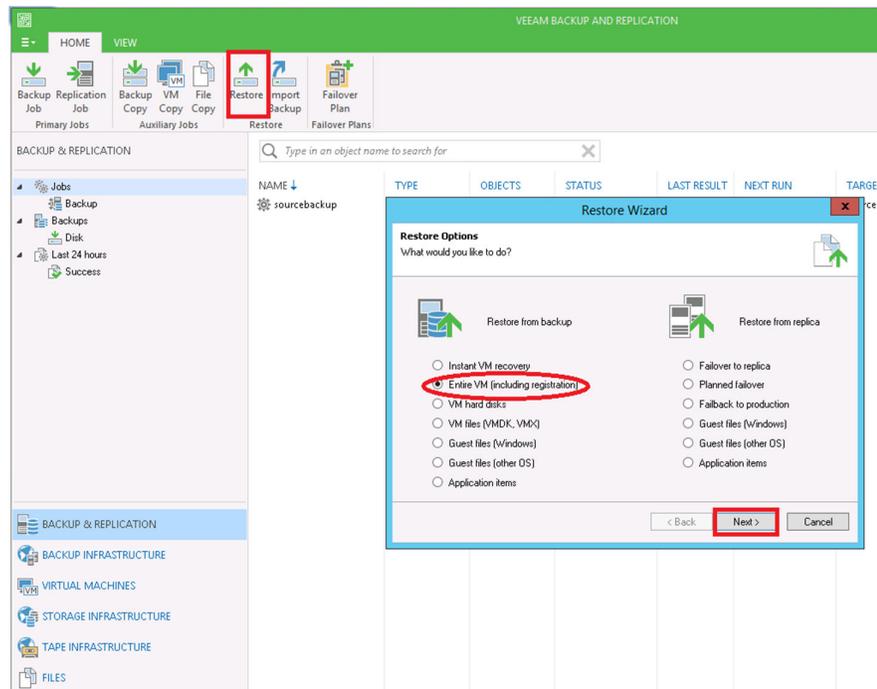
Restoring from the replication target

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.

To restore from the replication target

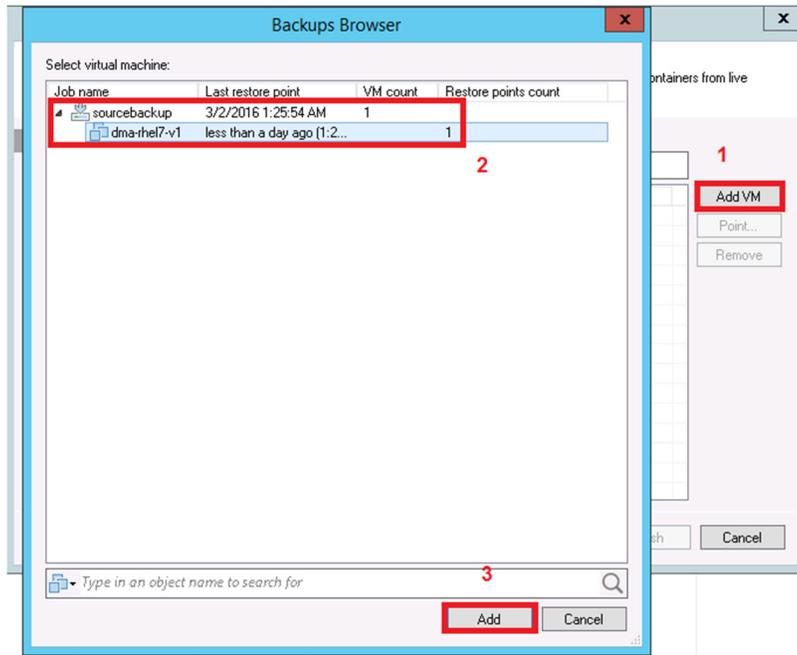
- 1 Add the target QoreStor system container to the Veeam repository. For instructions, see [Creating a CIFS container for use with Veeam](#) or [Creating an NFS container for use with Veeam](#). Update all backup jobs that use the source QoreStor system container as a repository and change them to use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

Figure 58: Veeam Restore Wizard



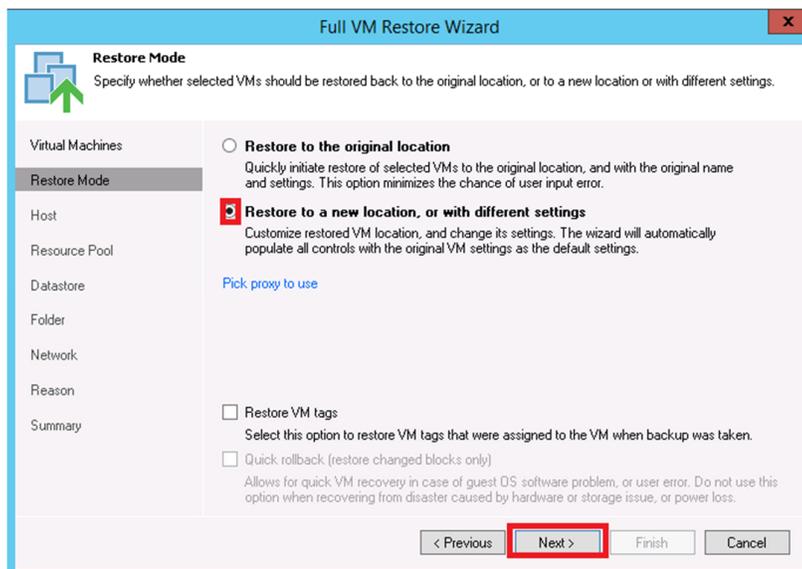
- 3 Click **Add VM** and select **From backup**. Select the VM you want to restore and click **Add**.

Figure 59: Veeam Backups Browser



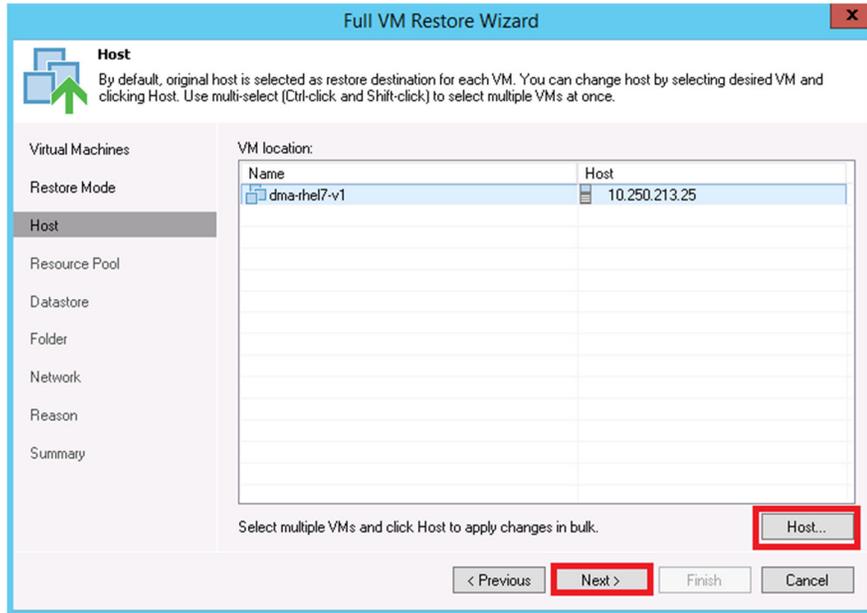
- 4 Select a Restore Mode and click **Next**.

Figure 60: Veeam Full VM Restore Wizard



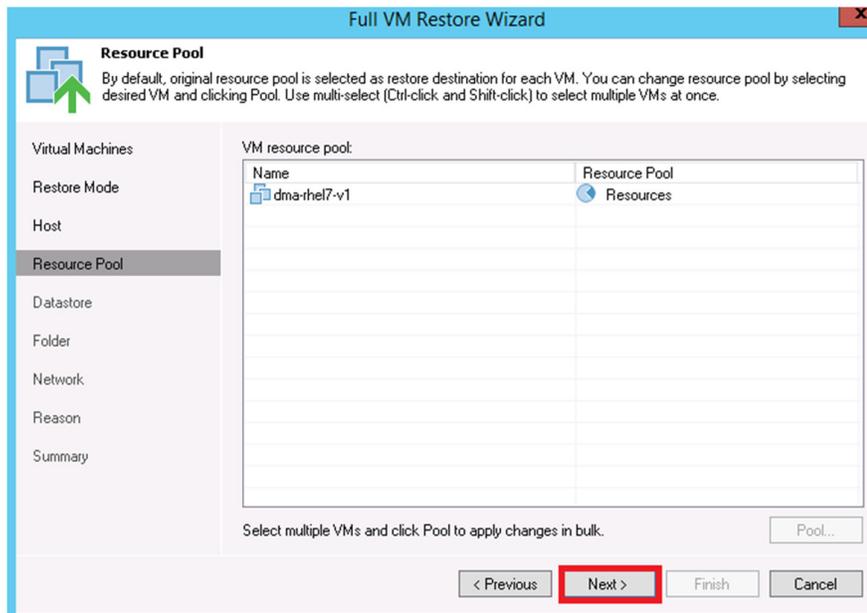
- 5 Provide the Host details as per requirement and click **Next**.

Figure 61: Veeam Full VM Restore Wizard – VM location



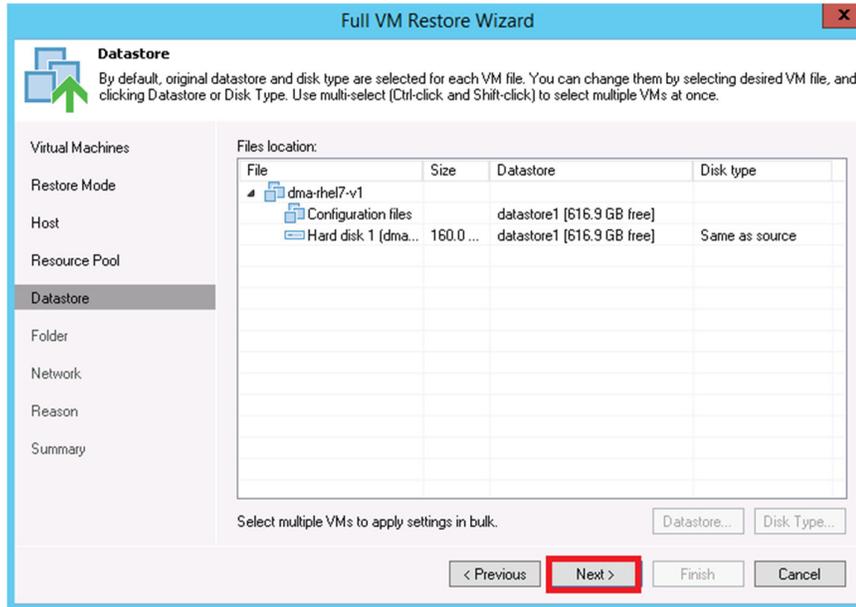
- 6 Select the resource pool and click **Next**.

Figure 62: Veeam Full VM Restore Wizard – Resource Pool



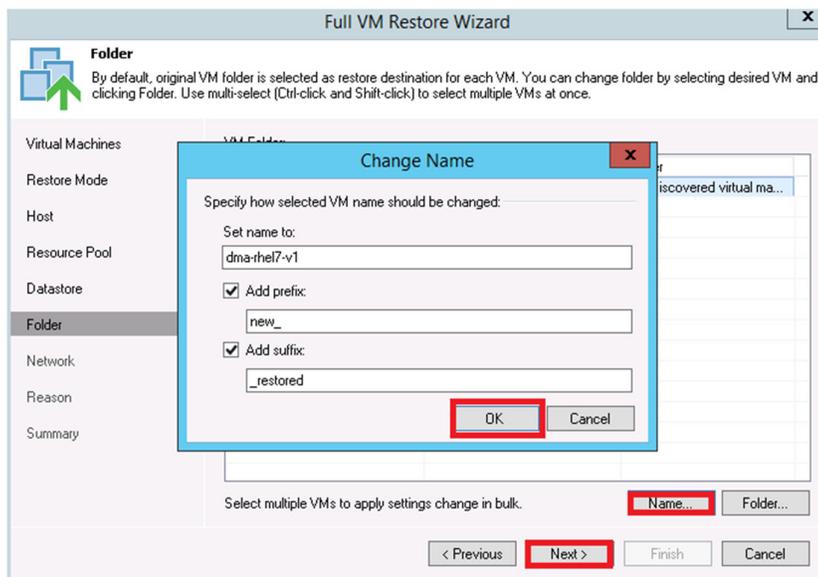
- 7 Select the data store and disk type, and then click **Next**.

Figure 63: Veeam Full VM Restore Wizard - Datastore



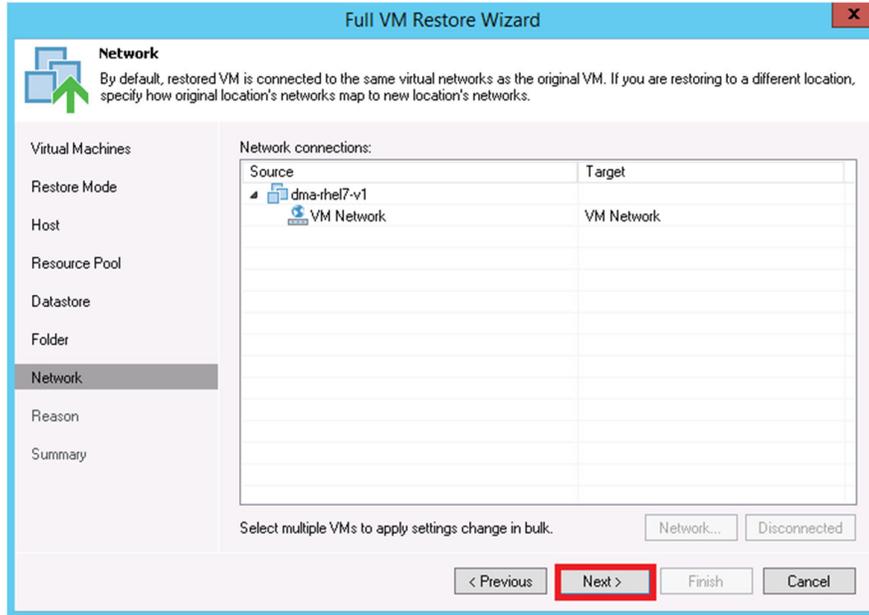
- 8 Provide the new name for the restored VM by clicking **Name**, entering a name, and clicking **OK**. Click **Next**.

Figure 64: Veeam Full VM Restore Wizard – Change Name



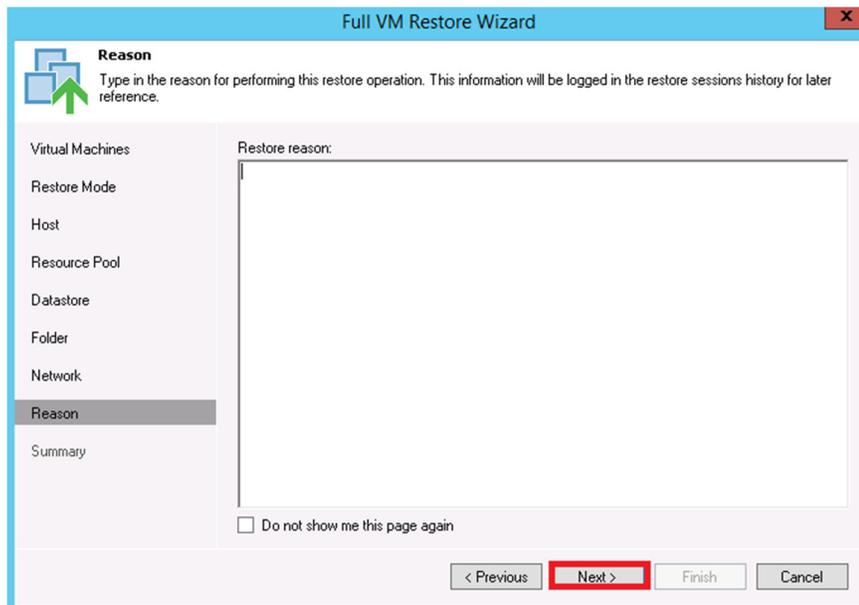
- 9 Select the network location and click **Next**.

Figure 65: Veeam Full VM Restore Wizard - Network



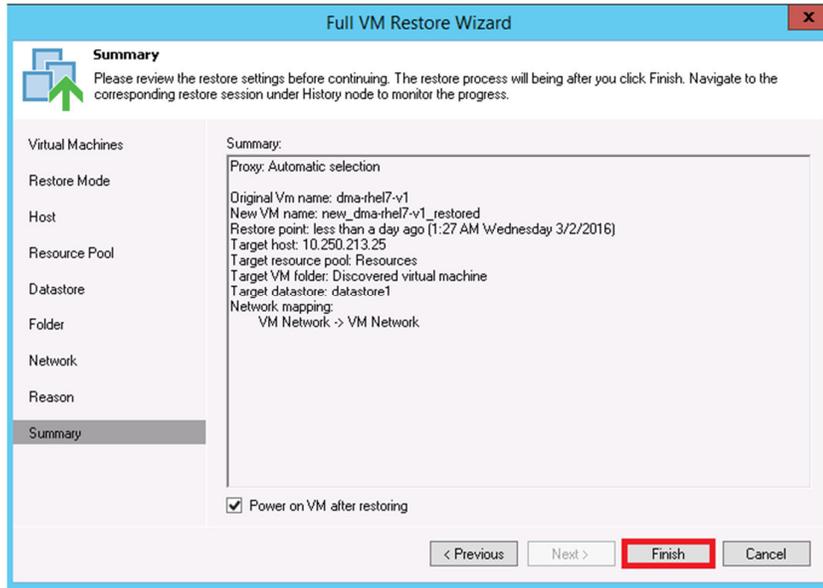
- 10 Provide the reason for the restore and click **Next**.

Figure 66: Veeam Full VM Restore Wizard - Reason



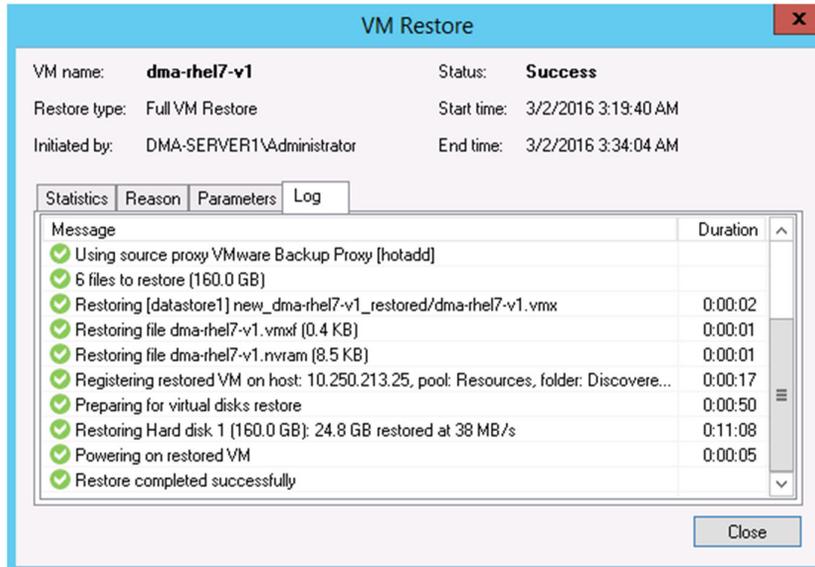
11 Click **Finish**.

Figure 67: Veeam Full VM Restore Wizard - Summary



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

Figure 68: Veeam Backup & Replication menu



QoreStor for Scale-Out Capacity

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

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 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.

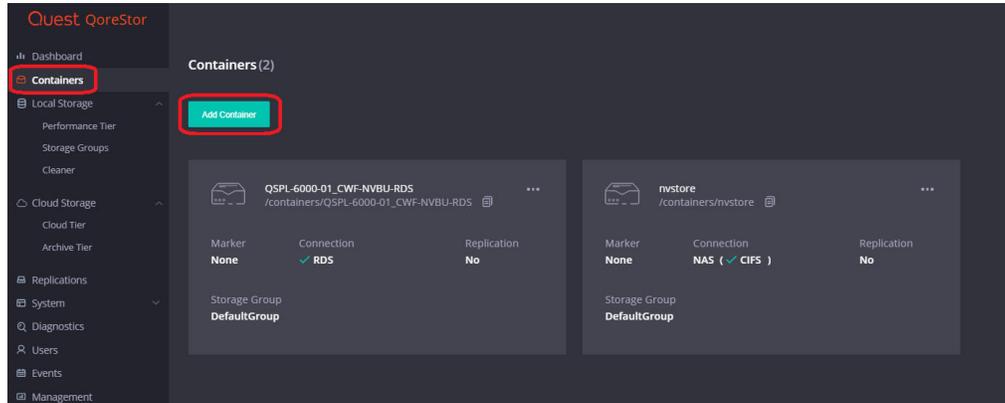
Scale-Out repositories work by first creating basic repositories. Then you create a scale-out repository adding the initial performance tiers and capacity tiers already added as basic repositories.

Creating an Object Container(S3) in QoreStor

To create an Object Container(S3) in QoreStor

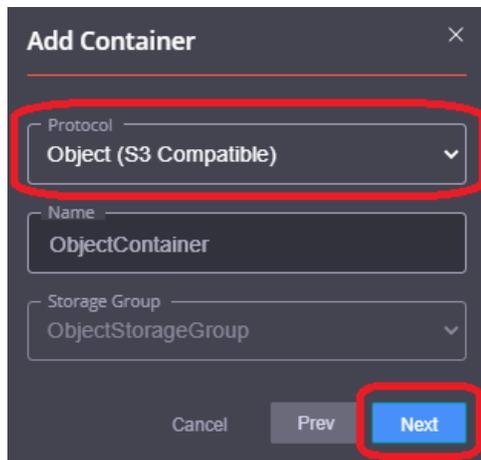
- 1 From the QoreStor UI select **Containers** then click **Add Container**.

Figure 69: QoreStor Containers page



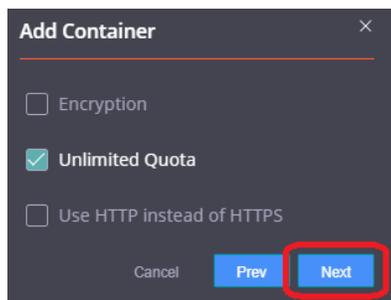
2. Select the **Protocol** drop down and set it to **Object (S3 Compatible)**, and then click **Next**.

Figure 70: Add Container window - Protocol



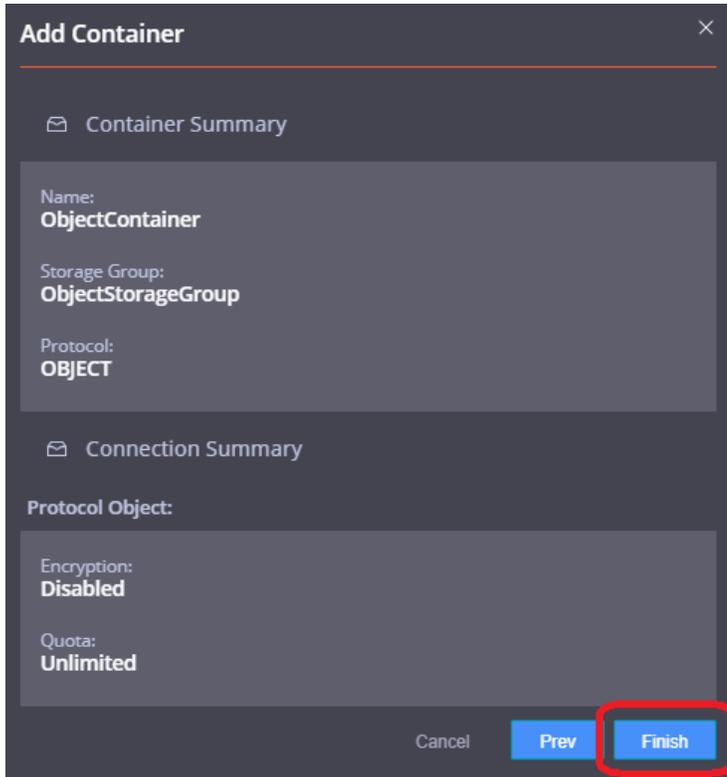
3. Select **Unlimited Quota**, and then click **Next**.

Figure 71: Add Container window – Unlimited Quota



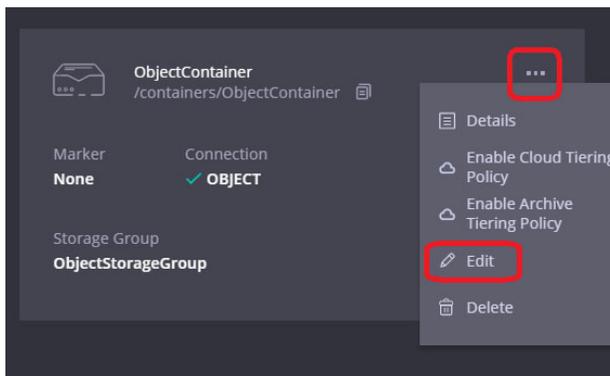
4. Verify the summary is correct and click **Finish**.

Figure 72: Add Container window - Summary



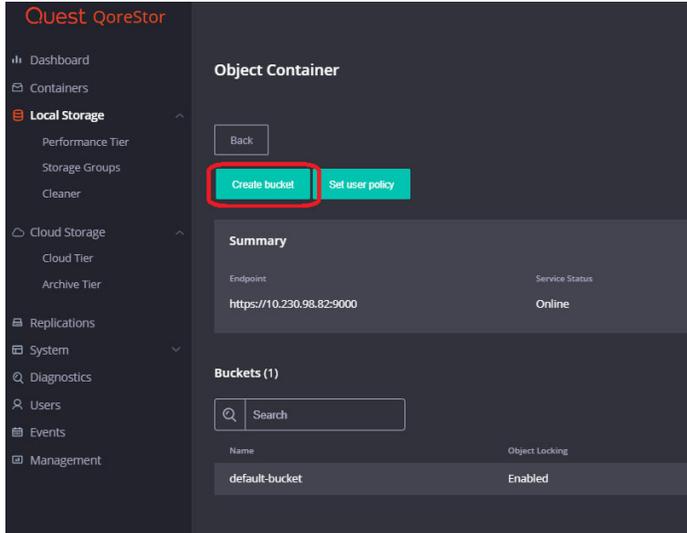
5. The Object Container is now created but we need to create a bucket other than the default. Click the **ellipsis** on the container and click **Edit**.

Figure 73: Object Container actions menu



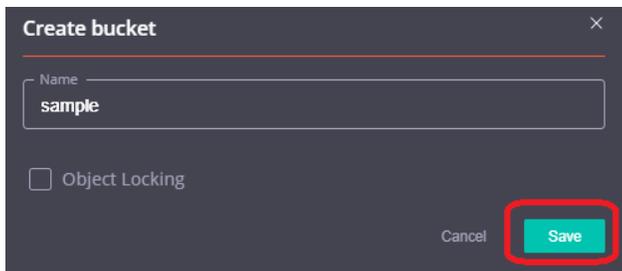
- 6 On the Object Container page click **Create bucket**.

Figure 74: QoreStor – Object Container page



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

Figure 75: Create bucket window

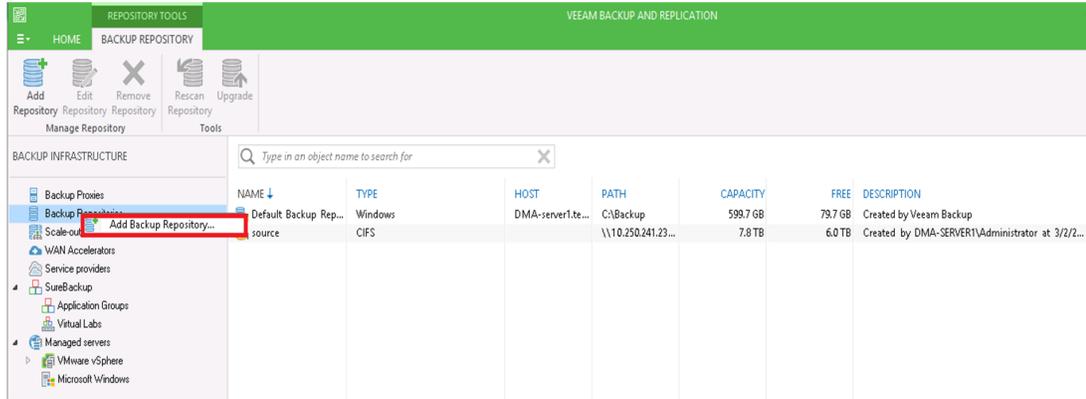


Adding the QoreStor Object Container(S3) as a repository in Veeam

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

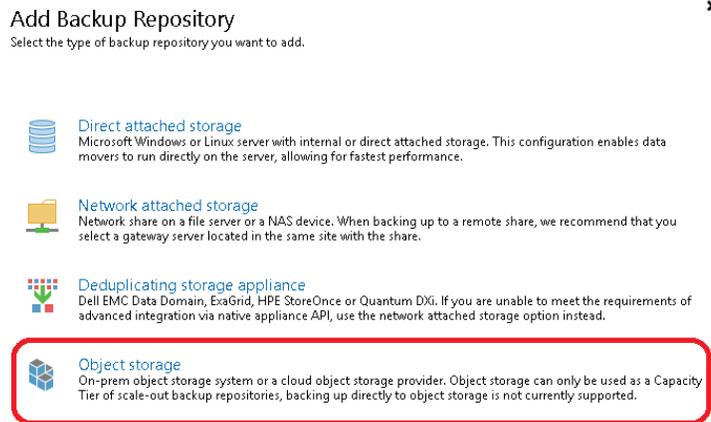
- 1 In the Backup Infrastructure section, right-click Backup Repositories, and select Add Backup Repository.

Figure 76: Veeam backup repositories



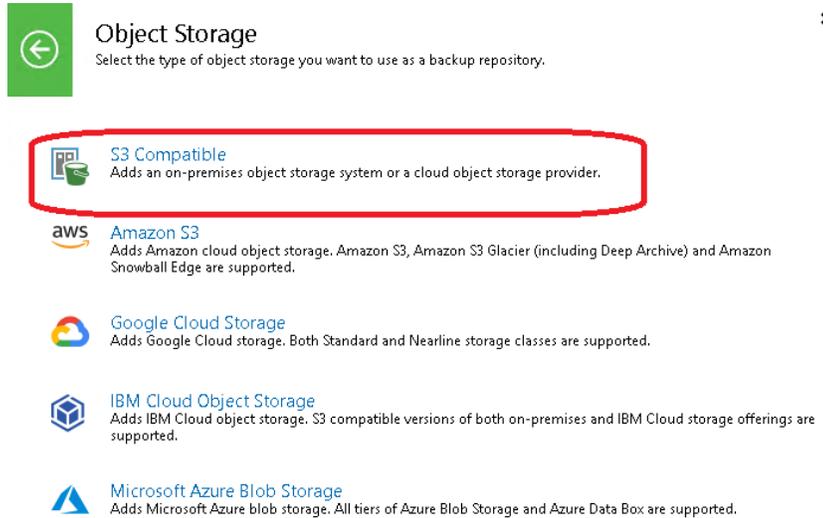
2 Click Object storage.

Figure 77: Veeam Add Backup Repository



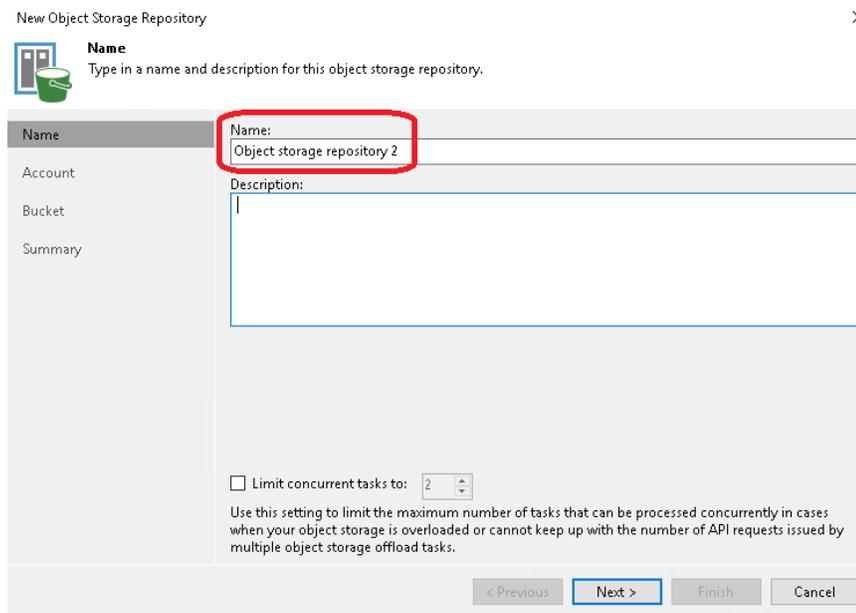
3 Click **S3 Compatible**.

Figure 78: Veeam backup repository object storage options



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

Figure 79: Veeam storage object name



- 5 Under Credentials, click **Add**.

Figure 80: Veeam storage object account

New Object Storage Repository

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

Name	Service point: https://hostname:9000
Account	Region: us-east-1
Bucket	Credentials: backup_user (last edited: 320 days ago) Add...
Summary	Manage cloud accounts

Use the following gateway server:
R720-40.systemtest.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).

Figure 81: Veeam storage object account credentials

Credentials

Access key: backup_user

Secret key: [masked]

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**.

Figure 82: Veeam storage object service point

New Object Storage Repository

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

Name: Service point: `https://hostname:9000`

Account: Region: `us-east-1`

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

Summary: [Manage cloud accounts](#)

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

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

Figure 83: Veeam Certificate Security Alert

Certificate Security Alert

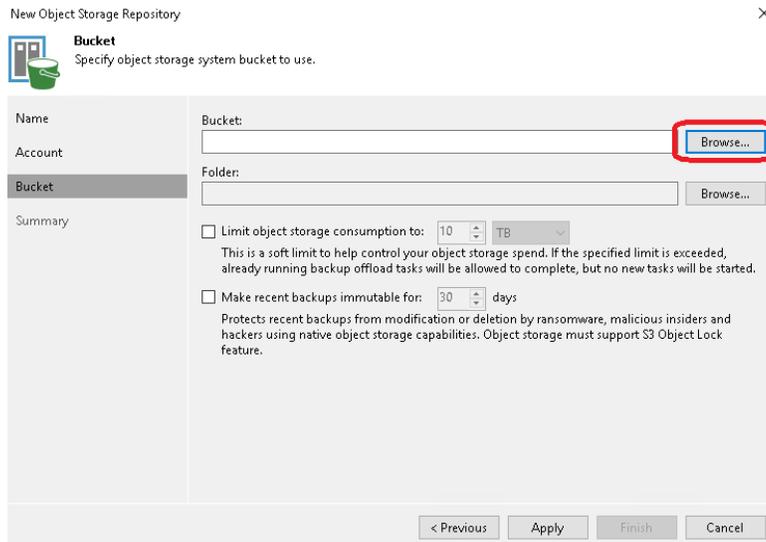
Site certificate cannot be verified. Continue anyway?

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

View... **Continue** Cancel

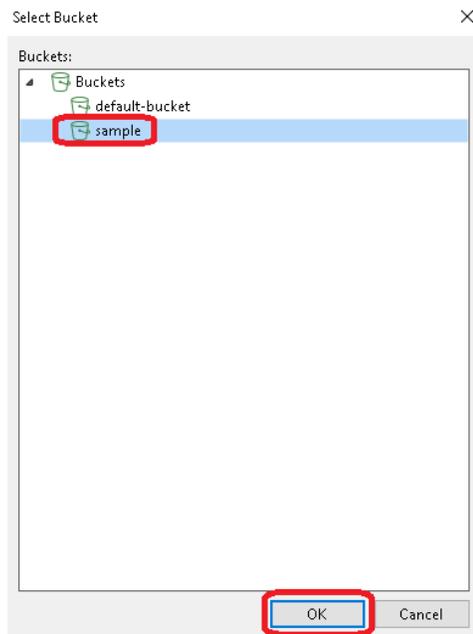
- 9 On the bucket page click **Browse...** under the bucket line.

Figure 84: Veeam storage object bucket



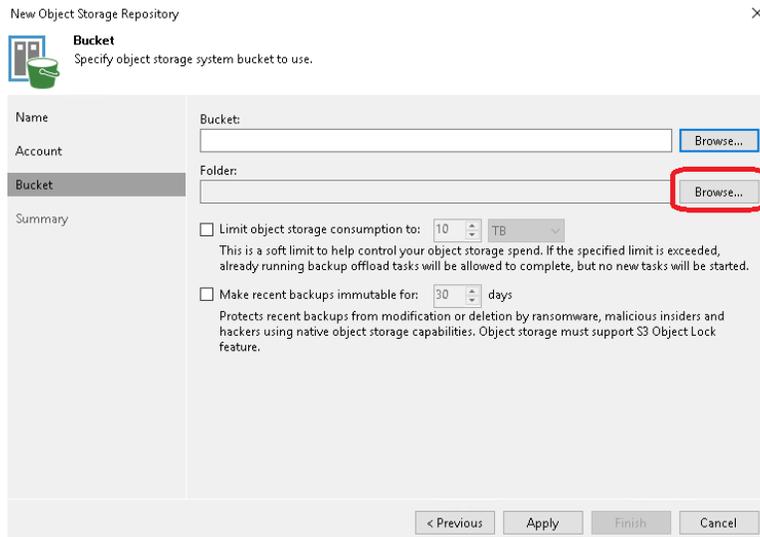
- 10 Select the bucket name created in the **Creating an Object Container(S3) in QoreStor** section of this guide. Click **OK**.

Veeam 85: Veeam storage object bucket options



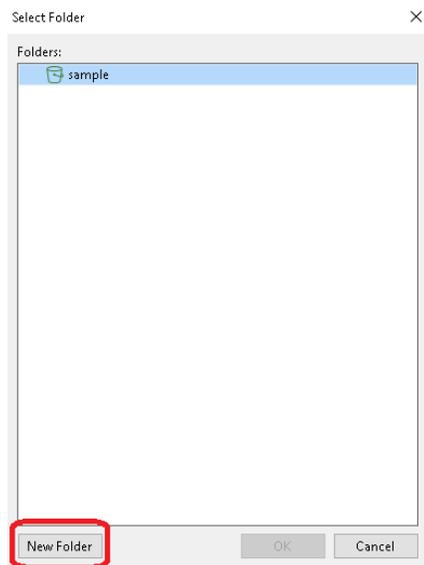
11 Back on the bucket page click **Browse...** under the folder line.

Figure 86: Veeam storage object bucket folder



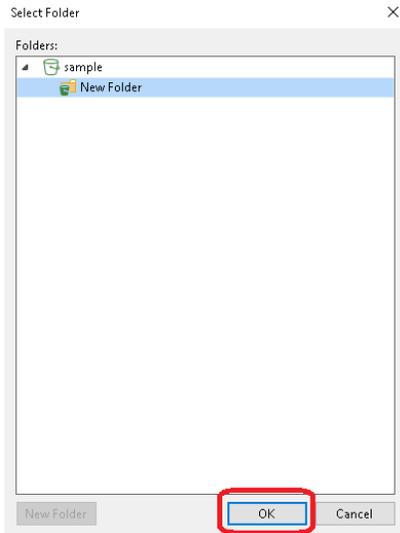
12 Click **New Folder** and define a folder name.

Figure 87: Veeam new storage object folder



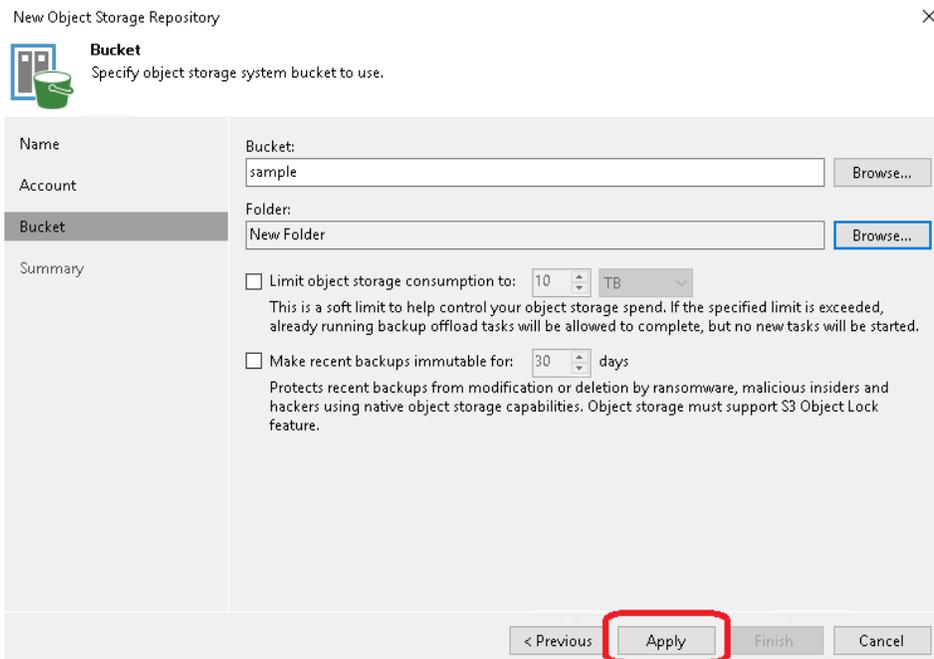
13 Select the newly created folder and click **OK**.

Figure 88: Veeam storage object with new folder



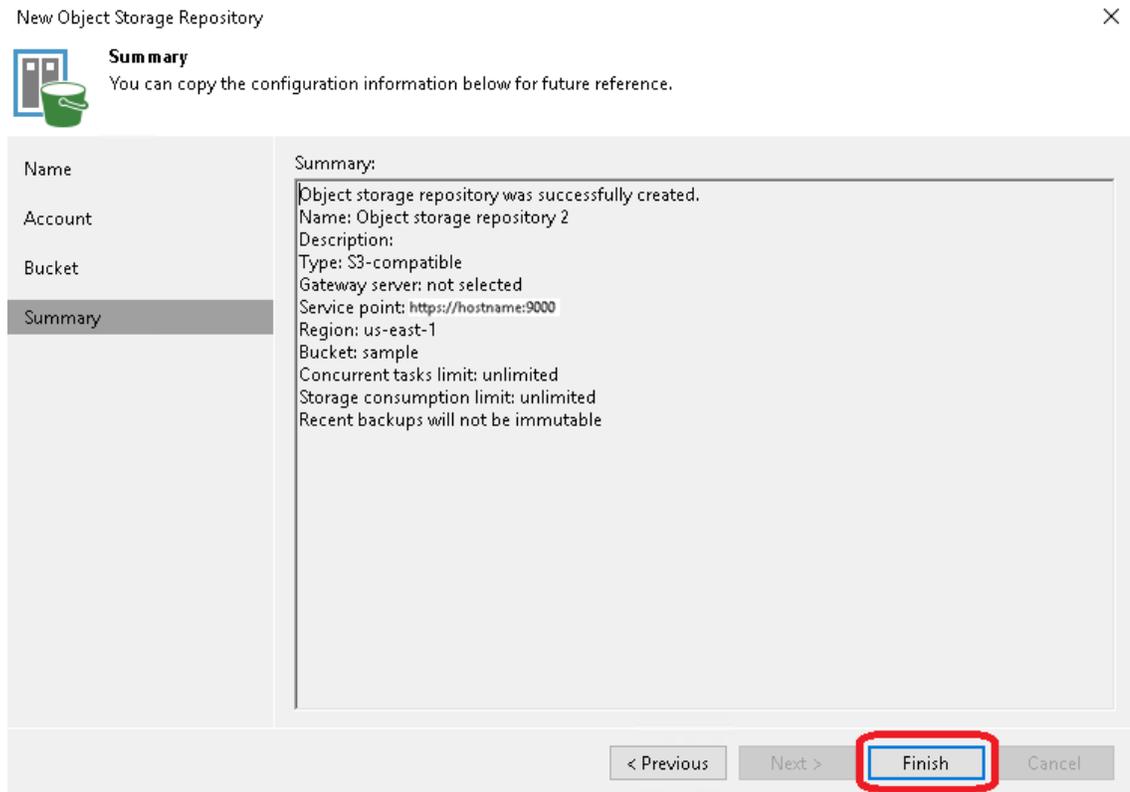
14 Back on the bucket page click **OK**.

Figure 89: Veeam storage object bucket Apply button



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

Figure 90: Veeam storage object summary

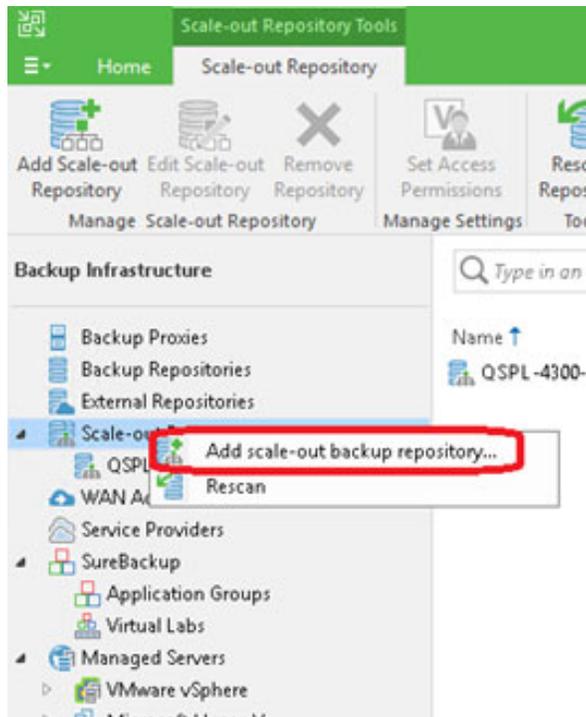


Adding the Object Container(S3) as a capacity tier to a Scale-Out repository

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

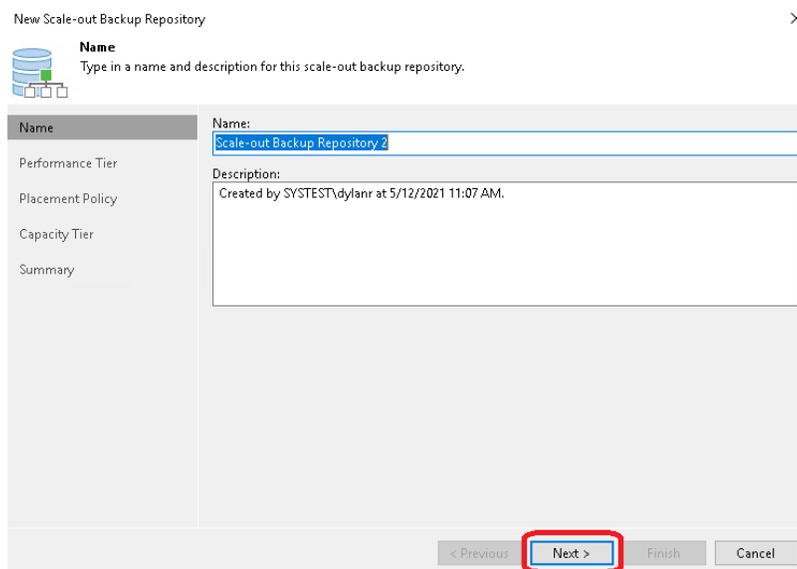
- 1 In the Backup Infrastructure section, right-click **Scale-out Repositories**, and select **Add Scale-out backup repository**.

Figure 91: Veeam Scale-out Repository page



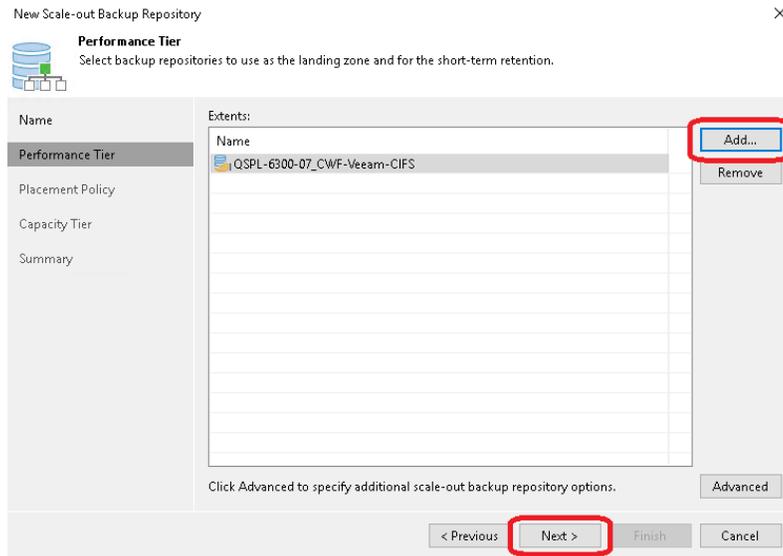
2 Click **Next**.

Figure 92: Veeam New Scale-out Backup Repository - Name



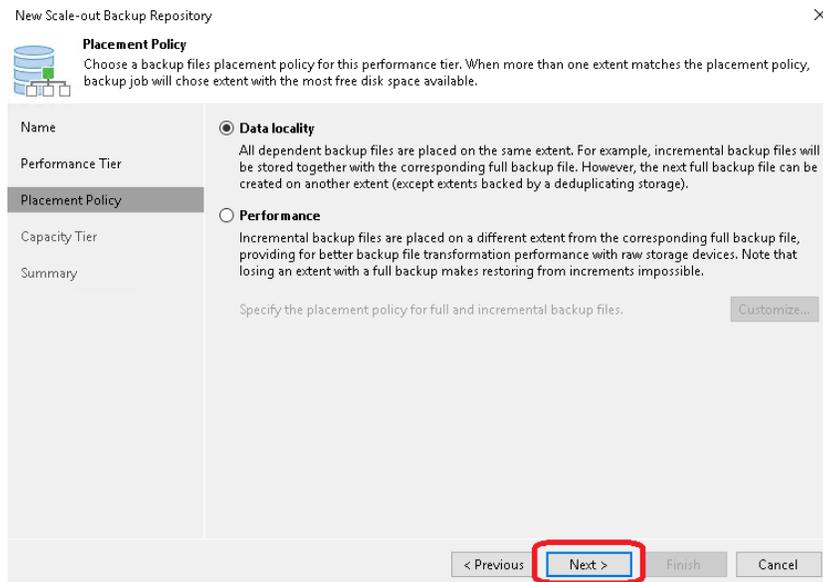
- 3 **Add** an existing spindled disk Repository or QoreStor Performance Tier based Repository to this page. Click **Next**.

Figure 93: Veeam New Scale-out Backup Repository – Performance Tier



- 4 Set your performance tier placement policy, setting this depends on the number of performance tier repositories added and resiliency of the backups required. Please reference Veeam documentation.

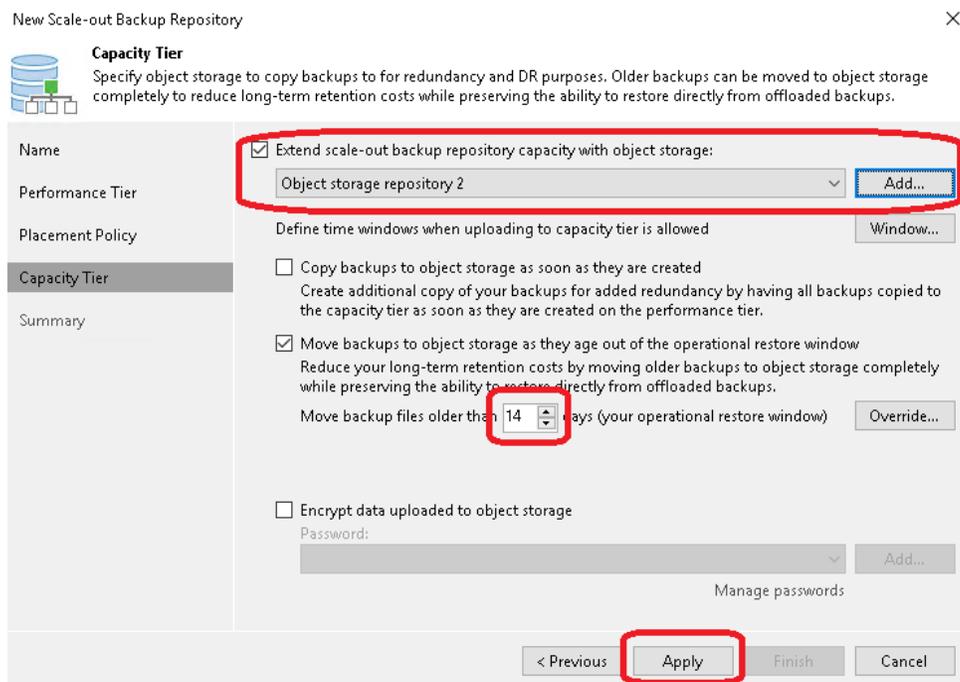
Figure 94: Veeam New Scale-out Backup Repository – Placement Policy



- 5 Select **Extend scale-out backup repository capacity with object storage**, and then do the following:
 - a. Select the object storage repository created from Adding the QoreStor Object Container(S3) as a repository in Veeam.
 - b. Set the retention age for the object repository, keeping in mind that restores are quicker from a performance tier.
 - c. Click **Apply**.

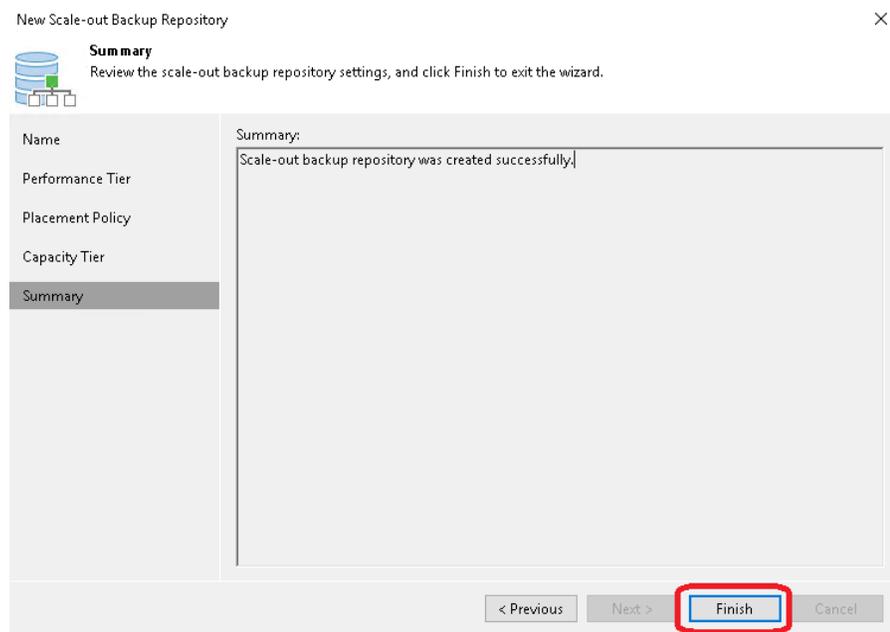
CAUTION: Do not configure Encryption in Veeam, this will cause QoreStor savings to be extremely low. Instead configure the Object Container to use encryption in QoreStor.

Figure 95: Veeam New Scale-out Backup Repository – Capacity Tier



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

Figure 96: Veeam New Scale-out Backup Repository – Summary



Instant Recovery

- Using Instant Recovery with QoreStor
- Instant Recovery with ESX
- Finalizing Instant Recovery

Using Instant Recovery with QoreStor

Veeam's Instant VM Recovery immediately restores a virtual machine (VM) back 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 by migrating it 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 NFS server.

Instant Recovery with ESX

- Instant Recovery with ESX requires completing the following procedures: Enabling Instant Recovery with ESX
- Performing Instant Recovery for ESX

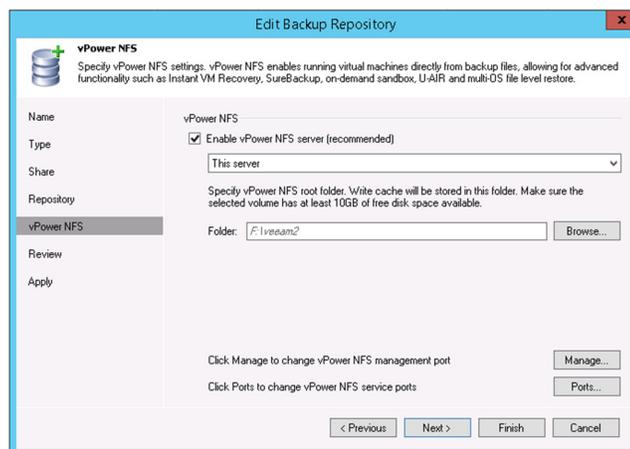
Enabling Instant Recovery with ESX

To enable Instant Recovery with ESX

Create a backup job for the required VM as described in Creating an NFS container for use with Veeam, but set the **vPower NFS Datastore** in the **vPower NFS** tab.

- 1 Select **Enable vPower NFS Server** on the **vPower NFS** tab and select the appropriate folder as the NFS Datastore.
- 2 Optionally, to configure NFS Datastore on a different Windows server, select it from the drop-down list and add the host and credentials.

Figure 97: Edit Backup Repository – vPower NFS

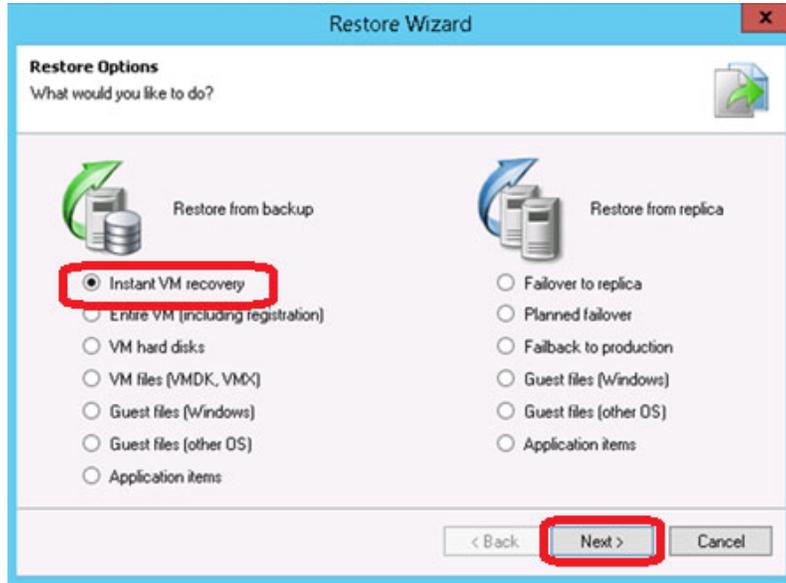


Performing Instant Recovery for ESX

To perform Instant Recovery for ESX

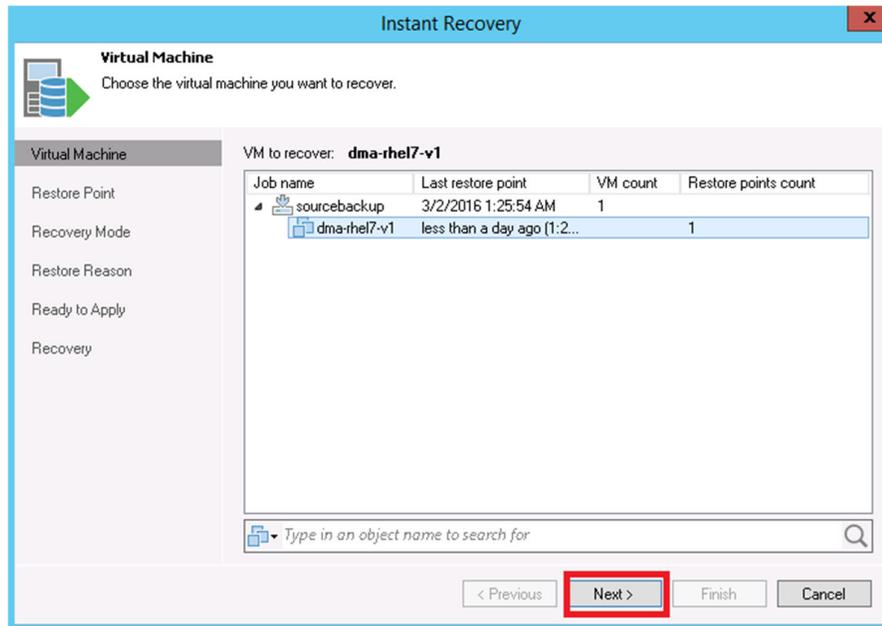
- 1 On the console for Veeam Server, click **Restore Wizard**, select **VMware**, and then select **Instant VM recovery**.

Figure 98: Restore Wizard – Restore Options



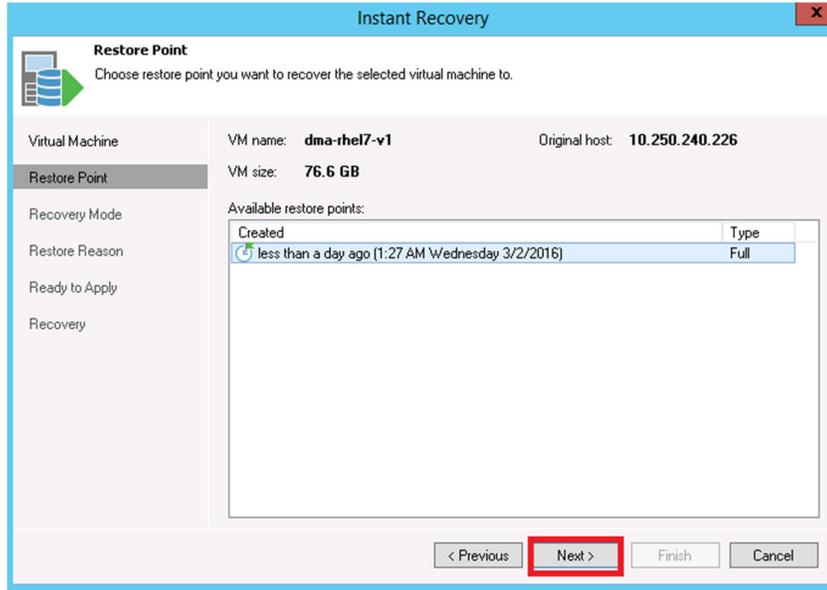
2. Select the virtual machine to be recovered and click **Next**.

Figure 99: Instant Recovery – Virtual Machine



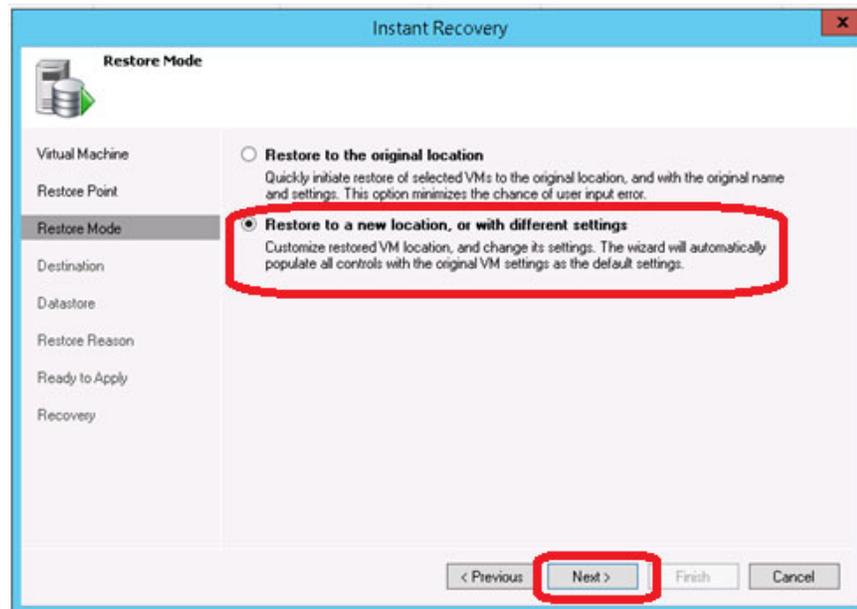
- 3 At the **Restore Point** step, select the restore point desired.

Figure 100: Instant Recovery – Restore Point



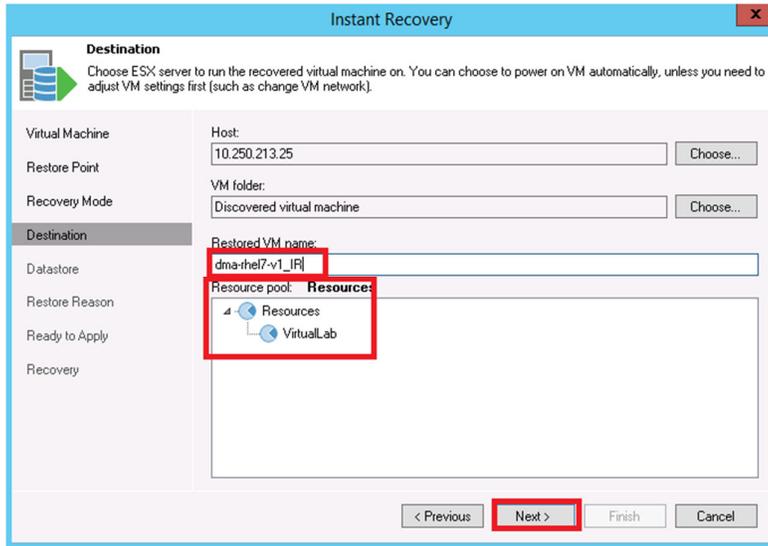
- 4 At the **Restore Mode** step, select **Restore to a new location, or with different settings**.

Figure 101: Instant Recovery – Restore Mode



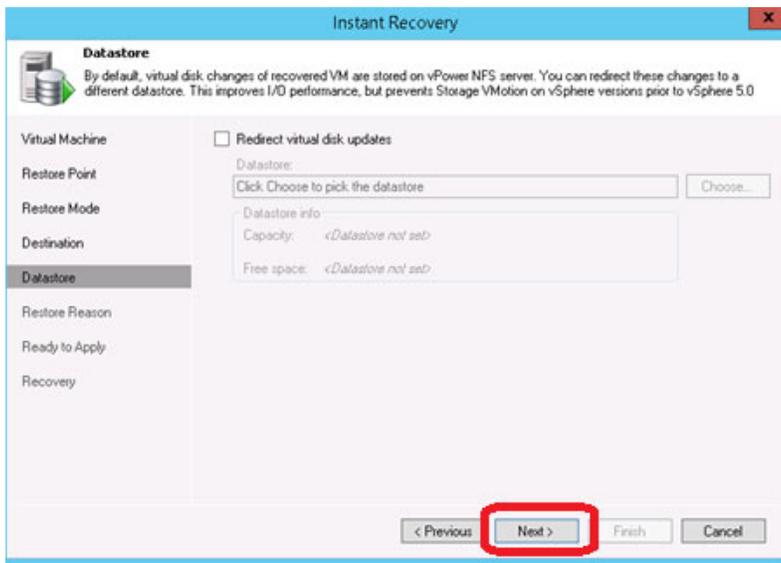
- 5 At the **Destination** step, select the ESX host on which the VM should be restored instantly. In the **Resource pool** box, select the resource pool to which the restored VM should belong.
- 6 In the **Restored VM name** field, set the desired VM name.

Figure 102: Instant Recovery - Destination



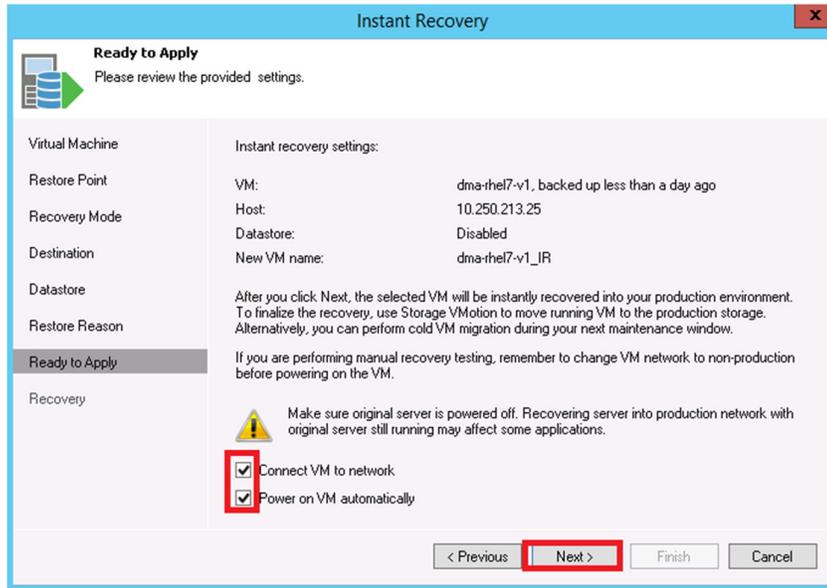
- 7 On the **Datastore** tab, leave the **Redirect virtual disk updates option** clear, which will let you use Storage vMotion to migrate the VM to production after the VM is recovered from the backup.

Figure 103: Instant Recovery - Datastore



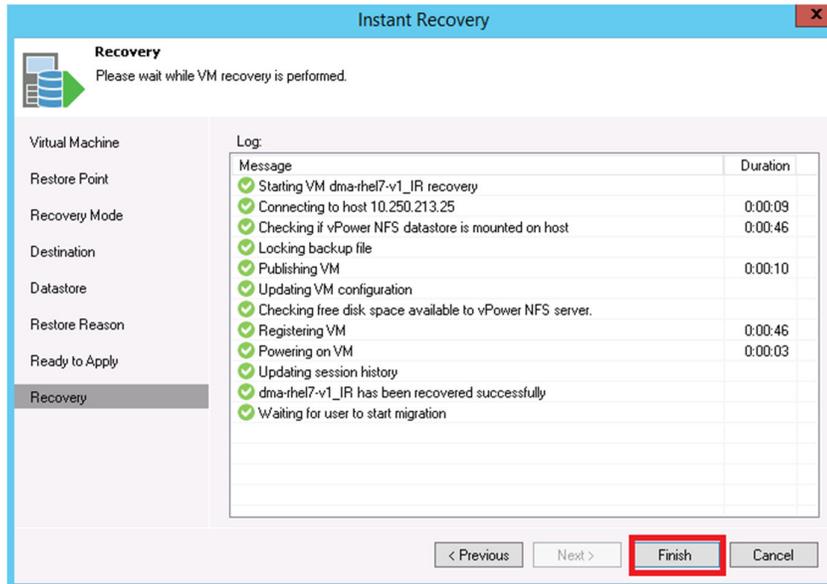
- 8 In **Ready to Apply** screen, select **Connect VM to network** and **Power on VM automatically**, and then click **Next**.

Figure 104: Instant Recovery – Ready to Apply



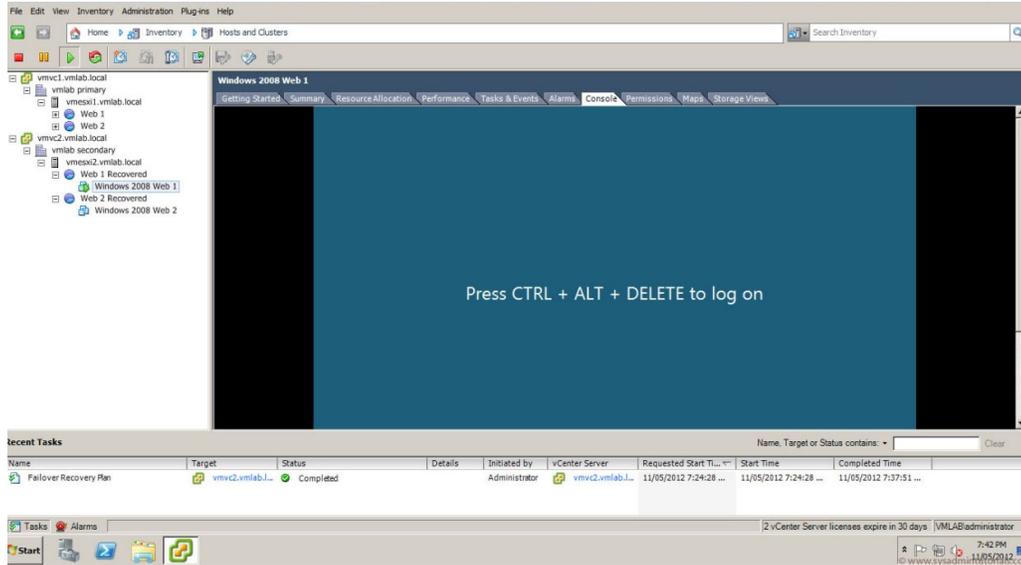
- 9 To start Instant VM Recovery, click **Finish**.

Figure 105: Instant Recovery - Recovery



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

Figure 106: vSphere client



11 In Veeam Backup & Replication, open **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.

Figure 107: Veeam Backup & Replication view



Instant Recovery with Hyper-V Server

Instant Recovery with ESX requires completing the following procedures:

- Enabling Instant Recovery with Hyper-V
- Performing Instant Recovery for Hyper-V

Enabling Instant Recovery with Hyper-V

To enable Instant Recovery with Hyper-V

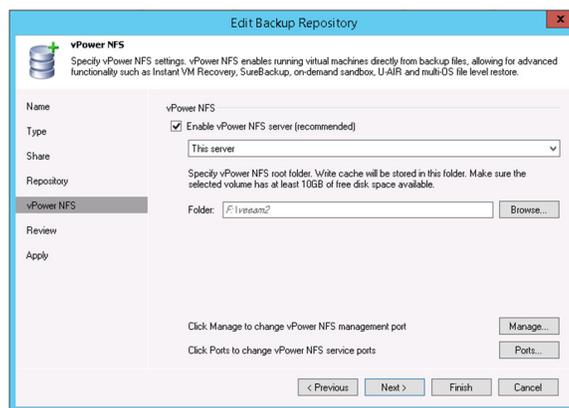
Create a backup job for the required VM as described in [Creating an NFS container for use with Veeam](#), but set the **vPower NFS Datastore** in the **vPower NFS** tab.

- 1 On the **vPower NFS** tab, select **Enable vPower NFS Server**.



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

Figure 108: Edit Backup Repository – vPower NFS

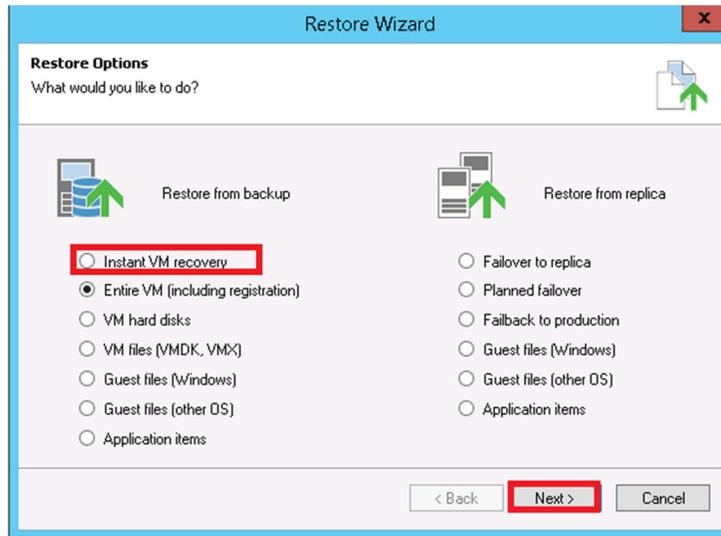


Performing Instant Recovery for Hyper-V

To perform Instant Recovery for Hyper-V

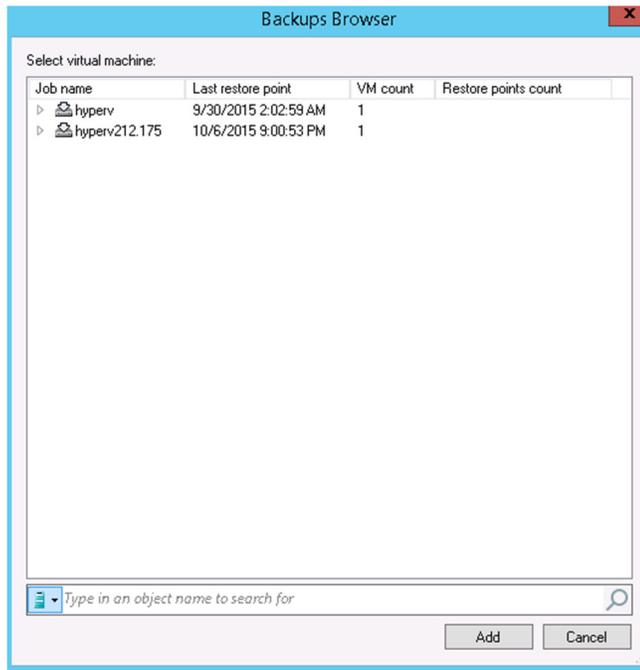
- 1 On the Veeam Backup and Replication console, click the **Restore Wizard**, select **Hyper-V**, and then select **Instant VM recovery**. Click **Next**.

Figure 109: Restore Wizard – Restore Options



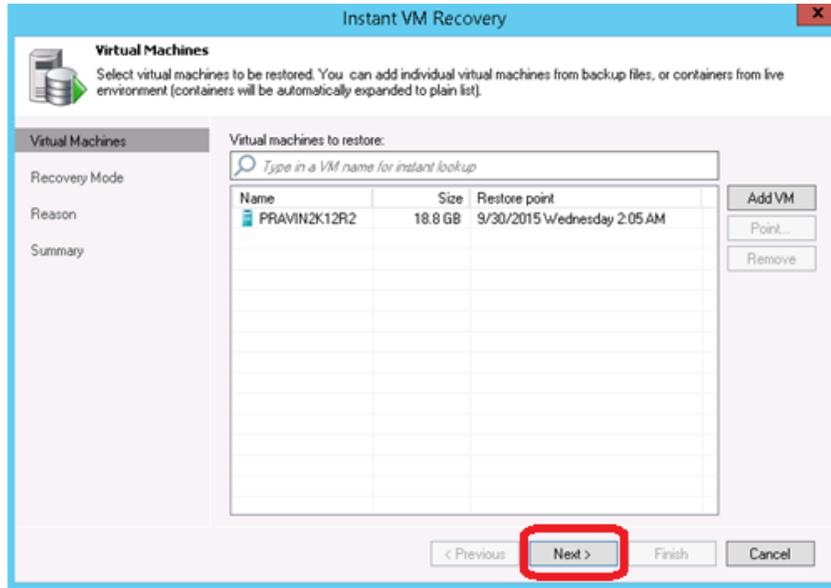
- 2 Select the virtual machine that you want to recover.

Figure 110: Backups Browser



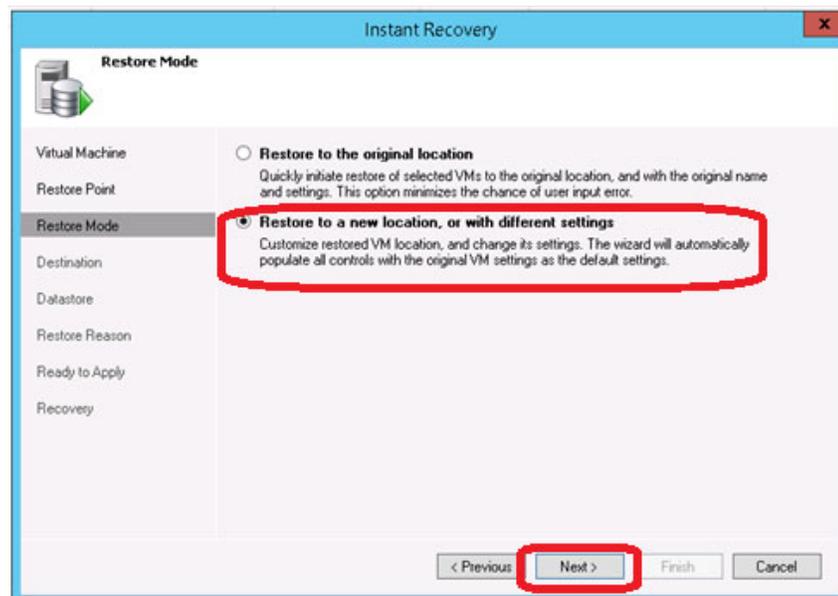
- 3 Select the restore point you want to use.

Figure 111: Instant VM Recovery – Virtual Machines



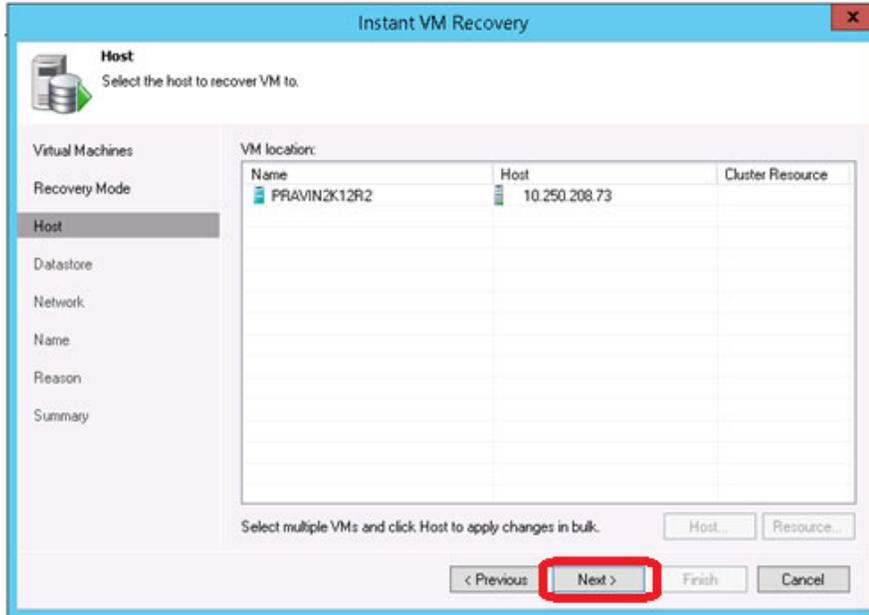
- 4 At the **Restore Mode** step, select **Restore to a new location, or with different settings**.

Figure 112: Instant Recovery – Restore Mode



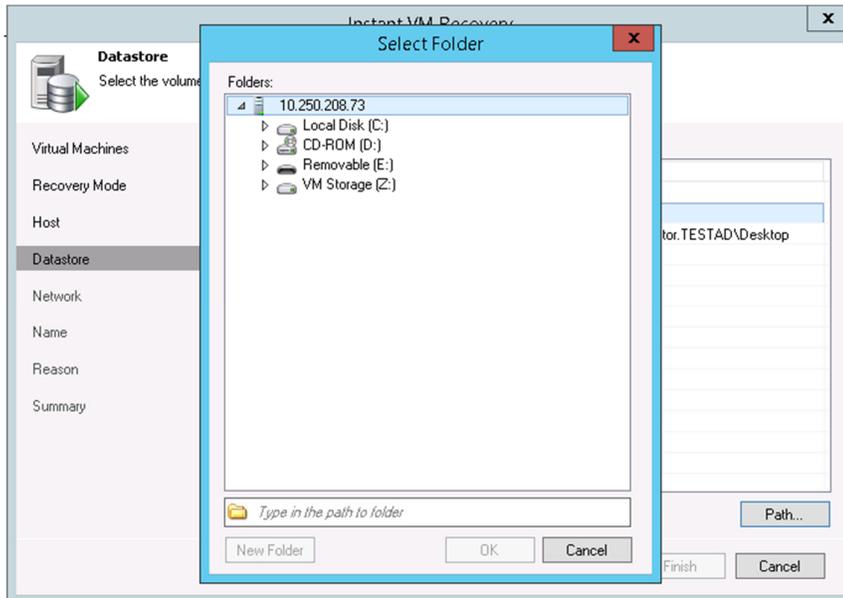
- 5 Select the Host to which your VM should be recovered.

Figure 113: Instant VM Recovery - Host



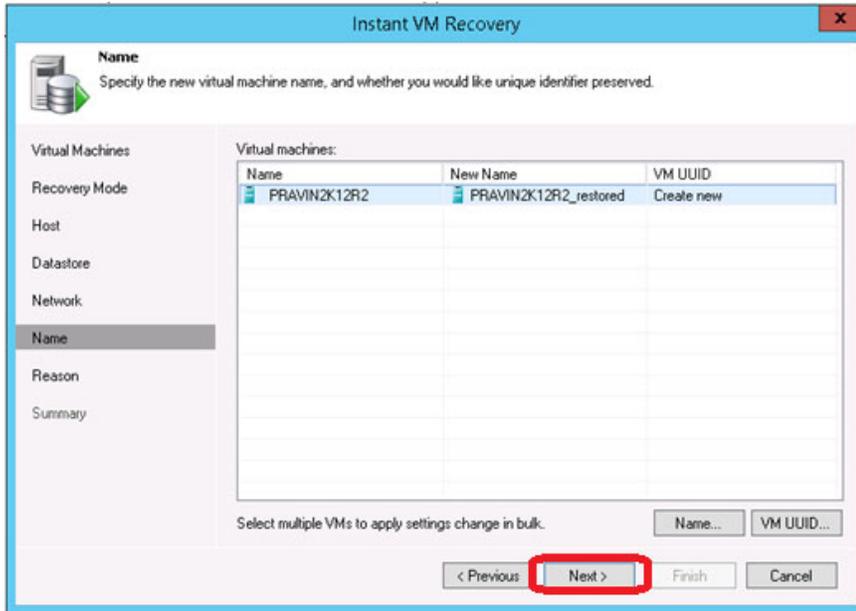
- 6 In the Datastore step, provide the location to temporarily store the cache data.

Figure 114: Select Folder window



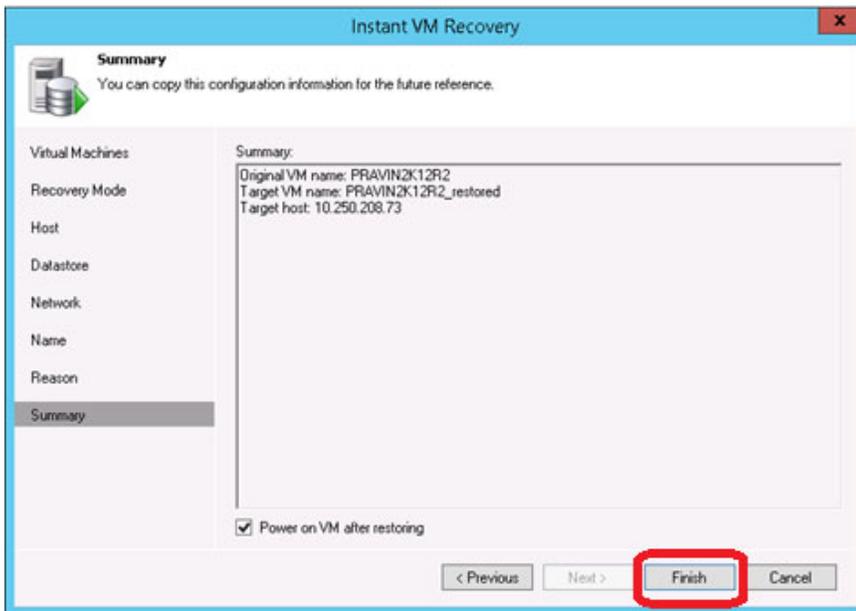
- 8 In the **Restored VM name** field, set the desired VM name.

Figure 117: Instant VM Recovery – Name



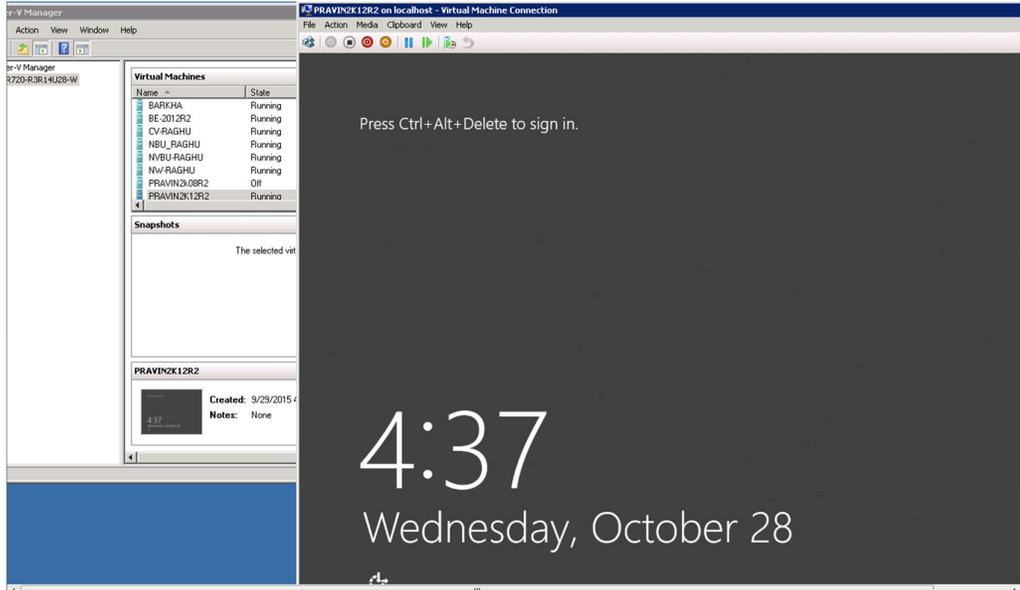
- 9 To start the recovery, click **Finish**.

Figure 118: Instant VM Recovery - Summary



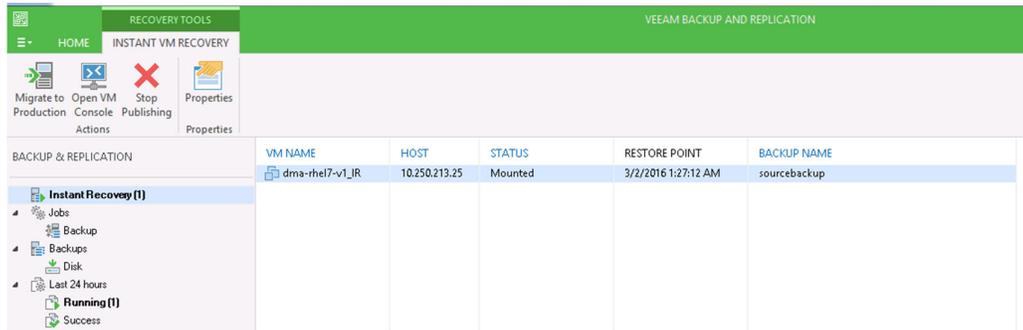
- Open Hyper-V Client and make sure that the restored VM is started on the host you selected.

Figure 119: Hyper-V Client



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

Figure 120: Veeam Backup & Replication view



Finalizing Instant Recovery

To finalize Instant Recovery, complete the following procedures:

1. Migrating the restored VM to production
2. Terminating the Instant VM Recovery session

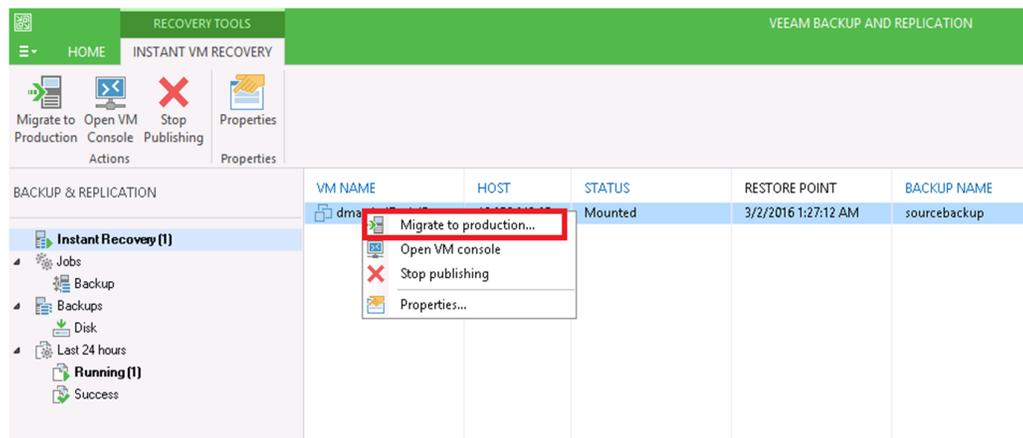
Migrating the restored VM to production

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 to production

- 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**.

Figure 121: Veeam migration 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 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

- Requirements of Fast Clone
- Configuring a new Fast Clone repository
- Reconfiguring an existing QoreStor repository for Fast Clone

Understanding Fast Clone

Fast clone allows for synthetic full backups of Hyper-V systems or Data Copy jobs with VM's on ReFS file system with less read performance impact to 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 exiting 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 VM's 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 and cover additional steps to reconfigure existing QoreStor repositories in Veeam.

To configure a new Fast Clone repository

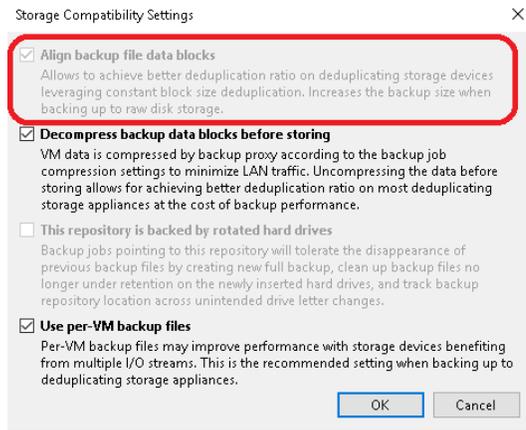
- 1 Open a command prompt to the QoreStore user and elevate to root.
- 2 Edit the `/etc/oca/customer.env` file using the following command:


```
vi /etc/oca/customer.env
```
- 3 Add the line `export OCA_ENABLE_SYNTHETIC_FULL=1` and save it.
- 4 Restart QoreStor services using the following command:


```
systemctl restart ocards
```
- 5 Wait for the QoreStor services to restart and for the system to become operational again.
- 6 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 you plan to use by following the steps in Installing Rapid CIFS on a Veeam Windows Proxy.

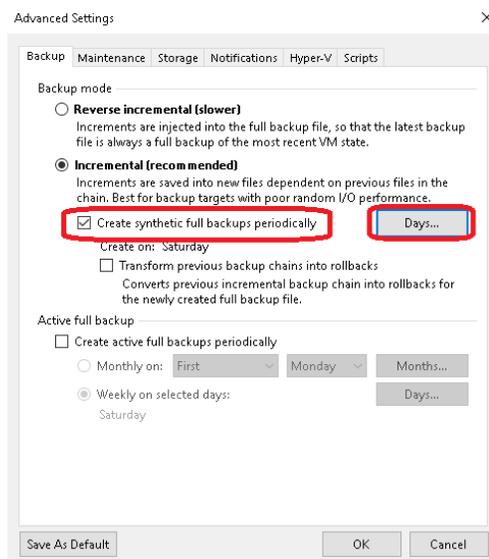
- 7 Create a new CIFS container and add it to Veeam as a repository by following the **Creating a CIFS container for use with Veeam** and **Adding the QoreStor CIFS container as a repository in Veeam** sections of this guide. Ensure the **Align backup file data blocks** option is checked when adding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

Figure 122: Storage Compatibility Settings



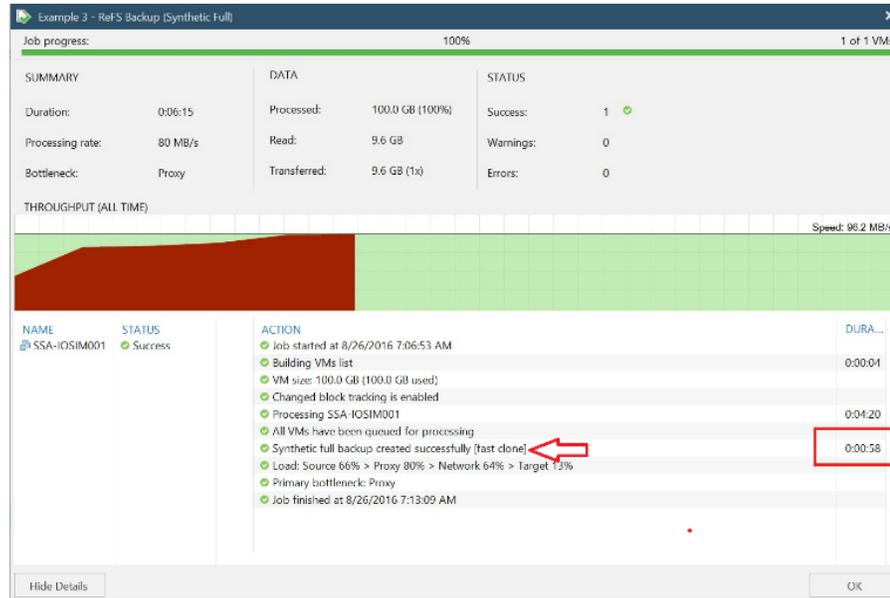
- 8 Create a new Hyper-V backup or Data Copy job following the **Creating a backup job with the QoreStor system as target** section of this guide ensuring to use the Synthetic full option in the job settings.

Figure 123: Advanced Settings



With the next Synthetic Full backup, you should see Fast Clone referenced in the job details.

Figure 124: Synthetic Full job progress



Reconfiguring an existing QoreStor repository for Fast Clone

This section includes additional steps needed to get an existing QoreStor repository recognized as supporting Fast Clone by Veeam. To achieve recognition, the existing repository will need to be removed and re-added to Veeam, which involves pointing existing jobs to other repositories or deleting them outright.

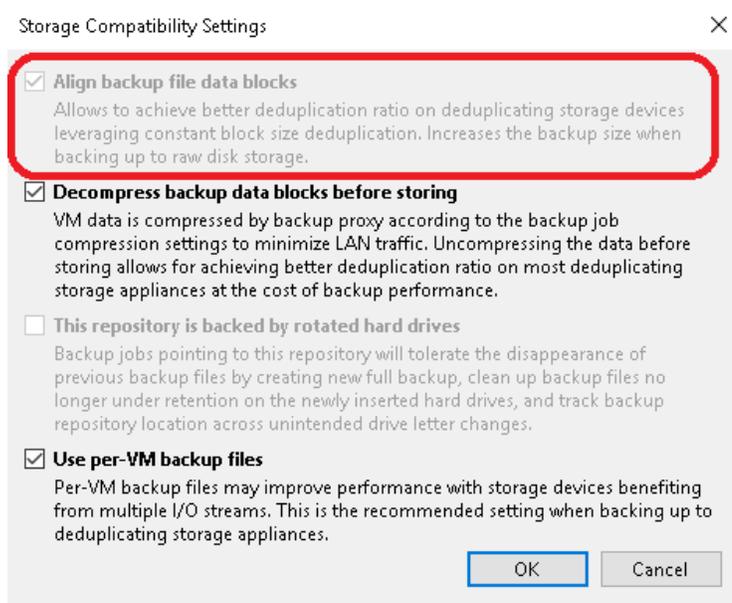
Warning: This is an advanced operation and should only be attempted by a customer comfortable with the Veeam product. Quest recommends creating a new Repository in the same Storage Group and leaving your existing repository in place rather than following these steps.

To reconfigure an existing QoreStor repository for Fast Clone

- 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, which in 9.5 U4 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, which 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. If Veeam recognizes Fast Clone support, the **Align backup file data blocks** option should be automatically checked and greyed out. If not, double check all previous steps.

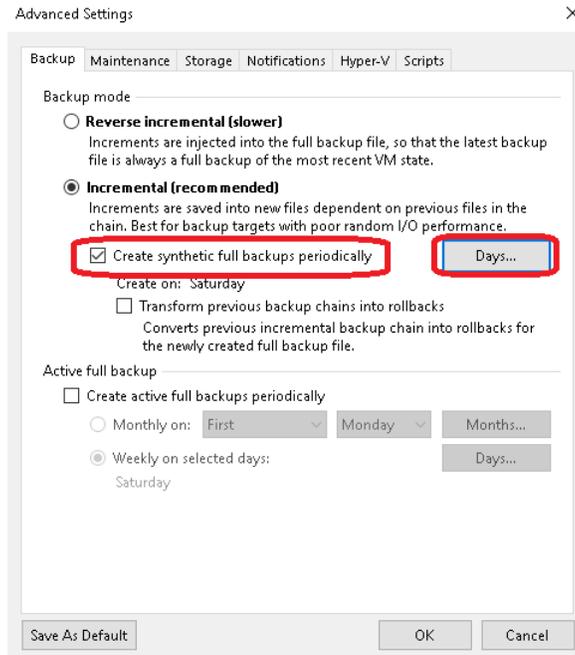
Figure 125: Storage Compatibility Settings



- 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.
 - 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.

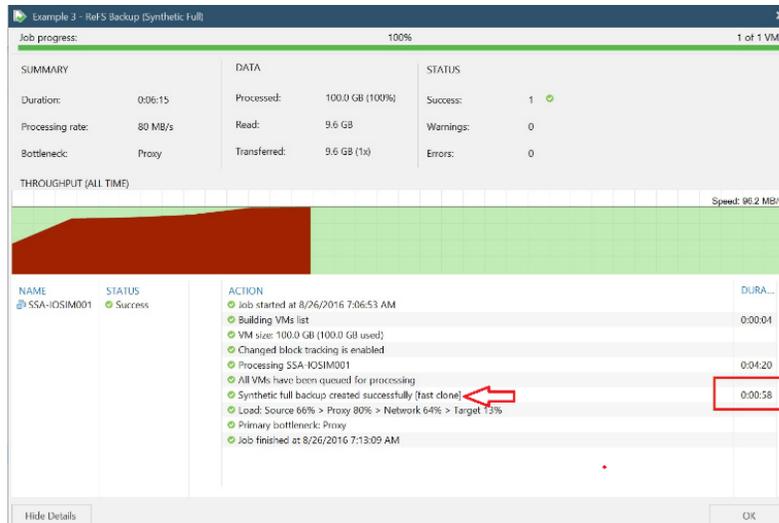
- Edit your cloned jobs to use the newly re-added repository. Ensure the Synthetic feature is selected in the job advanced options for every cloned job.

Figure 126: Advanced Settings



The next Synthetic Full you should see Fast Clone referenced in the job details.

Figure 127: Synthetic Full progress with Fast Clone



Performance Tier

- Understanding Performance Tier
- Setting up Performance Tier with QoreStor
- Optimizing Performance Tier via Sync Always option

Understanding Performance Tier

A Performance Tier allows you to define a set of faster disks as a Storage Group and created 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 during 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 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.

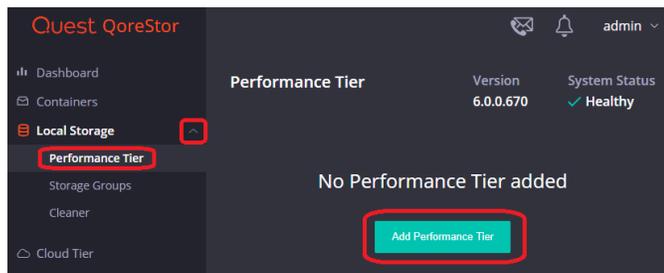
Setting up Performance Tier with QoreStor

This section does not cover adding a device, creating a partition, creating a XFS filesystem, or defining a mount point in detail. For information about those procedures, see the *QoreStor Installation Guide*.

To set up Performance Tier with QoreStor

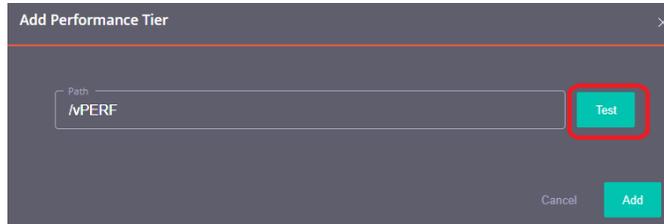
- 1 Cable and add the disks to the OS level. After the OS sees QoreStor as a device, create an aligned partition, an XFS file system, and a mount point defined in fstab that includes mount option requirements defined in the *QoreStor Installation Guide*.
- 2 After you add a file system path to the high-performance storage, 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**.

Figure 128: QoreStor Performance Tier page



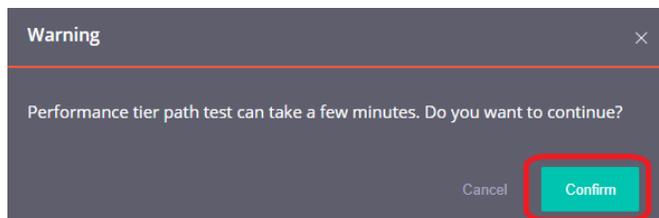
- 3 Enter the Performance Tier mount path and click **Test**.

Figure 129: Add Performance Tier test



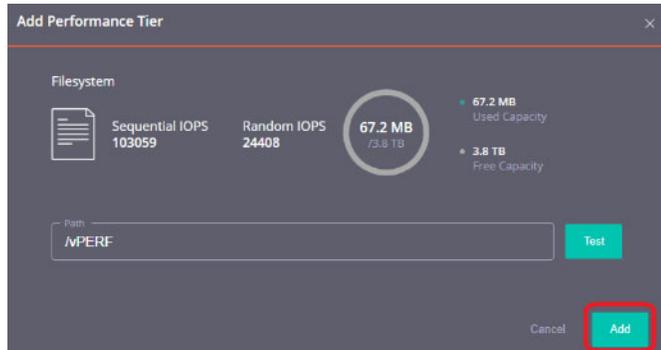
- 4 Click **Confirm**.

Figure 130: Performance Tier warning



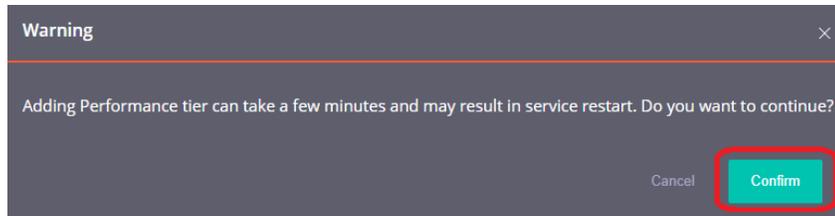
- 5 If the path gets the expected performance, click **Add**.

Figure 131: Add Performance Tier add



- 6 Click **Confirm** to finish adding the Performance Tier and restart QoreStor services.

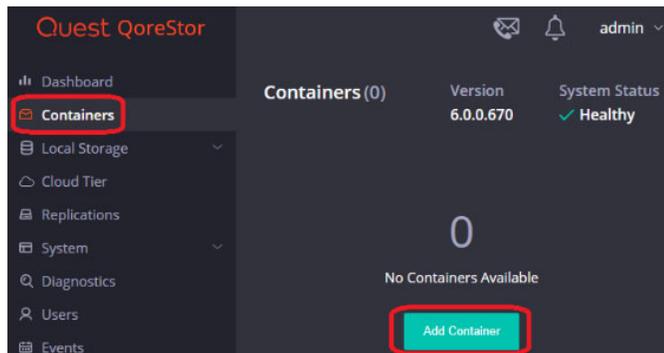
Figure 132: Adding Performance Tier warning



After adding the Performance Tier, you will be logged out. When you log in, the Performance Tier tab will now list a dashboard for the Performance Tier.

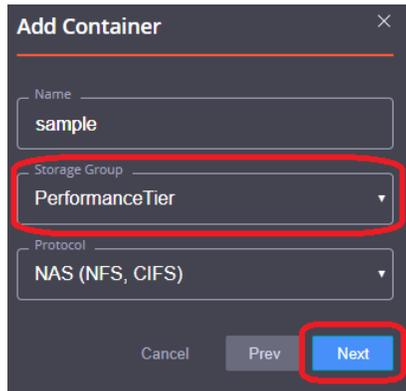
- 7 Navigate to the Containers tab and click **Add Container**.

Figure 133: QoreStor Containers page



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

Figure 134: Add Container window – Storage Group



The screenshot shows a dark-themed dialog box titled "Add Container" with a close button (X) in the top right corner. It contains three input fields: "Name" with the text "sample", "Storage Group" with a dropdown menu showing "Performance Tier", and "Protocol" with a dropdown menu showing "NAS (NFS, CIFS)". At the bottom, there are three buttons: "Cancel", "Prev", and "Next". The "Storage Group" dropdown and the "Next" button are both circled in red.

- 9 To finish configuring your Performance Tier container, follow the remaining steps in the Creating a CIFS container for use with Veeam and Adding the QoreStor CIFS container as a repository in Veeam sections of this guide.

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 option to Yes can decrease performance but adds more resiliency to writes in case of interruption of the QoreStor system before writes sync to disk. We do not recommend this option in cases where performance is a key factor.

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
- 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 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 cloud.



Warning: After you configure a container as Cloud Tier, the only way to remove it is to delete the container or contact Support to fully restore all data blocks from the cloud, which may involve a read cost from the cloud provider.



Warning: Consider your Veeam Job configuration and policy configuration when deploying Cloud Tier. Failure to do so could result in unexpected charged from the cloud provider or failing backup jobs. Please read this section in its entirety as well as checking 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 offers the following options:

- **Idle Time Before Cloud Migration** – Replicates stable data blocks that are idle for more than the selected amount of time to the cloud and locates the complete data blocks both On-Premises and in the cloud. All restores come from the On-Premises data block and do not induce any cost. Any attempted modification of files after this idle time results in Access Denied errors. For more information, see [Setting up Cloud Tier](#).
- **On-Prem Retention Age** – After the selected amount of time passes, data blocks that have replicated to the cloud are removed from On-Premises storage, after which any data reads, such as restore or [synthetic full backups](#), are from the cloud provider. This option can be slower and induce cost from the provider.
- **Folder Paths** – Allows for including or excluding specific paths from cloud tiering replication. This feature should not be necessary with Veeam.
- **File Extensions** – Allows for including or excluding specific file types from cloud tiering replication. Usually, this feature should not be necessary with Veeam.

In most cases with Veeam, you should only need to consider the **Idle Time Before Cloud Migration** and **On-Prem Retention Age** options.



NOTE: **Idle time** is especially important to consider with the Forever Forward Incremental and Forward Incremental with Synthetic Full Backups workflows.

The Policy Manager offers the following workflows:

- **Forever Forward Incremental** – Quest recommends against using Forever Forward Incremental jobs with Cloud Tiering. In this workflow, a full backup is taken initially and kept; every backup after this is an incremental backup. 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 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 charges 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 Backup, Veeam will read from the most recent Full Backup, as well as every incremental after it. All this data will be written into a new Full Backup file.

i IMPORTANT: Be sure that your On-Prem Retention Age setting is longer than your synthetic schedule. If it is not, then 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 an incremental, which is then injected directly into the full. After the injection, an incremental file is left with all the data removed from the full. These files are okay to tier to 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 must be longer than your scheduled Incremental Backup frequency. Failure to do so results in Access Denied errors.

Setting up Cloud Tier

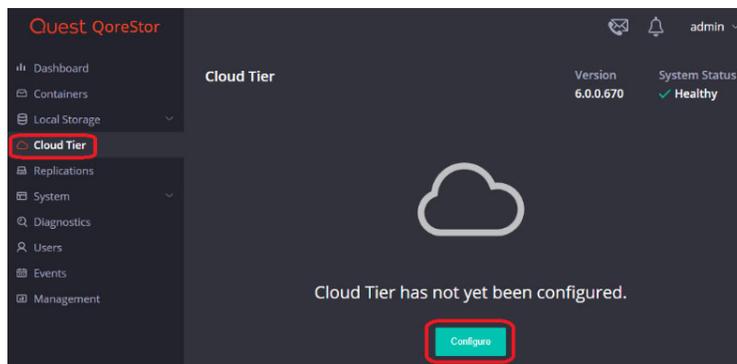
Before setting up Cloud Tier, gather the following information from your cloud provider:

- If using Azure, you will need your Connection String, which you can find on your Azure portal under your blob storage account.
- If using AWS, Wasabi or an S3 Compatible cloud provider, have your Access Key, Secret Key, Region, and Endpoint setting (if using a cloud emulator). These settings are available on your AWS console or from your cloud provider.

To set up Cloud Tier

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

Figure 135: QoreStor Cloud Tier page



- 2 For Azure, enter your Azure Container name, which 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. Click **Configure**.

i **NOTE:** The Azure Container name must be lower case and some symbols are not allowed. This is a limitation of Azure.

Figure 136: Configure Cloud Tier window

The screenshot shows the 'Configure Cloud Tier' window with the following fields and controls:

- Cloud Provider:** Azure Blob
- Need Help?** (link)
- Azure Container:** blobcontainer
- Connection String:** (empty text area)
- Cloud Tier Encryption:**
 - Passphrase:** (masked with dots)
 - Confirm Passphrase:** (masked with dots)
- Buttons:** Close, Configure

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



NOTE: The S3 Bucket name need to be lower case and some symbols are not allowed. This is a limitation of S3.

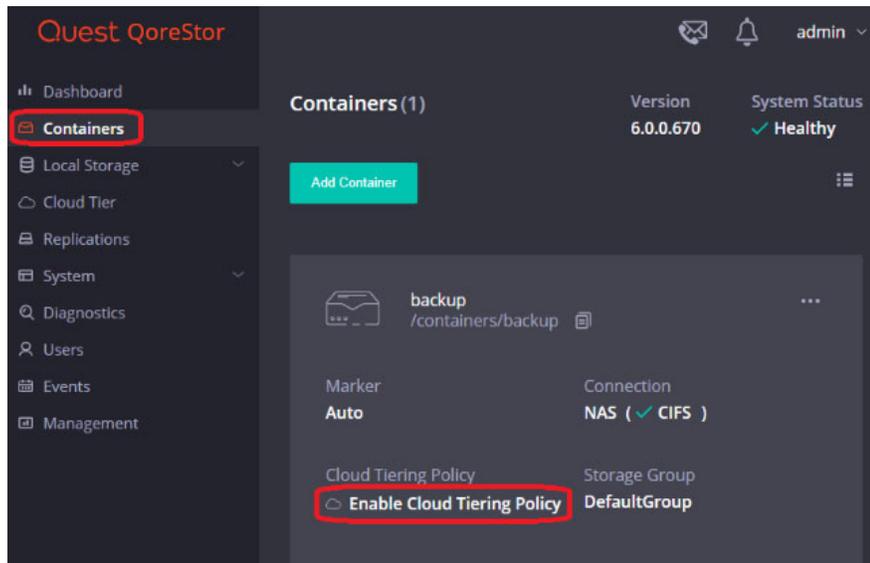
Figure 137: Configure Cloud Tier window

Cloud Tier should show as configured and the **Cloud Tier** tab populates with statistics.

- 4 Enable the Cloud Tiering Policy on individual containers.
- 5 Select the **Containers** tab and find or create a container, and then click **Enable Cloud Tiering Policy** on this container.

Warning: After you configure the container as Cloud Tier, the only way to remove it is 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.

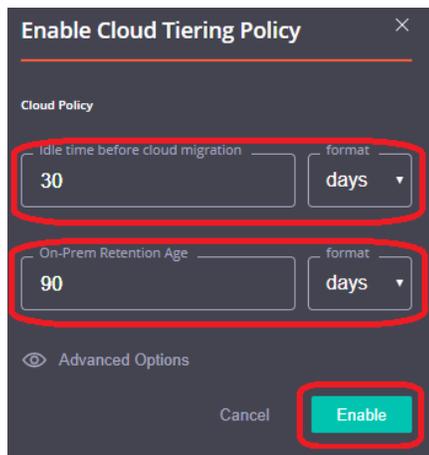
Figure 138: QoreStor Containers page



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

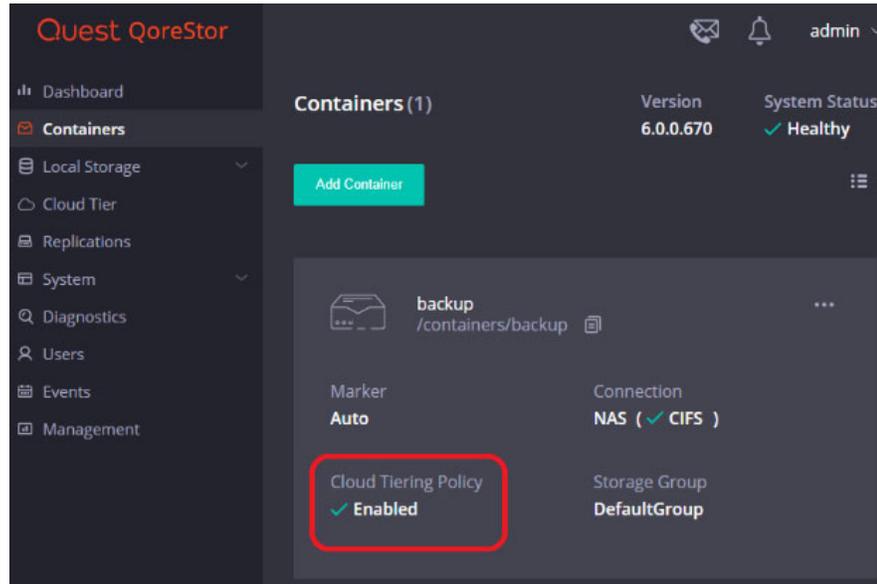
Warning: Before defining idle time and retention age, see the Important Considerations for Cloud Tier with Veeam.

Figure 139: Enable Cloud Tiering Policy window



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

Figure 140: QoreStor Containers page with Cloud Tiering Policy enabled



Archive Tier

QoreStor's Archive Tier feature lets you quickly and easily archive data to long-term Amazon S3 Glacier or Amazon S3 Glacier Deep Archive storage.

Important Considerations for Archive Tier with Veeam

When using Veeam and a supported protocol (Object container(S3)), you can write files to a QoreStor container and migrate them to your archive tier according to easily defined policies. QoreStor provides a policy engine that lets you 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 a 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

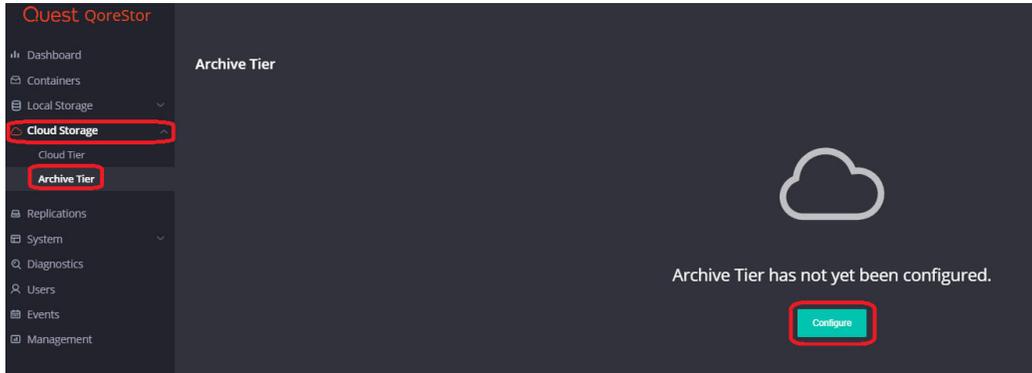
Setting up Archive Tier

Archive Tier is a feature that allows a QoreStor system to tier deduplicated blocks of files to an AWS glacier or deep archive using 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 premises in QoreStor, how long it's available both on premises and in archive simultaneously, and at what point is it only available in cloud. Archive tier restores are more difficult, so careful consideration should be given to how long the data should be available on prem before configuring archive tier.

To set up Archive Tier

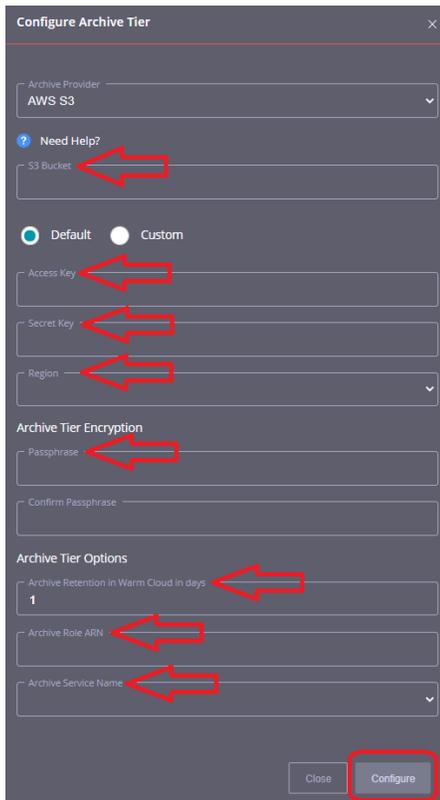
- 1 Open the QoreStor UI, expand the **Cloud Storage** section and select the **Archive Tier** page and click **Configure**.

Figure 141: QoreStor Archive Tier page



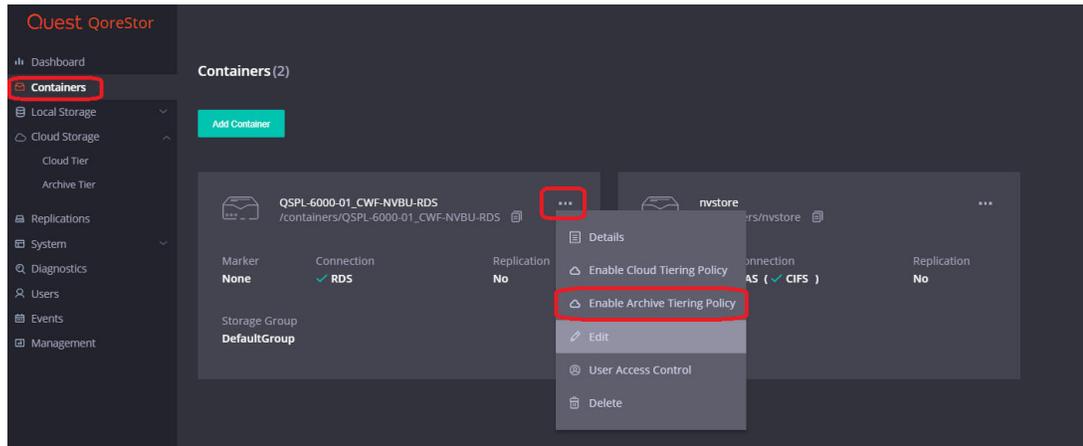
- 2 Enter the information from your AWS account including the **access key, secret, correct region, ARN role** and select an **Archive Service Name**. The **S3 bucket name** will be created and is character limited by the provider. Make sure to keep your **passphrase**; without this, the data is not recoverable in a Disaster Recovery scenario. Click **Configure**.

Figure 142: Configure Archive Tier window



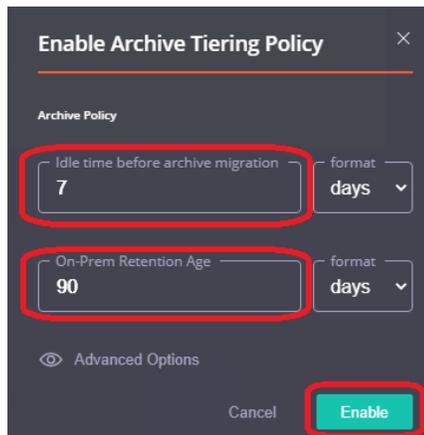
- To add an archive tiering policy to a specific container, navigate to the **Containers** page, select the **ellipsis** in the top right corner of the specific container, and click **Enabled Cloud Tiering Policy**.

Figure 143: QoreStor Containers page actions menu



- In the Enable Archive Tiering Policy window, define the policy by entering the **Idle time before archive migration**, which specifies the number of hours or days datablocks must be kept idle before being sent to the cloud; then enter the **On-Prem Retention age**, which specifies the number of hours or days files will be kept locally after they are sent to archive. Finally click **Enable**.

Figure 144: Enable Archive Tiering Policy window



Setting up the QoreStor system cleaner

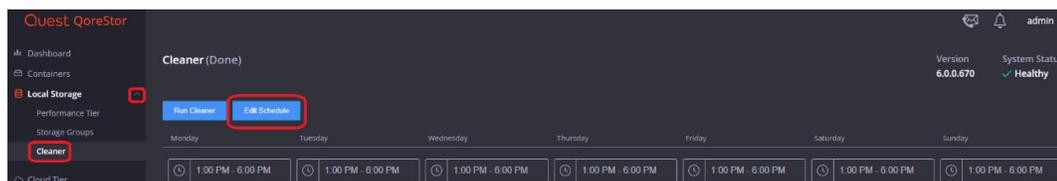
Quest recommends performing scheduled disk space reclamation operations as a method for recovering disk space from system containers in which files were deleted as a result of deduplication.

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 completed. For more information, see the *QoreStor Series Cleaner Best Practices* white paper.

To set up the QoreStor system cleaner

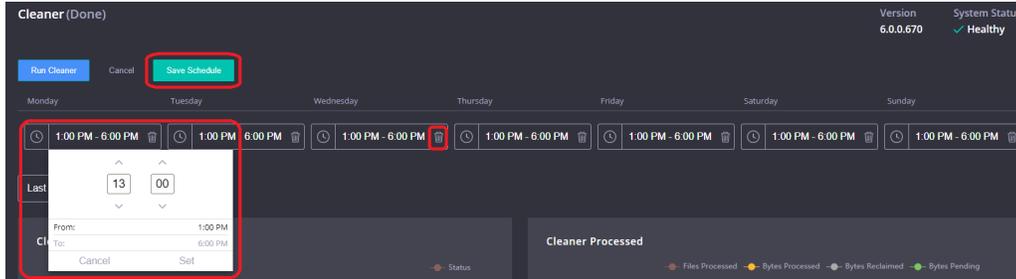
- 1 In the QoreStor system GUI, expand the **Local Storage** tab, click **Cleaner**, and then click **Edit Schedule**.

Figure 145: QoreStor Cleaner page



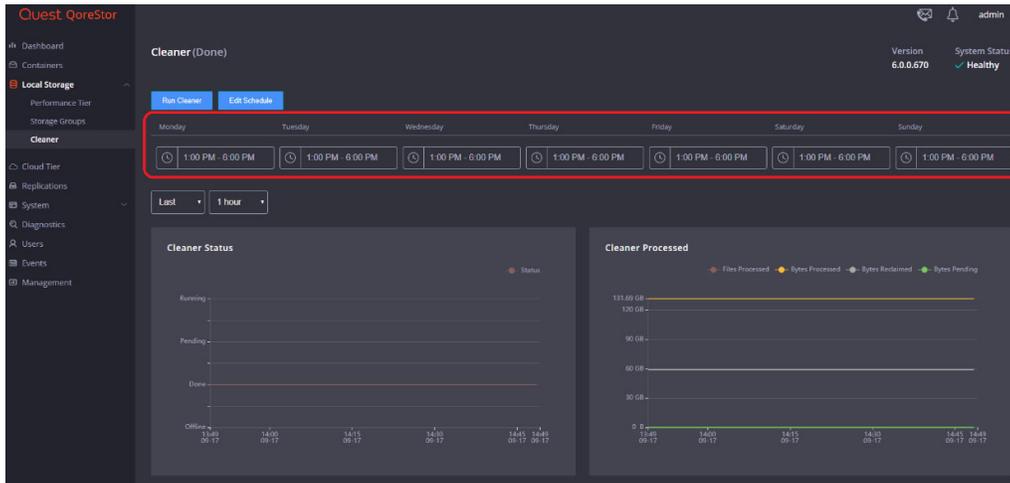
- 2 Define the schedule and click **Save Schedule**.

Figure 146: QoreStor cleaner schedule



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

Figure 147: QoreStor Cleaner page with schedule set



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.

Figure 148: QoreStor dashboard

