Veeam Backup for Hyper-V Operational Guide

Volume 2
Based on Version 11a
Focused on Microsoft Hyper-V

By:

Dave Kawula Cristal Kawula Emile Cabot Cary Sun

PUBLISHED BY

MVPDays Publishing http://www.mvpdays.com

Copyright © 2023 by MVPDays Publishing

All rights reserved. No part of this book may be reproduced or transmitted in any form or by any means without the publisher's prior written permission.

ISBN: TBD

Warning and Disclaimer

Every effort has been made to make this manual as complete and accurate as possible, but no warranty or fitness is implied. The information provided is on an "as is" basis. The authors and the publisher shall have neither liability nor responsibility to any person or entity concerning any loss or damages arising from the information contained in this book.

Feedback Information

We'd like to hear from you! If you have any comments about how we could improve the quality of this book, please don't hesitate to contact us by visiting www.checkyourlogs.net or emailing feedback@mvpdays.com.

Forward

Here is another book by Dave, Cristal, Cary and Emile; what a significant milestone!

Ask yourself one question: Why? There are so many technologies, but why do we use what we use? Why do we do what we do? The answer is how. It's how we use something. I like to explain sometimes compliance in this way. No product or technology is inherently compliant. It's how it is implemented and how it is audited. The same goes for technology implementations; it's about how we use them. The how is the why.

Operations are still cool. There are so many razzle-dazzle job titles and buzzwords in the market today, but in the end, Operations are Operations. DevOps, PlatformOps, SRE (Sire Reliability Engineer), Platform Engineering... I do not need to go on, but no technology will take care of itself across all disciplines. How it is used, implemented, monitored, etc., matters today. Technology still needs humans and their knowledge.

Expert advice is the difference. We all learn from each other. When taking on the next new challenge, where does one go first? We look for resources to consume. Blogs, books like this, and social profiles; the established experts are the trusted advisors in the technology space. Call it community, social sharing, or what you want; we all find ourselves going to the go-to experts of a particular space.

Above and Beyond. What Dave, Cristal, Cary and Emile put forth in this book is outstanding in their practicing advice for technology. They could easily focus on their professional responsibilities and keep them narrow. But writing a book is hard work! Editing a book is hard work! I've not discussed this with them, but I'm sure they aren't doing it for the money of writing a book. They write this book because they go above and beyond, share, and care.

I'm sure you will enjoy this book, and a big congratulations on this book, Dave, Cristal, Cary and Emile.

Best Regards,

Rick W. Vanover Microsoft MVP, VMware vExpert, Cisco Champion

Senior Director, Product Strategy - Veeam Software

Twitter: @RickVanover

About the Authors

Dave Kawula – Microsoft MVP

Dave is a Microsoft Most Valuable Professional (MVP) with over 20 years of experience in the IT industry. His background includes data communications networks within multi-server environments, and he has led architecture teams for virtualization, System Centers, Exchange, Active Directory, and Internet gateways. Very active within the Microsoft technical and consulting teams, Dave has provided deep-dive technical knowledge and subject matter expertise on various System centers and operating system topics.

Dave is well-known as an evangelist for Microsoft, 1E, and Veeam technologies. Locating Dave is easy as he speaks at conferences and sessions each year, including TechEd, Ignite, MVP Days Community Roadshow, and VeeamOn.

Recently Dave has been honoured to take on the role of Conference Co-Chair of TechMentor and Cyber Security & Ransomware Live with fellow MVP Sami Laiho. The lineup of speakers and attendees attending this conference over the past 20 years is fantastic. Come down to Redmond or Orlando in 2018 and meet him in person. Check out his speaking site at https://sessionize.com/dave-kawula/

He recently tied for 1st place out of 1800 speakers at the Microsoft Ignite Conference in Orlando.

As the founder and Managing Principal Consultant at TriCon Elite Consulting, Dave is a leading technology expert for local customers and large international enterprises, providing optimal guidance and methodologies to achieve and maintain an efficient infrastructure.

BLOG: www.checkyourlogs.net

Twitter: @DaveKawula





Cristal Kawula – Microsoft MVP

Cristal Kawula co-founded MVPDays Community Roadshow and #MVPHour live Twitter Chat. She was also a Technical Advisory board member and the President of TriCon Elite Consulting. Cristal is the only 2nd Woman worldwide to receive the prestigious Veeam Vanguard award.

Cristal speaks at Microsoft Ignite, MVPDays, and other local user groups. In addition, she has been instrumental in founding MVPDays Publishing and has helped author over 25 + books.

At conferences like Microsoft Ignite, she has led community meetups on Women in IT, Parenting in IT, Diversity in Tech, and becoming a Community Rockstar.

BLOG: http://www.checkyourlogs.net

Twitter: @supercristal1



Cary Sun – Microsoft MVP

Cary Sun is a CISCO CERTIFIED INTERNETWORK EXPERT (CCIE No.4531) and MCSE, MCIPT, Citrix CCA with over twenty years in the planning, design, and implementation of network technologies and Management and system integration. Background includes hands-on experience with multiplatform, all LAN/WAN topologies, network administration, E-mail and Internet systems, security products, PCs and Servers environment. Expertise is analyzing users' needs and coordinating system designs from concept through implementation. Exceptional analysis, organization, communication, and interpersonal skills. Demonstrated ability to work independently or as an integral part of a team to achieve objectives and goals. Specialties: CCIE /CCNA / MCSE / MCITP / MCTS / MCSA / Solution Expert / CCA

Cary is a very active blogger at checkyourlogs.net and is permanently available online for questions from the community. His passion for technology is contagious, improving everyone around him at what they do.

Blog:http://www.checkyourlogs.net

Twitter:@SifuSun





Emile Cabot – Microsoft MVP

Emile started in the industry during the mid-90s working at an ISP and designing celebrity websites. He has a solid operational background specializing in Systems Management and collaboration solutions. In addition, he has spent many years performing infrastructure analyses and solution implementations for organizations ranging from 20 to over 200,000 employees. Coupling his wealth of experience with a small partner network, Emile works very closely with TriCon Elite, 1E, and Veeam to deliver low-cost solutions with minimal infrastructure requirements.

He actively volunteers as a member of the Canadian Ski Patrol, providing over 250 hours each year for first aid services and public education at Castle Mountain Resort and in the community.

BLOG: http://www.checkyourlogs.net

Twitter: @ecabot



Contents

Forward	iii
About the Authors	vi
Dave Kawula – Microsoft MVP	vi
Cristal Kawula – Microsoft MVP	vii
Cary Sun – Microsoft MVP	viii
Emile Cabot – Microsoft MVP	ix
Contents	x
Introduction	13
Sample Files	13
Additional Resources	13
Chapter 1	14
Cloud Repositories	14
Azure Blob Object Storage	14
Adding Microsoft Azure Blob Object Storage Repositories	15
Adding Microsoft Azure Archive Storage Repositories without Azure Helpe appliance	
Adding Microsoft Azure Archive Storage Repositories with Azure Helper ap	•
Adding Local Directory and Azure Blob Object Storage as Scale-out Repositional Archive Tier	-

Adding Local Directory and Azure Blob Object Storage as Scale-out Repository with Archive Tier	
Chapter 2	170
Cloud Backup and Backup Copy	170
Creating a Backup job using Azure Blob repositories as Cloud Redundant Data	171
Creating a Backup Copy Jobs offload backups to Azure Blob without using the Arc	
Creating a Backup Copy Jobs offload backups to Azure Blob using the Archive Tie	
Chapter 3	243
Restore to Microsoft Azure	243
Restore On-premises VM to Microsoft Azure	243
Chapter 5	272
SureBackup	272
Creating Application Group	273
Configuring Basic Single-Host Virtual Labs	284
Configuring Advanced Single-Host Virtual Labs	294
Creating a SureBackup Job	311
Chapter 6	329
Creating On-Demand Sandbox	329
Configuring On-Demand Sandbox	330
Chapter 7	373
Reporting	373
Real-Time Job Statistics Report from Veeam Backup & Replication Console	374

	Job Session History Report from Backup and Replication Console	378
	Job and Job Session Report from Veeam Backup and Replication Console	382
	Generate Backup Reports from Veeam ONE	386
С	hapter 8	418
Jo	oin us at MVPDays and meet great MVPs like this in person	418
	Live Presentations	418
	Video Training	418
	Live Instructor-led Classes	418
	Consulting Services	419

Introduction

This book aims to showcase the fantastic expertise of our guest speakers of MVPDays Online. They have so much passion, expertise, and expert knowledge that it only seemed fitting to write it down in a book.

This book aims to show how to be operationally proficient using Veeam Backup and Replication, Veeam One and various other Veeam products and tools. We hope you find immense value in reviewing this guide and encourage you to share your operational knowledge and skills with others in the community.

Sample Files

All sample files for this book can be downloaded from www.checkyourlogs.net and www.github.com/mvpdays

Additional Resources

In addition to all the tips and tricks provided in this book, you can find extra resources like articles and video recordings on our blog http://www.checkyourlogs.net

Chapter 1

Cloud Repositories

Veeam Backup and Replication supports the object storage repositories listed below.

- Amazon S3, Amazon S3 Glacier and AWS Snowball Edge
- S3 compatible
- Google Cloud
- IBM Cloud
- Microsoft Azure Blob, Azure Archive Storage and Azure Data Box

This chapter will walk through the initial configurations of the Microsoft object storage repository. These include:

- Microsoft Azure Blob Object Storage Repositories
- Microsoft Azure Archive Storage Repositories without Azure Helper appliance
- Microsoft Azure Archive Storage Repositories with Azure Helper appliance
- Local Directory and Azure Blob Object Storage as Scale-out Repository without Archive Tier
- Local Directory and Azure Blob Object Storage as Scale-out Repository with Archive Tier

These steps must be configured before setting up Backup Jobs.

Azure Blob Object Storage

The backup infrastructure of Veeam Backup & Replication consists of Backup proxy servers, Backup Repositories, Object storage repositories, Scale-out Repositories, External Repositories, WAN Accelerators, Managed Servers...

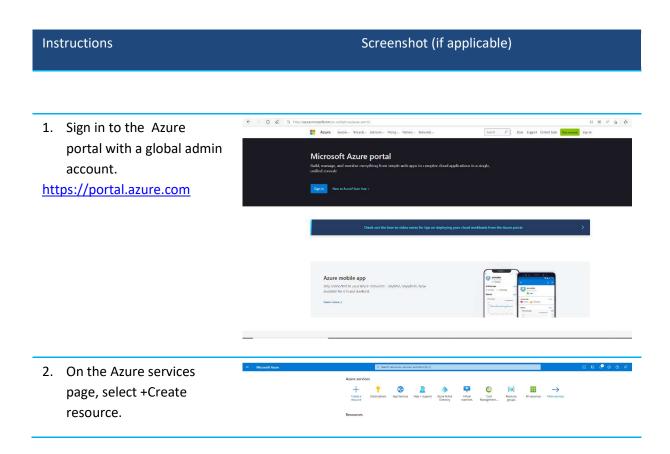
You can co-install Veeam Backup & Replication components on the same physical or virtual machine or set them up separately for a more scalable approach.

Adding Microsoft Azure Blob Object Storage Repositories

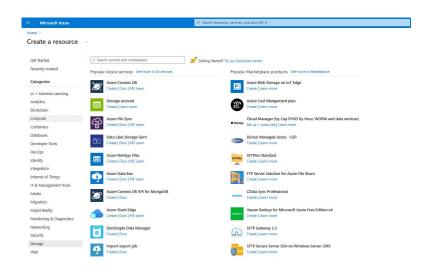
Veeam Backup&Replication supports different types of storage accounts.

Storage account type	Supported services	Supported performance tiers	Supported access tiers
General-purpose V2	Blob (block blobs	Standard	Excellent: for
Blob Storage	only)		infrequently accessed
			data.
			Hot: for frequently
			accessed data.
			Archive: for rarely
			accessed data. It can
			be set only on the
			blob level and
			supported in Archive
			Tier object storage
			systems.
			Veeam Backup &
			Replication will use
			the access tier you
			select as the default.

General-purpose V1	Blob (block blobs	Standard	N/A
	only)		
BlockBlob Storage	Blob (block blobs	Premium	N/A
	only)		



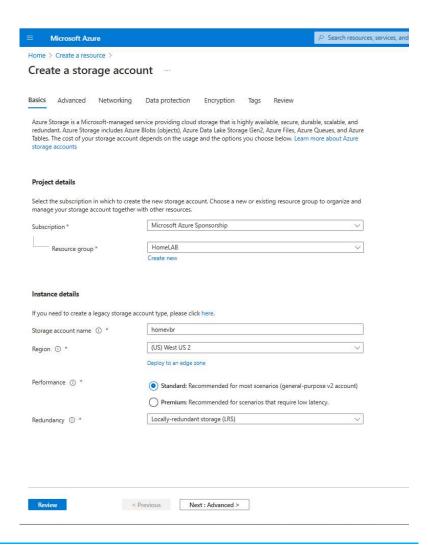
 Select Storage on the Create a resource page, and click Storage account.



4. On the Create storage account page, select Basics, configure as follows, and click Next: Advanced.

Section	Field	Required or optional	Description
Project details	Subscription	Required	Select the subscription for the new storage account.
Project details	Resource group	Required	Create a new resource group for this storage account, or select an existing one. For more information, see Resource groups.
Instance details	Storage account name	Required	Choose a unique name for your storage account. Storage account names must be between 3 and 24 characters in length and may contain numbers and lowercase letters only.
Instance details	Region	Required	Select the appropriate region for your storage account. For more information, see Regions and Availability Zones in Azure.
			Not all regions are supported for all types of storage accounts or redundancy configurations. For more information, see Azure Storage redundancy.
			The choice of region can have a billing impact. For more information, see Storage account billing.
Instance details	Performance	Required	Select Standard performance for general-purpose v2 storage accounts (default). This type of account is recommended by Microsoft for most scenarios. For more information, see Types of storage accounts.
			Select Premium for scenarios requiring low latency. After selecting Premium, select the type of premium storage account to create. The following types of premium storage accounts are available:
			Block blobs
			File shares
			Page blobs
Instance details	Redundancy	Required	Select your desired redundancy configuration. Not all redundancy options are available for all types of storage accounts in all regions. For more information about redundancy configurations, see Azure Storage redundancy.
			If you select a geo-redundant configuration (GRS or GZRS), your data is replicated to a data center in a different region. For read access to data in the secondary region, select Make read access to data available in the event of regional unavailability.

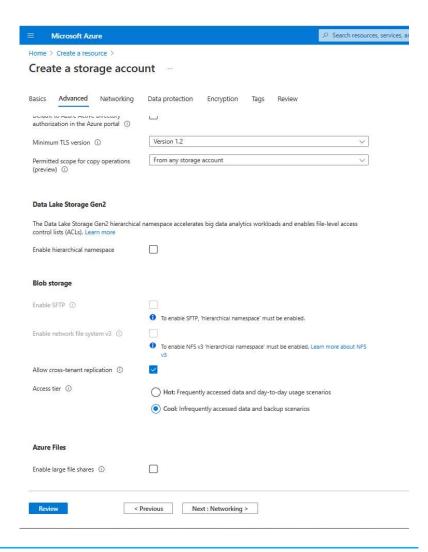
5. My settings are as screen capture.



6. On the Advanced page, configure and click Next: Networking.

Section	Field	Required or optional	Description
Security	Require secure transfer for REST API operations	Optional	Require secure transfer to ensure that incoming requests to this storage account are made only via HTTPS (default). Recommended for optimal security. For more information, see Require secure transfer to ensure secure connections.
Security	Enable blob public access	Optional	When enabled, this setting allows a user with the appropriate permissions to enable anonymous public access to a container in the storage account (default). Disabling this setting prevents all anonymous public access to the storage account. For more information, see Prevent anonymous public read access to containers and blobs.
			Enabling blob public access does not make blob data available for public access unless the user takes the additional step to explicitly configure the container's public access setting.
Security	Enable storage account key access	Optional	When enabled, this setting allows clients to authorize requests to the storage account using either the account access keys or an Azure Active Directory (Azure AD) account (default). Disabling this setting prevents authorization with the account access keys. For more information, see Prevent Shared Key authorization for an Azure Storage account.
Security	Default to Azure Active Directory authorization in the Azure portal	Optional	When enabled, the Azure portal authorizes data operations with the user's Azure AD credentials by default. If the user does not have the appropriate permissions assigned via Azure role-based access control (Azure RBAC) to perform data operations, then the portal will use the account access keys for data access instead. The user can also choose to switch to using the account access keys. For more information, see Default to Azure AD authorization in the Azure portal.
Security	Minimum TLS version	Required	Select the minimum version of Transport Layer Security (TLS) for incoming requests to the storage account. The default value is TLS version 1.2. When set to the default value, incoming requests made using TLS 1.0 or TLS 1.1 are rejected. For more information, see Enforce a minimum required version of Transport Layer Security (TLS) for requests to a storage account.
Data Lake Storage Gen2	Enable hierarchical namespace	Optional	To use this storage account for Azure Data Lake Storage Gen2 workloads, configure a hierarchical namespace. For more information, see Introduction to Azure Data Lake Storage Gen2.
Blob storage	Enable SFTP	Optional	Enable the use of Secure File Transfer Protocol (SFTP) to securely transfer of data over the internet. For more information, see Secure File Transfer (SFTP) protocol support in Azure Blob Storage.
Blob storage	Enable network file share (NFS) v3	Optional	NFS v3 provides Linux file system compatibility at object storage scale enables Linux clients to mount a container in Biob storage from an Azure Virtual Machine (VM) or a computer on-premises. For more information, see Network File System (NFS) 3.0 protocol support in Azure Biob storage.
Blob storage	Allow cross-tenant replication	Required	By default, users with appropriate permissions can configure object replication across Azure AD tenants. To prevent replication across tenants, deselect this option. For more information, see Prevent replication across Azure AD tenants.
Blob storage	Access tier	Required	Blob access tiers enable you to store blob data in the most cost-effective manner, based on usage. Select the hot tier (default) for frequently accessed data. Select the cool tier for infrequently accessed data. For more information, see Hot, Cool, and Archive access tiers for blob data.
Azure Files	Enable large file shares	Optional	Available only for standard file shares with the LRS or ZRS redundancies.

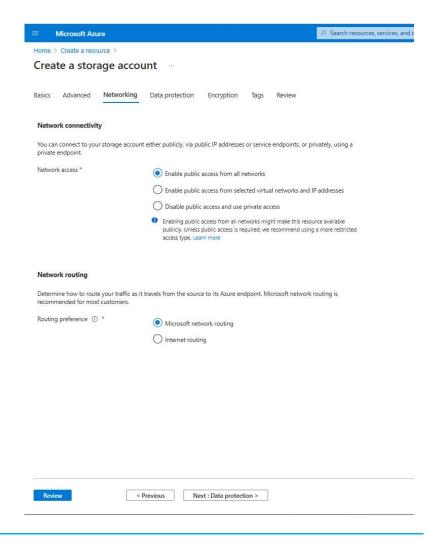
7. My settings are as screen capture.



8. On the Networking page, configure as follow and then click Next: Data protection.

Section	Field	Required or optional	Description
Network connectivity	Connectivity method	Required	By default, incoming network traffic is routed to the public endpoint for your storage account. You can specify that traffic must be routed to the public endpoint through an Azure virtual network. You can also configure private endpoints for your storage account. For more information, see Use private endpoints for Azure Storage.
Network routing	Routing preference	Required	The network routing preference specifies how network traffic is routed to the public endpoint of your storage account from clients over the internet. By default, a new storage account uses Microsoft network routing. You can also choose to route network traffic through the POP closest to the storage account, which may lower networking costs. For more information, see Network routing preference for Azure Storage.

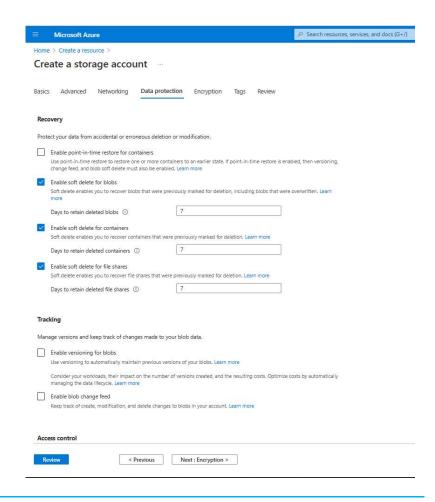
9. My settings are as screen capture.



10. On the Data protection page, configure as follow and then click Next:
Encryption.

Section	Field	Required or optional	Description
Recovery	Enable point-in- time restore for containers	Optional	Point-in-time restore provides protection against accidental deletion or corruption by enabling you to restore block blob data to an earlier state. For more information, see Point-in-time restore for block blobs.
			Enabling point-in-time restore also enables blob versioning, blob soft delete, and blob change feed. These prerequisite features may have a cost impact. For more information, see Pricing and billing for point-in-time restore.
Recovery	Enable soft delete for blobs	Optional	Blob soft delete protects an individual blob, snapshot, or version from accidental deletes or overwrites by maintaining the deleted data in the system for a specified retention period. During the retention period, you can restore a soft-deleted object to its state at the time it was deleted. For more information, see Soft delete for blobs.
			Microsoft recommends enabling blob soft delete for your storage accounts and setting a minimum retention period of seven days.
Recovery	Enable soft delete for containers	Optional	Container soft delete protects a container and its contents from accidental deletes by maintaining the deleted data in the system for a specified retention period. During the retention period, you can restore a soft-deleted container to its state at the time it was deleted. For more information, see Soft delete for containers (preview).
			Microsoft recommends enabling container soft delete for your storage accounts and setting a minimum retention period of seven days.
Recovery	Enable soft delete for file shares	Optional	Soft delete for file shares protects a file share and its contents from accidental deletes by maintaining the deleted data in the system for a specified retention period. During the retention period, you can restore a soft-deleted file share to its state at the time it was deleted. For more information, see Prevent accidental deletion of Azure file shares.
			Microsoft recommends enabling soft delete for file shares for Azure Files workloads and setting a minimum retention period of seven days.
Tracking	Enable versioning for blobs	Optional	Blob versioning automatically saves the state of a blob in a previous version when the blob is overwritten. For more information, see Blob versioning.
			Microsoft recommends enabling blob versioning for optimal data protection for the storage account.
Tracking	Enable blob change feed	Optional	The blob change feed provides transaction logs of all changes to all blobs in your storage account, as well as to their metadata. For more information, see Change feed support in Azure Blob Storage.
Access control	Enable version- level immutability support	Optional	Enable support for immutability policies that are scoped to the blob version. If this option is selected, then after you create the storage account, you can configure a default time-based retention policy for the account or for the container, which blob versions within the account or container will inherit by default. For more information, see Enable version-level immutability support on a storage account.

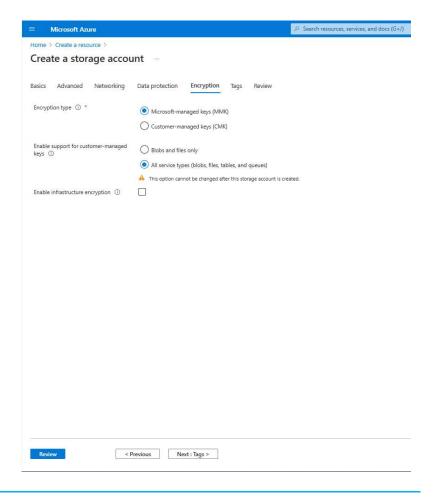
11. In my case, my settings are as screen capture.



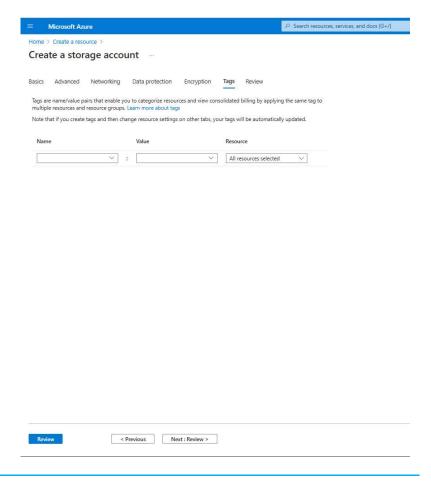
12. On the Encryption page, configure and click Next: Tags.

Field	Required or optional	Description
Encryption type	Required	By default, data in the storage account is encrypted by using Microsoft-managed keys. You can rely on Microsoft-managed keys for the encryption of your data, or you can manage encryption with your own keys. For more information, see Azure. Storage encryption for data at rest.
Enable support for customer- managed keys	Required	By default, customer managed keys can be used to encrypt only blobs and files. Set this option to All service types (blobs, files, tables, and queues) to enable support for customer-managed keys for all services. You are not required to use customer-managed keys if you choose this option. For more information, see Customer-managed keys for Azure Storage encryption.
Encryption key	Required if Encryption type field is set to Customer- managed keys.	If you choose Select a key vault and key, you are presented with the option to navigate to the key vault and key that you wish to use. If you choose Enter key from URI, then you are presented with a field to enter the key URI and the subscription.
User-assigned identity	Required if Encryption type field is set to Customer- managed keys.	If you are configuring customer-managed keys at create time for the storage account, you must provide a user-assigned identity to use for authorizing access to the key vault.
Enable infrastructure encryption	Optional	By default, infrastructure encryption is not enabled. Enable infrastructure encryption to encrypt your data at both the service level and the infrastructure level. For more information, see Create a storage account with infrastructure encryption enabled for double encryption of data.

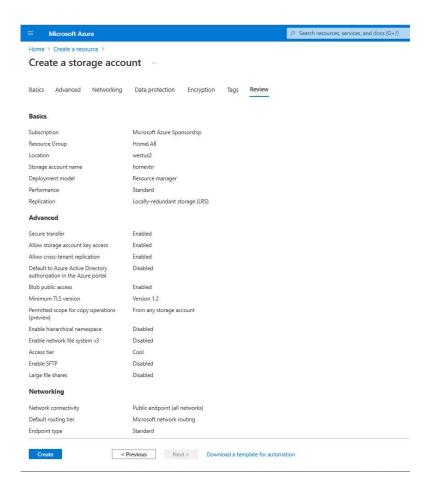
13. In my case, my settings are as screen capture.



14. On the Tags page, configure as follow and click Next: Review.You can specify ResourceManager tags on the Tags tab to help organize your Azure resources.



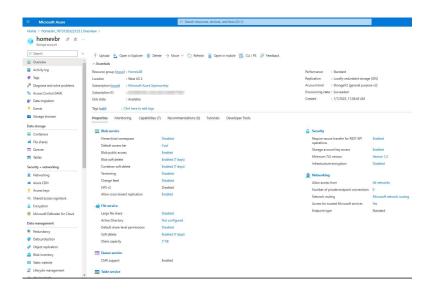
15. On the Review page, click Create.



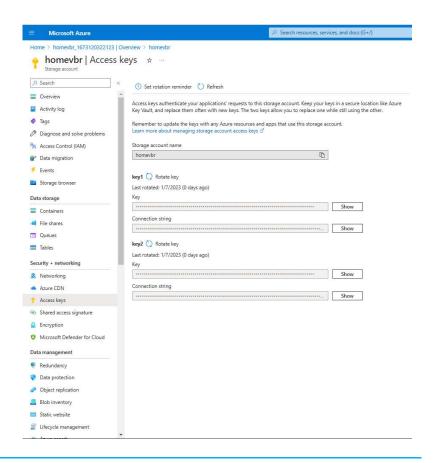
16. It may need a few mins to create the new storage account, and click Go to the resource.



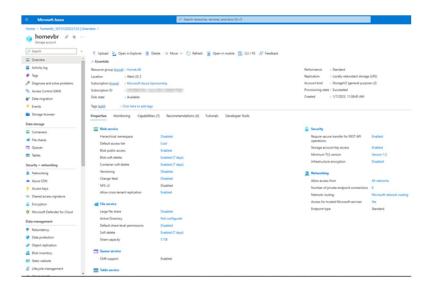
17. On the newly created Storage account page, select Access keys.



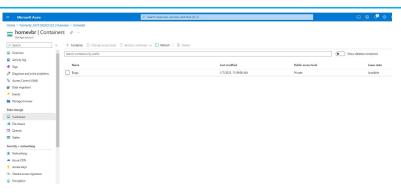
18. On the Access keys page, select Show keys and copy the Storage account name and key1. We need them for Veeam storage repository settings later.



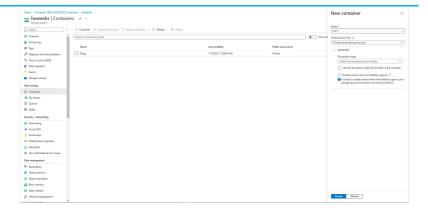
19. On the newly created Storage account page, select Containers.



20. On the Containers page, click +Container.

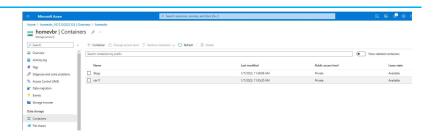


- 21. On the new container page, enter a name for your new container, select Private (no anonymous access) as Public access level and then click Create.
- 22. Veeam Backup and Replication v11 do not support Immutable



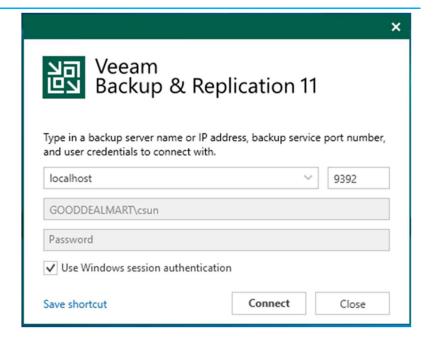
storage with versioning for azure blob storage.

23. Verify the new container created.

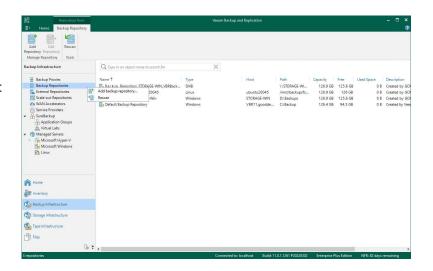


- 24. Log in to the Veeam

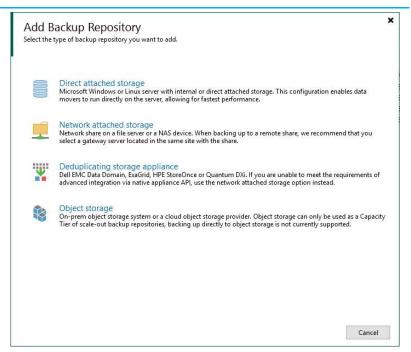
 Backup and replication
 manager server.
- 25. Open the Veeam Backup & Replication Console, and click Connect.



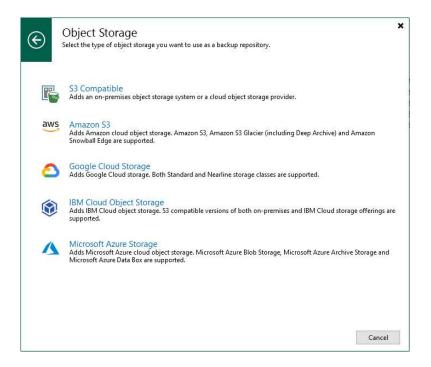
- 26. On the Home page, select Backup Infrastructure.
- 27. On the Backup
 Infrastructure page, select
 Backup Repositories,
 right-click Backup
 Repositories, and select
 Add backup repository.



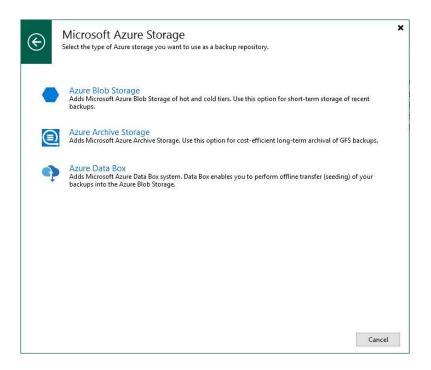
28. On the Add Backup
Repository page, select
Object storage.



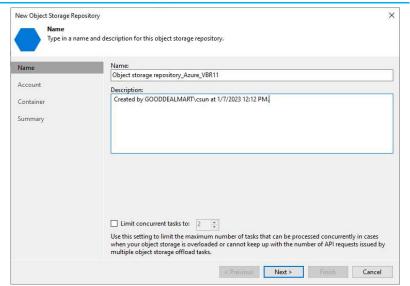
29. On the Object Storage page, select Microsoft Azure Storage.



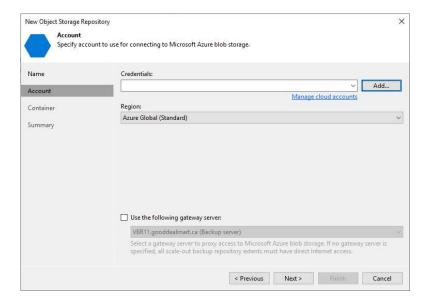
30. On the Microsoft Azure Storage page, select Azure Blob Storage.



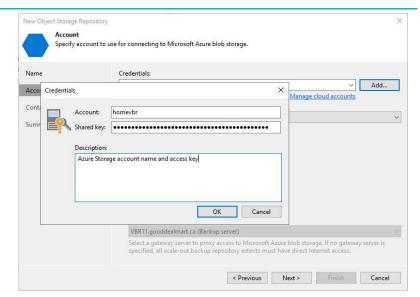
- 31. On the Name page, specify in the Name field.
- 32. In the Description field, describe future references.
- 33. Select the Limit concurrent tasks to N check box if you need to limit the maximum number of functions that can be processed at once
- 34. Click Next.



Click Add to create a
 Manage cloud account on the Account page.



36. On the Credentials page, paste the Azure storage account name as Account and paste the key1 as Shared key (you can find them from the Access key session of the Azure storage account), and click OK.

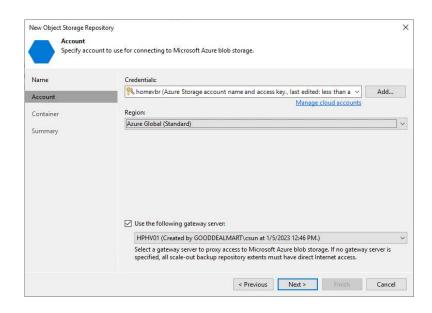


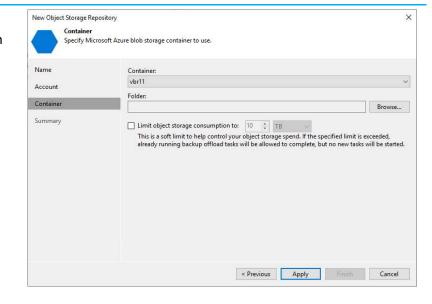
- On the Account page, select Azure Global (Standard) as Region.
- 38. Select Use the following gateway server checkbox and choose a server from the list.

Note:

If you do not select the Use the following gateway server check box, you must ensure all scale-out repository extents have direct internet access.

- 39. Click Next.
- 40. Select a container from the Container drop-down list on the Container page.
- 41. In the Folder field, click Browse.

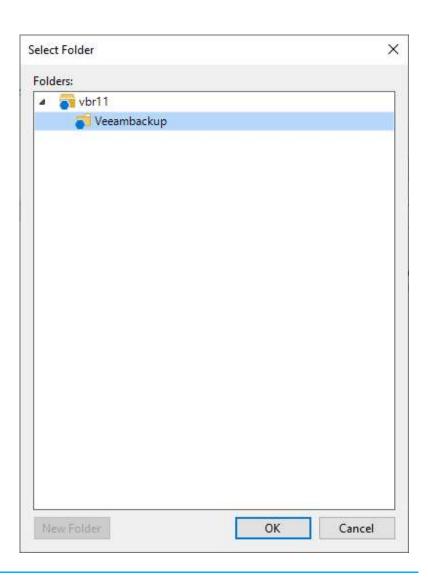




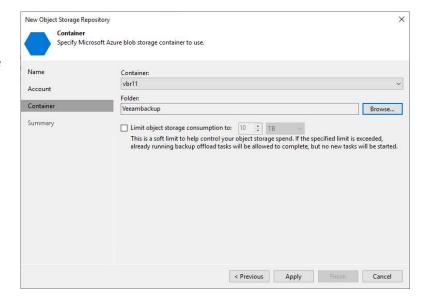
42. On the Folders page, select the container and click New Folder.



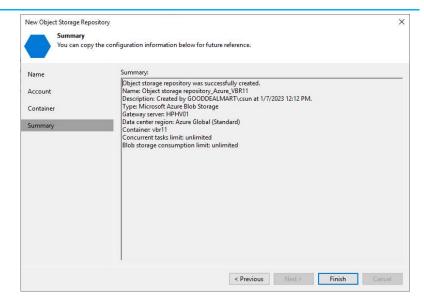
43. Enter the name for the new folder, select the folder and click OK.



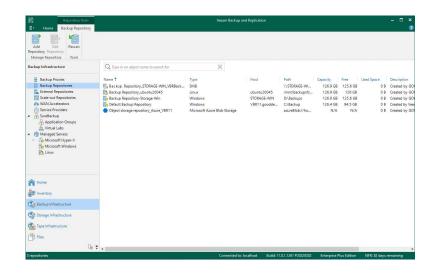
- 44. Select the Limit object storage consumption to check box and provide the value in TB or PB if you need to control the object storage spend.
- 45. Click Apply.



46. On the Summary page, click Finish.



47. Verify that the Backup Repository has been added



Adding Microsoft Azure Archive Storage Repositories without Azure Helper appliance

Veeam Backup&Replication supports different types of storage accounts.

Storage account type	Supported services	Supported performance tiers	Supported access tiers
General-purpose V2	Blob (block blobs only)	Standard	Excellent: for infrequently accessed
Blob Storage	,,		data.
			Hot: for frequently accessed data.
			Archive: for rarely accessed data. It can
			be set only on the blob level and
			supported in Archive Tier object storage systems.
			Veeam Backup & Replication will use
			the access tier you select as the default.
General-purpose V1	Blob (block blobs only)	Standard	N/A

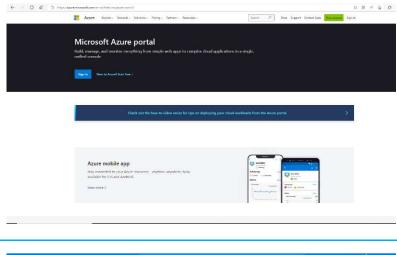
BlockBlobStorage	Blob (block blobs only)	Premium	N/A

Instructions

Screenshot (if applicable)

1. Sign in to the Azure portal with a global admin account.

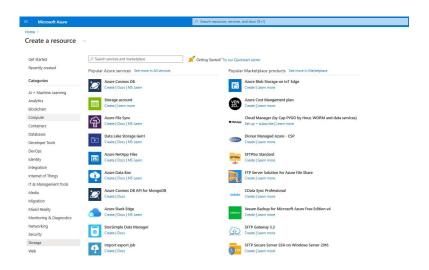
https://portal.azure.com



2. On the Azure services page, select +Create resource.



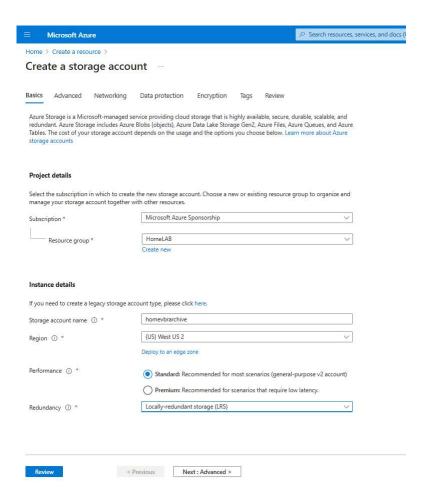
3. Select Storage on the Create a resource page, and click Storage account.



4. On the Create storage account page, select Basics, configure as follows, and click Next: Advanced.

Section	Field	Required or optional	Description
Project details	Subscription	Required	Select the subscription for the new storage account.
Project details	Resource group	Required	Create a new resource group for this storage account, or select an existing one. For more information, see Resource groups.
Instance details	Storage account name	Required	Choose a unique name for your storage account. Storage account names must be between 3 and 24 characters in length and may contain numbers and lowercase letters only.
Instance details	Region	Required	Select the appropriate region for your storage account. For more information, see Regions and Availability Zones in Azure.
			Not all regions are supported for all types of storage accounts or redundancy configurations. For more information, see Azure Storage redundancy.
			The choice of region can have a billing impact. For more information, see Storage account billing.
Instance Performance details	Performance	Required	Select Standard performance for general-purpose v2 storage accounts (default). This type of account is recommended by Microsoft for most scenarios. For more information, see Types of storage accounts.
			Select Premium for scenarios requiring low latency. After selecting Premium, select the type of premium storage account to create. The following types of premium storage accounts are available:
			Block blobs
			File shares
			Page blobs
Instance details	Redundancy	Required	Select your desired redundancy configuration. Not all redundancy options are available for all types of storage accounts in all regions. For more information about redundancy configurations, see Azure Storage redundancy.
			If you select a geo-redundant configuration (GRS or GZRS), your data is replicated to a data center in a different region. For read access to data in the secondary region, select Make read access to data available in the event of regional unavailability.

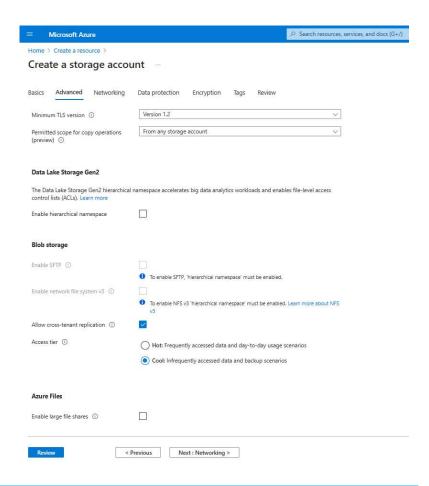
5. My settings are as screen capture.



6. On the Advanced page, configure and click Next: Networking.

Section	Field	Required or optional	Description
Security	Require secure transfer for REST API operations	Optional	Require secure transfer to ensure that incoming requests to this storage account are made only via HTTPS (default). Recommended for optimal security. For more information, see Require secure transfer to ensure secure connections.
Security	Enable blob public access	Optional	When enabled, this setting allows a user with the appropriate permissions to enable anonymous public access to a container in the storage account (default). Disabling this setting prevents all anonymous public access to the storage account. For more information, see Prevent anonymous public read access to containers and blobs.
			Enabling blob public access does not make blob data available for public access unless the user takes the additional step to explicitly configure the container's public access setting.
Security	Enable storage account key access	Optional	When enabled, this setting allows clients to authorize requests to the storage account using either the account access keys or an Azure Active Directory (Azure AD) account (default). Disabling this setting prevents authorization with the account access keys. For more information, see Prevent Shared Key authorization for an Azure Storage account.
Security	Default to Azure Active Directory authorization in the Azure portal	Optional	When enabled, the Azure portal authorizes data operations with the user's Azure AD credentials by default. If the user does not have the appropriate permissions assigned via Azure role-based access control (Azure RBAC) to perform data operations, then the portal will use the account access keys for data access instead. The user can also choose to switch to using the account access keys. For more information, see Default to Azure AD authorization in the Azure portal.
Security	Minimum TLS version	Required	Select the minimum version of Transport Layer Security (TLS) for incoming requests to the storage account. The default value is TLS version 1.2. When set to the default value, incoming requests made using TLS 1.0 or TLS 1.1 are rejected. For more information, see Enforce a minimum required version of Transport Layer Security (TLS) for requests to a storage account.
Data Lake Storage Gen2	Enable hierarchical namespace	Optional	To use this storage account for Azure Data Lake Storage Gen2 workloads, configure a hierarchical namespace. For more information, see Introduction to Azure Data Lake Storage Gen2.
Blob storage	Enable SFTP	Optional	Enable the use of Secure File Transfer Protocol (SFTP) to securely transfer of data over the internet. For more information, see Secure File Transfer (SFTP) protocol support in Azure Blob Storage.
Blob storage	Enable network file share (NFS) v3	Optional	NFS v3 provides Linux file system compatibility at object storage scale enables Linux clients to mount a container in Biob storage from an Azure Virtual Machine (VM) or a computer on-premises. For more information, see Network File System (NFS) 3.0 protocol support in Azure Biob storage.
Blob storage	Allow cross-tenant replication	Required	By default, users with appropriate permissions can configure object replication across Azure AD tenants. To prevent replication across tenants, deselect this option. For more information, see Prevent replication across Azure AD tenants.
Blob storage	Access tier	Required	Blob access tiers enable you to store blob data in the most cost-effective manner, based on usage. Select the hot tier (default) for frequently accessed data. Select the cool tier for infrequently accessed data. For more information, see Hot, Cool, and Archive access tiers for blob data.
Azure Files	Enable large file shares	Optional	Available only for standard file shares with the LRS or ZRS redundancies.

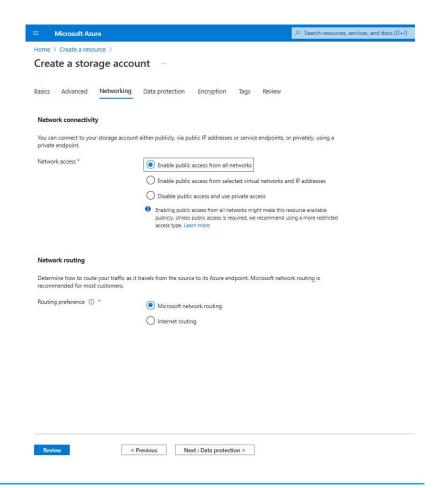
7. My settings are as screen capture.



 On the Networking page, configure as follow and then click Next: Data protection.

Section	Field	Required or optional	Description
Network connectivity	Connectivity method	Required	By default, incoming network traffic is routed to the public endpoint for your storage account. You can specify that traffic must be routed to the public endpoint through an Azure virtual network. You can also configure private endpoints for your storage account. For more information, see Use private endpoints for Azure Storage
Network routing	Routing preference	Required	The network routing preference specifies how network traffic is routed to the public endpoint of your storage account from clients over the internet. By default, a new storage account uses Microsoft network routing. You can also choose to route network traffic through the POP closest to the storage account, which may lower networking costs. For more information, see Network routing preference for Azure Storage.

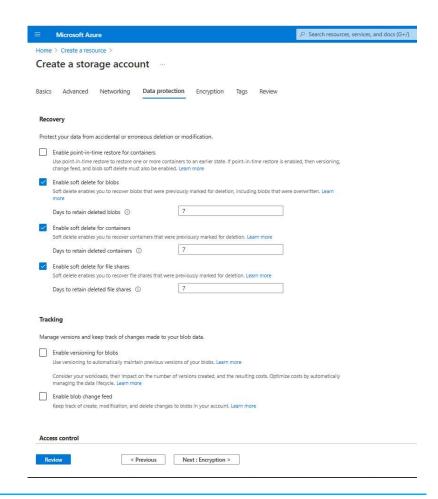
9. My settings are as screen capture.



10. On the Data protection page, configure as follow and then click Next: Encryption.

Section	Field	Required or optional	Description
Recovery	Enable point-in- time restore for containers	Optional	Point-in-time restore provides protection against accidental deletion or corruption by enabling you to restore block blob data to an earlier state. For more information, see Point-in-time restore for block blobs.
			Enabling point-in-time restore also enables blob versioning, blob soft delete, and blob change feed. These prerequisite features may have a cost impact. For more information, see Pricing and billing for point-in-time restore.
The second second	Enable soft delete for blobs	Optional	Blob soft delete protects an individual blob, snapshot, or version from accidental deletes or overwrites by maintaining the deleted data in the system for a specified retention period. During the retention period, you can restore a soft-deleted object to its state at the time it was deleted. For more information, see Soft delete for blobs.
			Microsoft recommends enabling blob soft delete for your storage accounts and setting a minimum retention period of seven days.
Recovery	Enable soft delete for containers	Optional	Container soft delete protects a container and its contents from accidental deletes by maintaining the deleted data in the system for a specified retention period. During the retention period, you can restore a soft-deleted container to its state at the time it was deleted. For more information, see Soft delete for containers (preview).
			Microsoft recommends enabling container soft delete for your storage accounts and setting a minimum retention period of seven days.
Recovery	Enable soft delete for file shares	Optional	Soft delete for file shares protects a file share and its contents from accidental deletes by maintaining the deleted data in the system for a specified retention period. During the retention period, you can restore a soft-deleted file share to its state at the time it was deleted. For more information, see Prevent accidental deletion of Azure file shares.
		Microsoft recommends enabling soft delete for file shares for Azure Files workloads and setting a minimum retention period of seven days.	
Tracking	Enable versioning for blobs	Optional	Blob versioning automatically saves the state of a blob in a previous version when the blob is overwritten. For more information, see Blob versioning.
			Microsoft recommends enabling blob versioning for optimal data protection for the storage account.
Tracking	Enable blob change feed	Optional	The blob change feed provides transaction logs of all changes to all blobs in your storage account, as well as to their metadata. For more information, see Change feed support in Azure Blob Storage.
Access	Enable version- level immutability support	Optional	Enable support for immutability policies that are scoped to the blob version. If this option is selected, then after you create the storage account, you can configure a default time-based retention policy for the account or for the container, which blob versions within the account or container will inherit by default. For more information, see Enable version-level immutability support on a storage account.

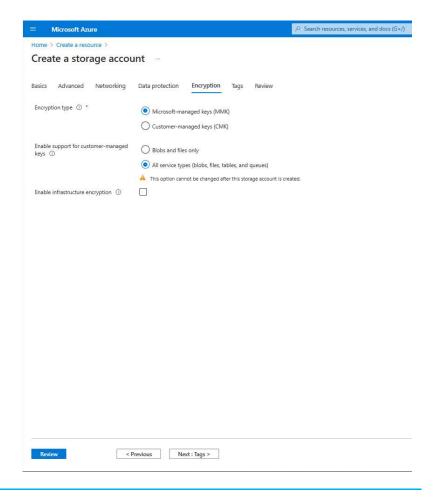
11. In my case, my settings are as screen capture.



12. On the Encryption page, configure and click Next: Tags.

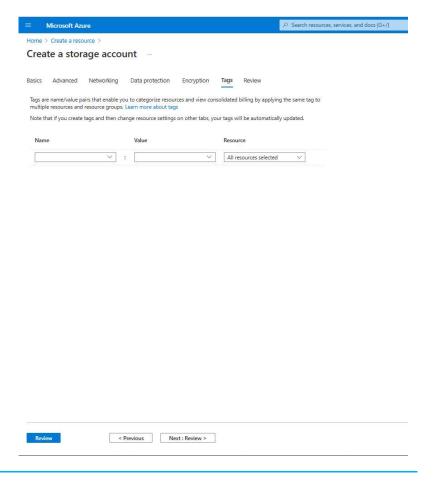
Field	Required or optional	Description
Encryption type	Required	By default, data in the storage account is encrypted by using Microsoft-managed keys. You can rely on Microsoft-managed keys for the encryption of your data, or you can manage encryption with your own keys. For more information, see Azure- Storage encryption for data at rest.
Enable support for customer- managed keys	Required	By default, customer managed keys can be used to encrypt only blobs and files. Set this option to All service types (blobs, files, tables, and queues) to enable support for customer-managed keys for all services. You are not required to use customer-managed keys if you choose this option. For more information, see Customer-managed keys for Azure Storage encryption.
Encryption key	Required if Encryption type field is set to Customer- managed keys.	If you choose Select a key vault and key, you are presented with the option to navigate to the key vault and key that you wish to use. If you choose Enter key from URI, then you are presented with a field to enter the key URI and the subscription.
User-assigned identity	Required if Encryption type field is set to Customer- managed keys.	If you are configuring customer-managed keys at create time for the storage account, you must provide a user-assigned identity to use for authorizing access to the key vault.
Enable infrastructure encryption	Optional	By default, infrastructure encryption is not enabled. Enable infrastructure encryption to encrypt your data at both the service level and the infrastructure level. For more information, see Create a storage account with infrastructure encryption enabled for double encryption of data.

13. In my case, my settings are as screen capture.

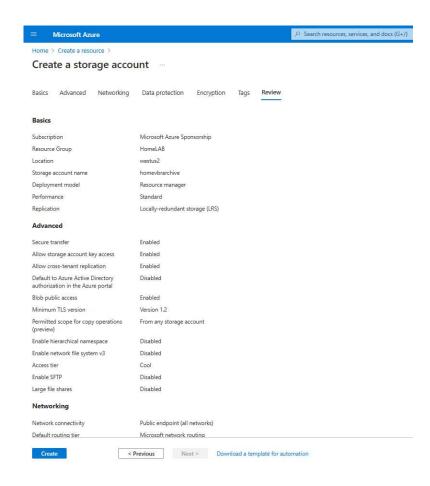


14. On the Tags page, configure as follow and click Next: Review.

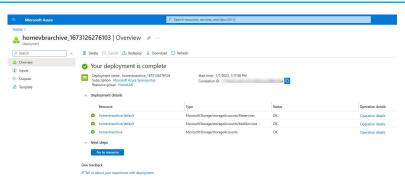
You can specify Resource Manager tags on the Tags tab to help organize your Azure resources.



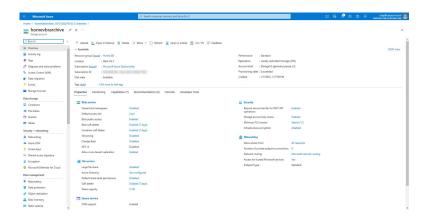
15. On the Review page, click Create.



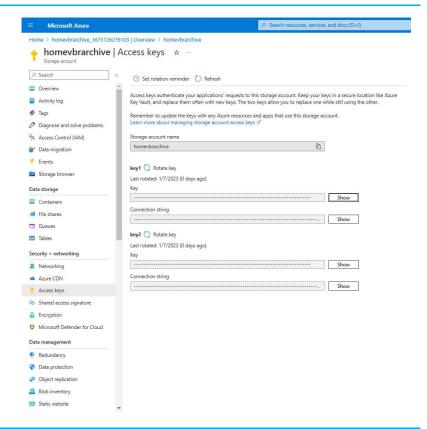
16. It may need a few mins to create the new storage account, and click Go to the resource.



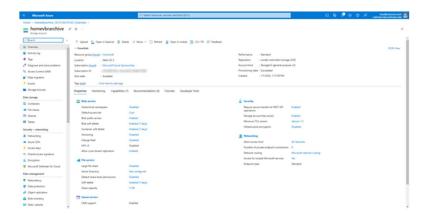
17. On the newly created Storage account page, select Access keys.



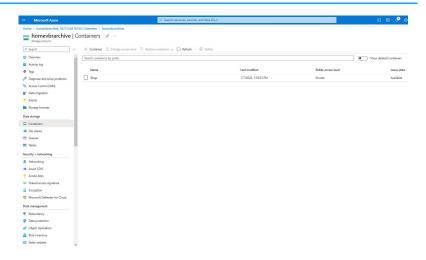
18. On the Access keys page, select Show keys and copy the Storage account name and key of key1. We need them for Veeam storage repository settings later.



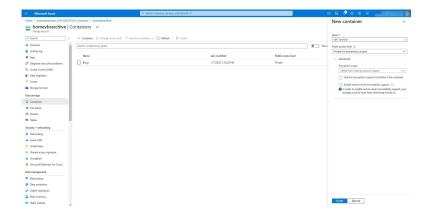
19. On the newly created Storage account page, select Containers.



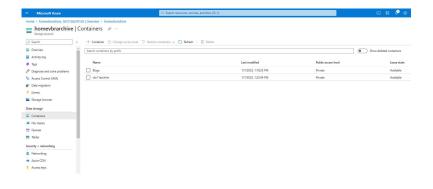
20. On the Containers page, click +Container.



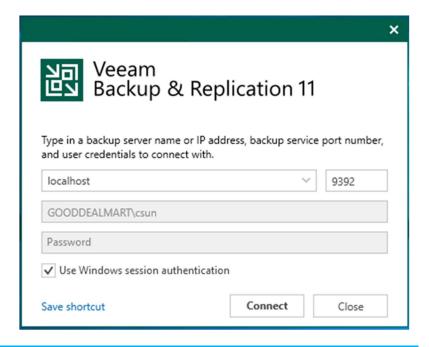
- 21. On the new container page, enter a name for your new container, select Private (no anonymous access) as Public access level and then click Create.
- 22. Veeam Backup and Replication v11 do not support Immutable storage with versioning for azure blob storage.



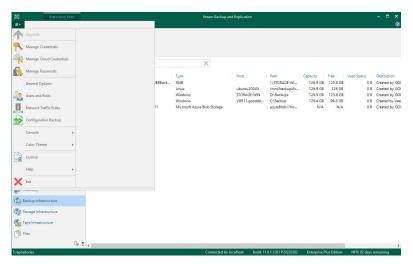
23. Verify the new container created.



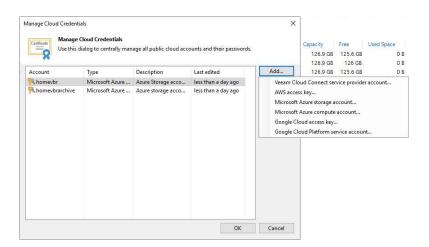
- 24. Log in to the Veeam Backup and replication manager server.
- 25. Open the Veeam Backup & Replication Console, and click Connect.



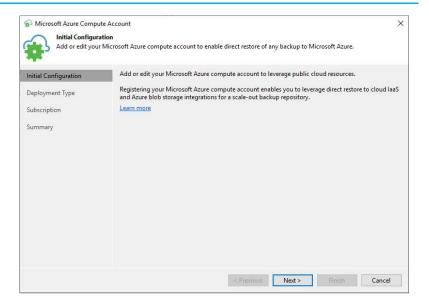
26. Select Manage Cloud Credentials from the main menu.



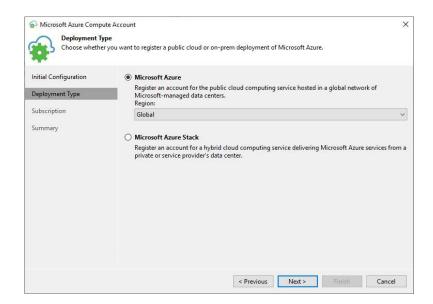
27. Click Add on the Manage Cloud Credentials page and select Microsoft Azure compute account.



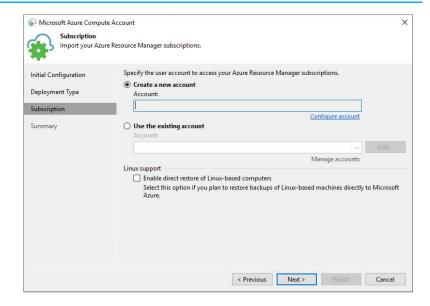
28. On the Initial
Configuration page, click
Next.



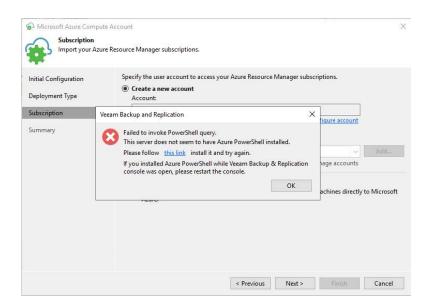
29. On the Deployment Type page, select Microsoft Azure, select a Microsoft Azure region from the Region drop-down list, and click Next.



30. On the Subscription page, select a new account and click Configure account.



- 31. Microsoft Azure
 PowerShell is installed on
 the machine that runs the
 Veeam Backup &
 Replication console. The
 Veeam Backup &
 Replication will display a
 warning if Microsoft
 Azure PowerShell is not
 installed.
- 32. In the warning window, click this link.



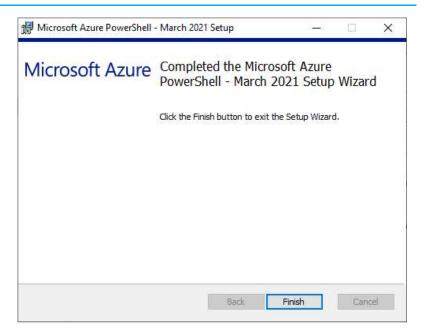
33. On the Microsoft Azure
PowerShell setup page,
select I accept the terms
in the License Agreement
checkbox and click Install.



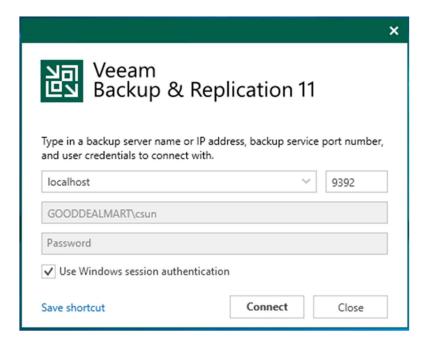
34. On the User Access Control page, click Yes.



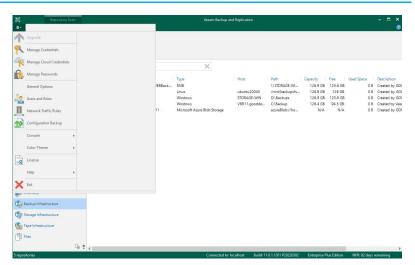
- 35. On the Microsoft Azure PowerShell setup page, click Finish.
- 36. Close the Veeam Backup & Replication console after completing the installation.
- 37. Reboot the machine to allow the Veeam Backup & Replication console to detect the newly installed version of Microsoft Azure PowerShell.



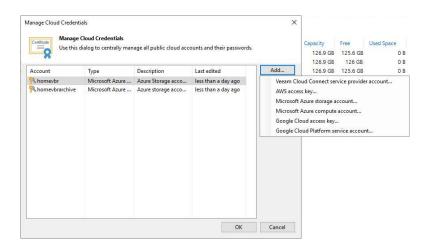
- 38. Log in to the Veeam Backup and replication manager server.
- 39. Open the Veeam Backup & Replication Console, and click Connect.



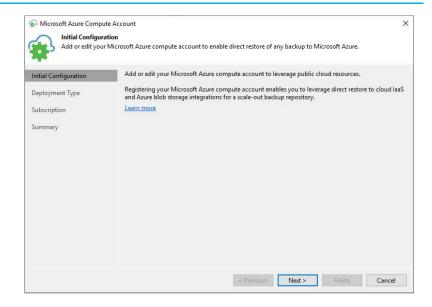
40. Select Manage Cloud Credentials from the main menu.



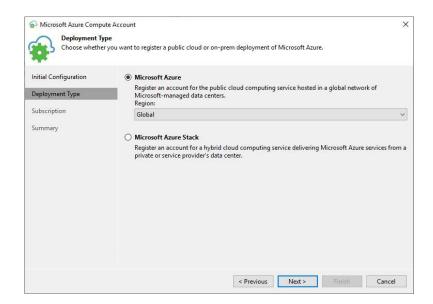
41. Click Add on the Manage Cloud Credentials page and select Microsoft Azure compute account.



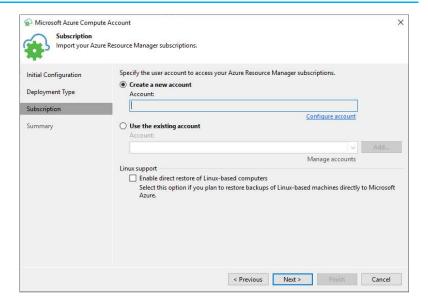
42. On the Initial
Configuration page, click
Next.



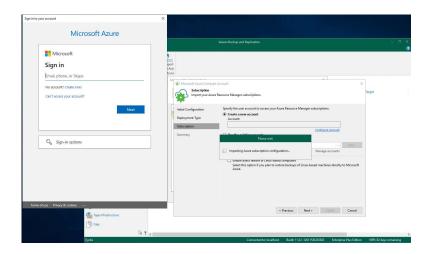
43. On the Deployment Type page, select Microsoft Azure, select a Microsoft Azure region from the Region drop-down list, and click Next.



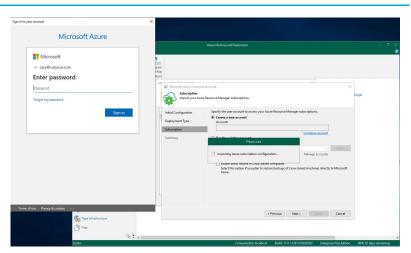
44. On the Subscription page, select a new account and click Configure account.



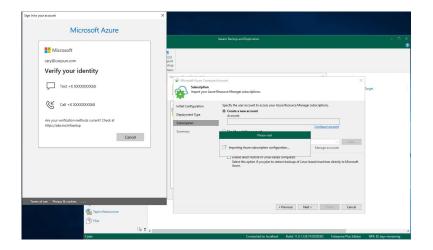
45. Enter the user account on the Microsoft Azure Signin page and click Next.



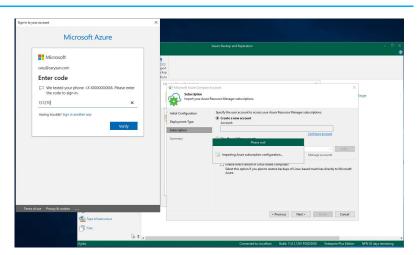
46. Type the password on the Enter password page, and click Sign in.



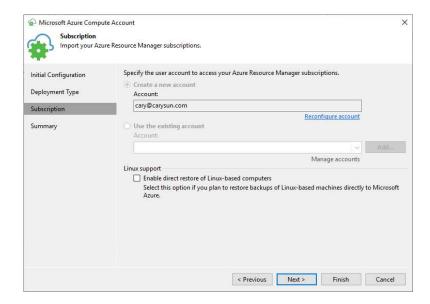
47. Select the verification method on the Verify your identity page.



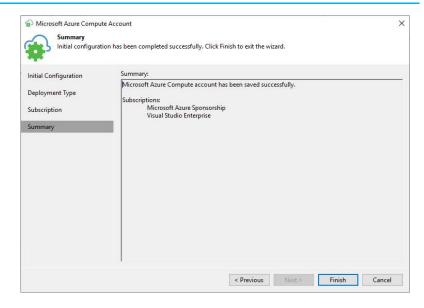
48. On the Enter code page, enter the texted code and click Verify.



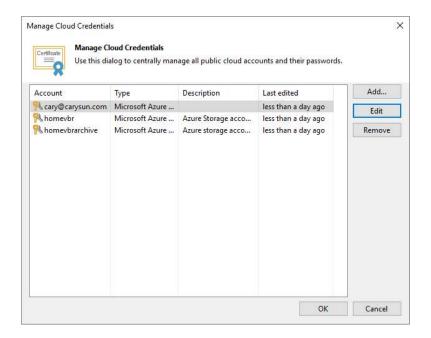
49. On the Subscription page, click Next if you don't need to enable direct restore of Linux-base computers.



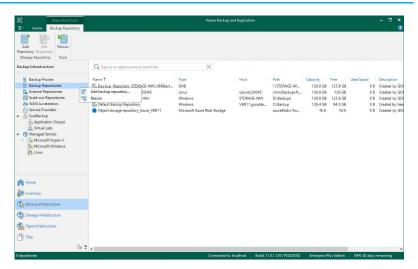
50. On the Summary page, click Finish.



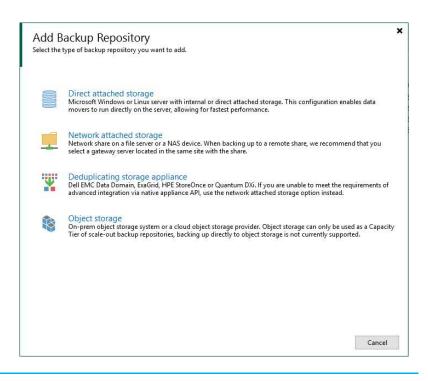
51. On the Manage Cloud Credentials page, click OK.



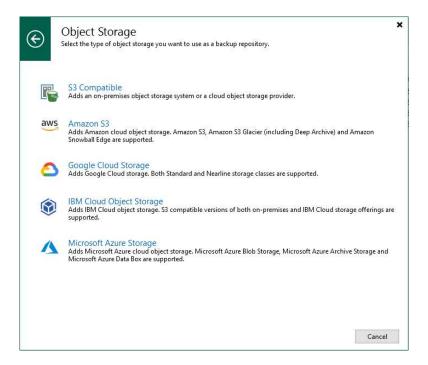
- 52. On the Home page, select Backup Infrastructure.
- 53. On the Backup
 Infrastructure page, select
 Backup Repositories,
 right-click Backup
 Repositories, and select
 Add backup repository.



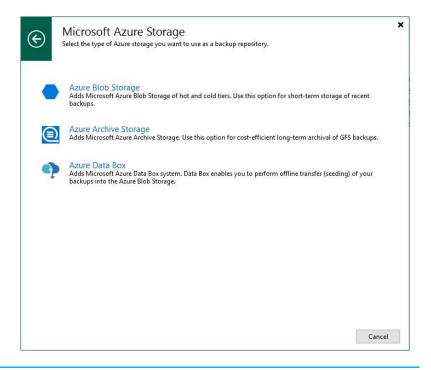
54. On the Add Backup Repository page, select Object storage.



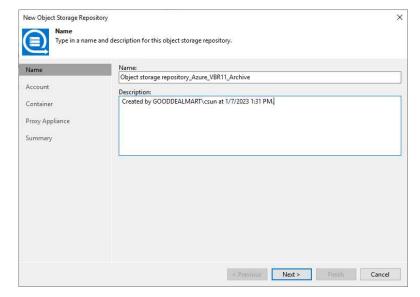
55. On the Object Storage page, select Microsoft Azure Storage.



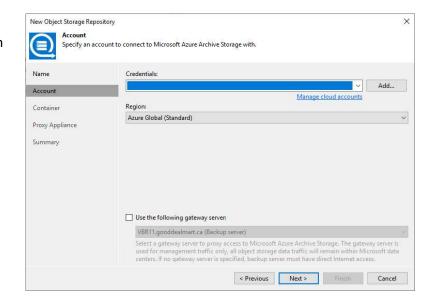
56. On the Microsoft Azure Storage page, select Azure Blob Storage.



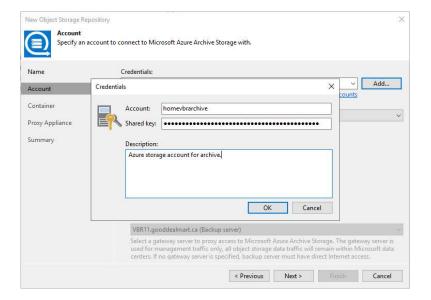
- 57. On the Name page, specify in the Name field.
- 58. In the Description field, describe future references.
- 59. Click Next.



60. Click Add to create a Manage cloud account on the Account page.



61. On the Credentials page, paste the Azure storage account name as Account and paste the key1 as Shared key (you can find them from the Access key session of the Azure storage account), and click OK.



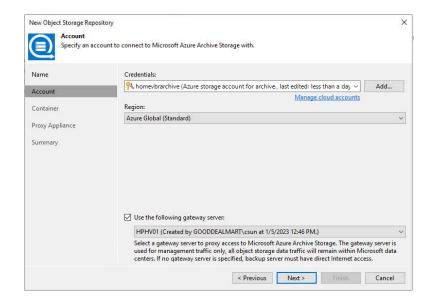
- 62. Select Azure Global (Standard) as Region on the Account page.
- 63. Select Use the following gateway server checkbox and choose a server from the list.

Note:

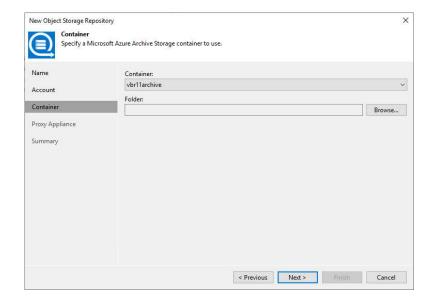
You can select any Microsoft Windows or Linux server added to your backup infrastructure with an internet connection, except Linux servers with the hardened repository role.

If you do not select the Use the following gateway server check box, you must make sure that the backup server has direct internet access

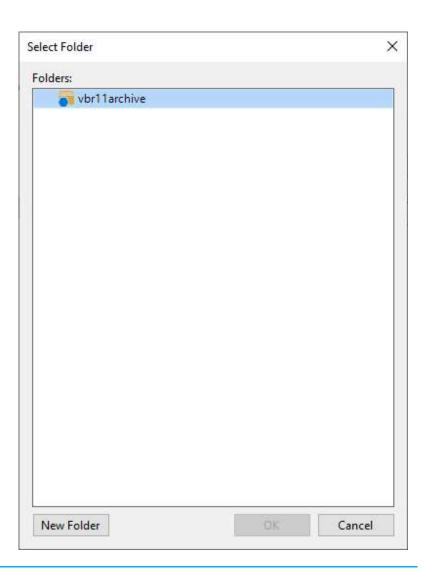
64. Click Next.



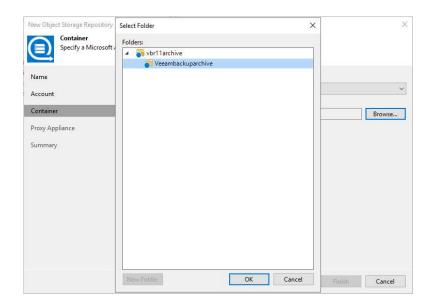
- 65. Select a container from the Container drop-down list on the Container page.
- 66. In the Folder field, click Browse.



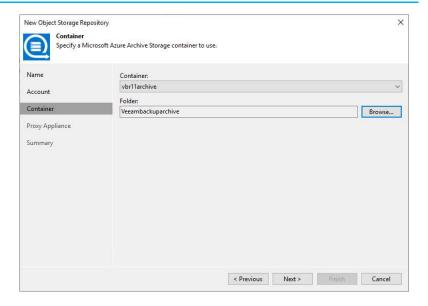
67. On the Folders page, select the container and click New Folder.



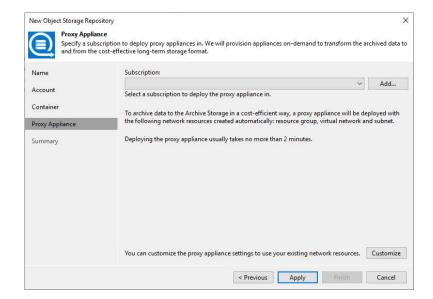
68. Enter the name for the new folder, select the folder and click OK.



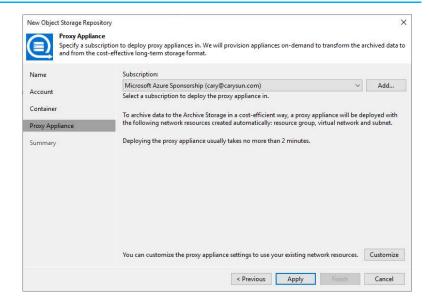
69. On the Container page, click Next.



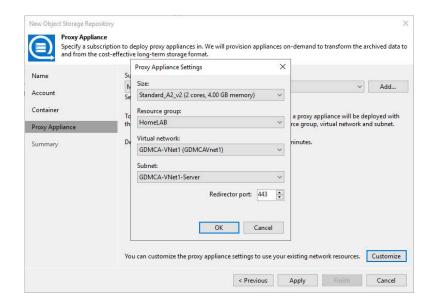
70. Select your Microsoft
Azure subscription
credentials from the
Subscription drop-down
list on the Proxy
Appliance page.



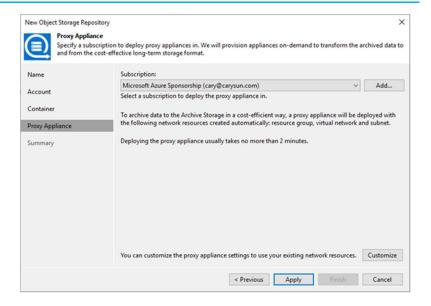
71. On the Proxy Appliance page, click Customise to customize the proxy appliance.



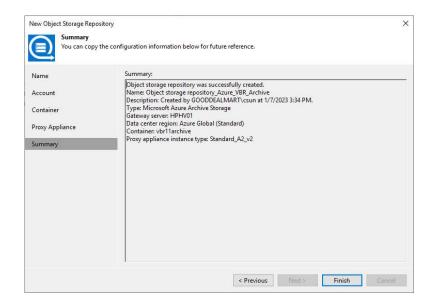
72. On the Proxy Appliance Settings page, specify the information for the proxy appliance and click OK.



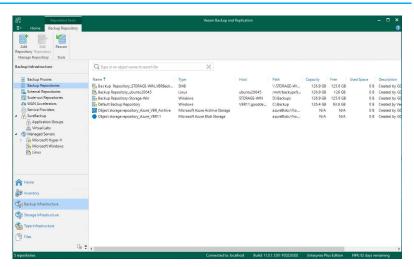
73. On the Proxy Appliance page, click Apply.



74. On the Summary page, click Finish.



75. Verify that the Backup Repository has been added.



Adding Microsoft Azure Archive Storage Repositories with Azure Helper appliance

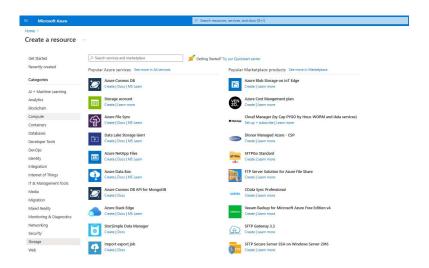
Veeam Backup&Replication supports different types of storage accounts.

Storage account type	Supported services	Supported performance tiers	Supported access tiers
General-purpose V2 BlobStorage	Blob (block blobs only)	Standard	Excellent: for infrequently accessed data.
			Hot: for frequently accessed data.
			Archive: for rarely accessed data. It can be set only on the blob level and supported in Archive Tier object storage systems.
			Veeam Backup & Replication will use the access tier you select as the default.
General-purpose V1	Blob (block blobs only)	Standard	N/A
BlockBlobStorage Blob (block blobs only)		Premium	N/A

resource.

1. Sign in Azure portal with a global admin account. https://portal.azure.com 2. On the Azure services page, select +Create

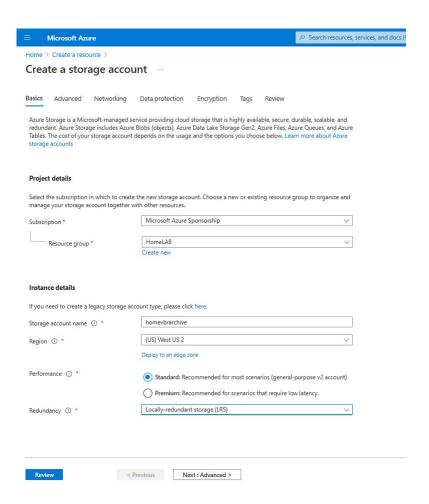
3. Select Storage on the Create a resource page, and click Storage account.



4. On the Create storage account page, select Basics, configure as follows, and click Next: Advanced.

Section	Field	Required or optional	Description
Project details	Subscription	Required	Select the subscription for the new storage account.
Project details	Resource group	Required	Create a new resource group for this storage account, or select an existing one. For more information, see Resource groups.
Instance details	Storage account name	Required	Choose a unique name for your storage account. Storage account names must be between 3 and 24 characters in length and may contain numbers and lowercase letters only.
Instance details	Region	Required	Select the appropriate region for your storage account. For more information, see Regions and Availability Zones in Azure.
			Not all regions are supported for all types of storage accounts or redundancy configurations. For more information, see Azure Storage redundancy.
			The choice of region can have a billing impact. For more information, see Storage account billing.
Instance Perform details	Performance	Required	Select Standard performance for general-purpose v2 storage accounts (default). This type of account is recommended by Microsoft for most scenarios. For more information, see Types of storage accounts.
			Select Premium for scenarios requiring low latency. After selecting Premium, select the type of premium storage account to create. The following types of premium storage accounts are available:
			Block blobs
			File shares
			Page blobs
Instance details	Redundancy	Required	Select your desired redundancy configuration. Not all redundancy options are available for all types of storage accounts in all regions. For more information about redundancy configurations, see Azure Storage redundancy.
			If you select a geo-redundant configuration (GRS or GZRS), your data is replicated to a data center in a different region. For read access to data in the secondary region, select Make read access to data available in the event of regional unavailability.

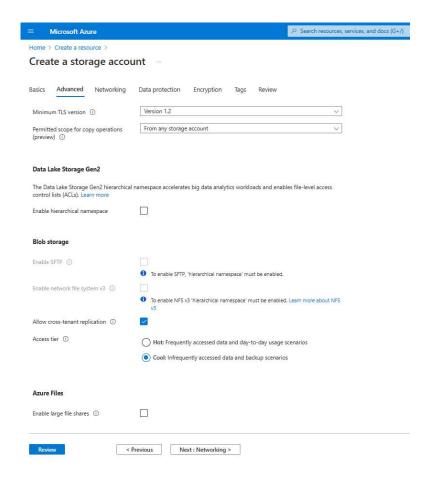
5. My settings are as screen capture.



6. On the Advanced page, configure and click Next: Networking.

Section	Field	Required or optional	Description
Security	Require secure transfer for REST API operations	Optional	Require secure transfer to ensure that incoming requests to this storage account are made only via HTTPS (default). Recommended for optimal security. For more information, see Require secure transfer to ensure secure connections.
Security	Enable blob public access	Optional	When enabled, this setting allows a user with the appropriate permissions to enable anonymous public access to a container in the storage account (default). Disabling this setting prevents all anonymous public access to the storage account. For more information, see Prevent anonymous public read access to containers and blobs.
			Enabling blob public access does not make blob data available for public access unless the user takes the additional step to explicitly configure the container's public access setting.
Security	Enable storage account key access	Optional	When enabled, this setting allows clients to authorize requests to the storage account using either the account access keys or an Azure Active Directory (Azure AD) account (default). Disabling this setting prevents authorization with the account access keys. For more information, see Prevent Shared Key authorization for an Azure Storage account.
Security	Default to Azure Active Directory authorization in the Azure portal	Optional	When enabled, the Azure portal authorizes data operations with the user's Azure AD credentials by default. If the user does not have the appropriate permissions assigned via Azure role-based access control (Azure RBAC) to perform data operations, then the portal will use the account access keys for data access instead. The user can also choose to switch to using the account access keys. For more information, see Default to Azure AD authorization in the Azure portal.
Security	Minimum TLS version	Required	Select the minimum version of Transport Layer Security (TLS) for incoming requests to the storage account. The default value is TLS version 1.2. When set to the default value, incoming requests made using TLS 1.0 or TLS 1.1 are rejected. For more information, see Enforce a minimum required version of Transport Layer Security (TLS) for requests to a storage account.
Data Lake Storage Gen2	Enable hierarchical namespace	Optional	To use this storage account for Azure Data Lake Storage Gen2 workloads, configure a hierarchical namespace. For more information, see Introduction to Azure Data Lake Storage Gen2.
Blob storage	Enable SFTP	Optional	Enable the use of Secure File Transfer Protocol (SFTP) to securely transfer of data over the internet. For more information, see Secure File Transfer (SFTP) protocol support in Azure Blob Storage.
Blob storage	Enable network file share (NFS) v3	Optional	NFS v3 provides Linux file system compatibility at object storage scale enables Linux clients to mount a container in Biob storage from an Azure Virtual Machine (VM) or a computer on-premises. For more information, see Network File System (NFS) 3.0 protocol support in Azure Biob storage.
Blob storage	Allow cross-tenant replication	Required	By default, users with appropriate permissions can configure object replication across Azure AD tenants. To prevent replication across tenants, deselect this option. For more information, see Prevent replication across Azure AD tenants.
Blob storage	Access tier	Required	Blob access tiers enable you to store blob data in the most cost-effective manner, based on usage. Select the hot tier (default) for frequently accessed data. Select the cool tier for infrequently accessed data. For more information, see Hot, Cool, and Archive access tiers for blob data.
Azure Files	Enable large file shares	Optional	Available only for standard file shares with the LRS or ZRS redundancies.

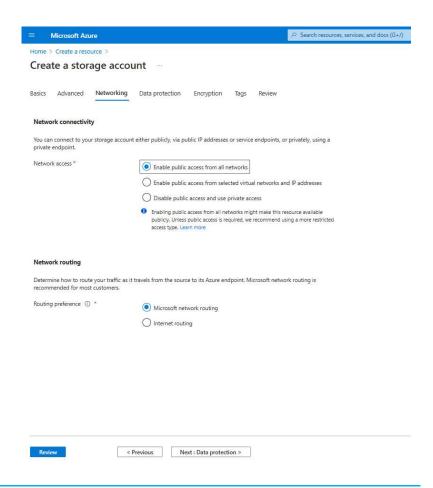
7. My settings are as screen capture.



 On the Networking page, configure as follow and then click Next: Data protection.

Section	Field	Required or optional	Description
Network connectivity	Connectivity method	Required	By default, incoming network traffic is routed to the public endpoint for your storage account. You can specify that traffic must be routed to the public endpoint through an Azure virtual network. You can also configure private endpoints for your storage account. For more information, see Use private endpoints for Azure Storage
Network routing	Routing preference	Required	The network routing preference specifies how network traffic is routed to the public endpoint of your storage account from clients over the internet. By default, a new storage account uses Microsoft network routing. You can also choose to route network traffic through the POP closest to the storage account, which may lower networking costs. For more information, see Network routing preference for Azure Storage.

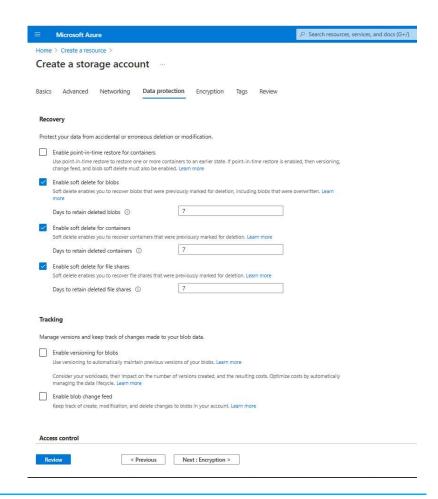
9. My settings are as screen capture.



10. On the Data protection page, configure as follow and then click Next:
Encryption.

Section	Field	Required or optional	Description
Recovery	Enable point-in- time restore for containers	Optional	Point-in-time restore provides protection against accidental deletion or corruption by enabling you to restore block blob data to an earlier state. For more information, see Point-in-time restore for block blobs.
			Enabling point-in-time restore also enables blob versioning, blob soft delete, and blob change feed. These prerequisite features may have a cost impact. For more information, see Pricing and billing for point-in-time restore.
Recovery	Enable soft delete for blobs	Optional	Blob soft delete protects an individual blob, snapshot, or version from accidental deletes or overwrites by maintaining the deleted data in the system for a specified retention period. During the retention period, you can restore a soft-deleted object to its state at the time it was deleted. For more information, see Soft delete for blobs.
			Microsoft recommends enabling blob soft delete for your storage accounts and setting a minimum retention period of seven days.
Recovery	Enable soft delete for containers	Optional	Container soft delete protects a container and its contents from accidental deletes by maintaining the deleted data in the system for a specified retention period. During the retention period, you can restore a soft-deleted container to its state at the time it was deleted. For more information, see Soft delete for containers (preview).
			Microsoft recommends enabling container soft delete for your storage accounts and setting a minimum retention period of seven days.
STATE OF THE PARTY	Enable soft delete for file shares	Optional	Soft delete for file shares protects a file share and its contents from accidental deletes by maintaining the deleted data in the system for a specified retention period. During the retention period, you can restore a soft-deleted file share to its state at the time it was deleted. For more information, see Prevent accidental deletion of Azure file shares.
			Microsoft recommends enabling soft delete for file shares for Azure Files workloads and setting a minimum retention period of seven days.
Tracking	Enable versioning for blobs	Optional	Blob versioning automatically saves the state of a blob in a previous version when the blob is overwritten. For more information, see Blob versioning.
			Microsoft recommends enabling blob versioning for optimal data protection for the storage account.
Tracking	Enable blob change feed	Optional	The blob change feed provides transaction logs of all changes to all blobs in your storage account, as well as to their metadata. For more information, see Change feed support in Azure Blob Storage.
Access	Enable version- level immutability support	Optional	Enable support for immutability policies that are scoped to the blob version. If this option is selected, then after you create the storage account, you can configure a default time-based retention policy for the account or for the container, which blob versions within the account or container will inherit by default. For more information, see Enable version-level immutability support on a storage account.

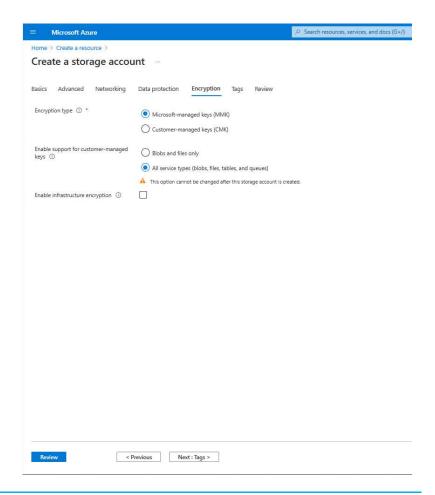
11. In my case, my settings are as screen capture.



12. On the Encryption page, configure and click Next: Tags.

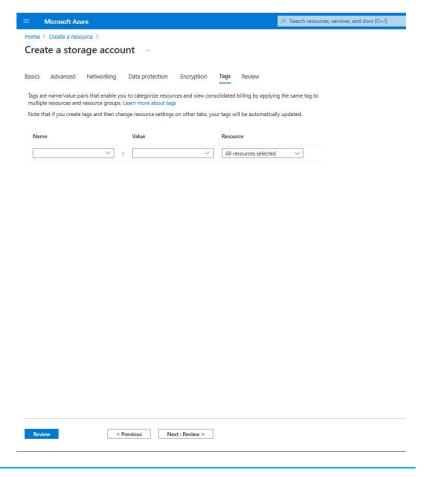
Field	Required or optional	Description
Encryption type	Required	By default, data in the storage account is encrypted by using Microsoft-managed keys. You can rely on Microsoft-managed keys for the encryption of your data, or you can manage encryption with your own keys. For more information, see Azure Storage encryption for data at rest.
Enable support for customer- managed keys	Required	By default, customer managed keys can be used to encrypt only blobs and files. Set this option to All service types (blobs, files, tables, and queues) to enable support for customer-managed keys for all services. You are not required to use customer-managed keys if you choose this option. For more information, see Customer-managed keys for Azure Storage encryption.
Encryption key	Required if Encryption type field is set to Customer- managed keys.	If you choose Select a key vault and key, you are presented with the option to navigate to the key vault and key that you wish to use. If you choose Enter key from URI, then you are presented with a field to enter the key URI and the subscription.
User-assigned identity	Required if Encryption type field is set to Customer- managed keys.	If you are configuring customer-managed keys at create time for the storage account, you must provide a user-assigned identity to use for authorizing access to the key vault.
Enable infrastructure encryption	Optional	By default, infrastructure encryption is not enabled. Enable infrastructure encryption to encrypt your data at both the service level and the infrastructure level. For more information, see Create a storage account with infrastructure encryption enabled for double encryption of data.

13. In my case, my settings are as screen capture.

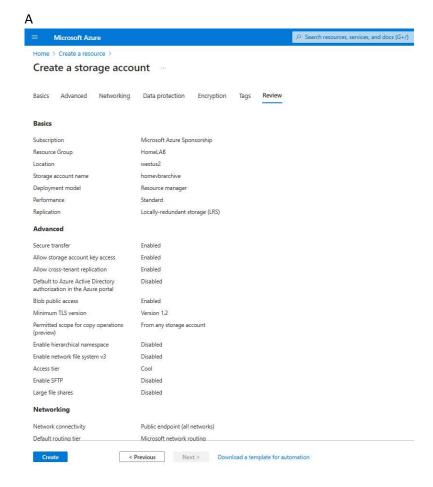


14. On the Tags page, configure as follow and click Next: Review.

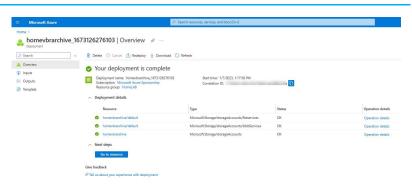
You can specify Resource Manager tags on the Tags tab to help organize your Azure resources.



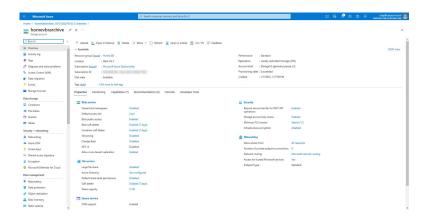
15. On the Review page, click Create.



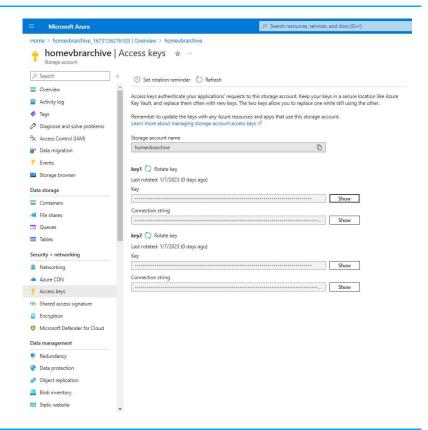
16. Creating the new storage account and clicking Go to the resource may take a few minutes.



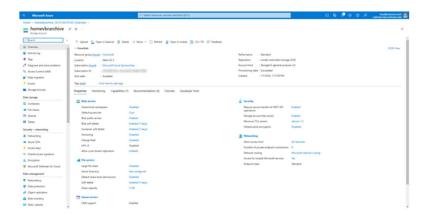
17. On the newly created Storage account page, select Access keys.



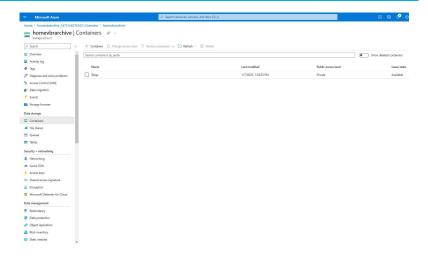
18. On the Access keys page, select Show keys and copy the Storage account name and key of key1. We need them for Veeam storage repository settings later.



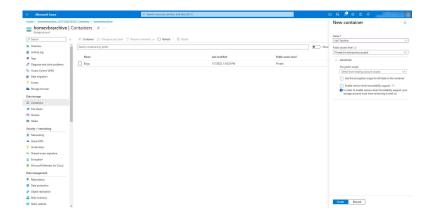
19. On the newly created Storage account page, select Containers.



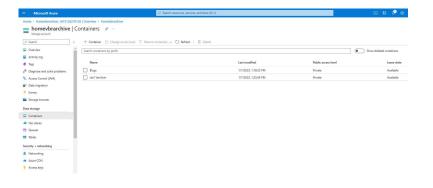
20. On the Containers page, click +Container.



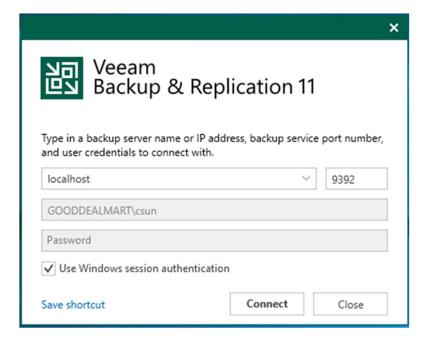
- 21. On the new container page, enter a name for your new container, select Private (no anonymous access) as Public access level and then click Create.
- 22. Veeam Backup and Replication v11 do not support Immutable storage with versioning for azure blob storage.



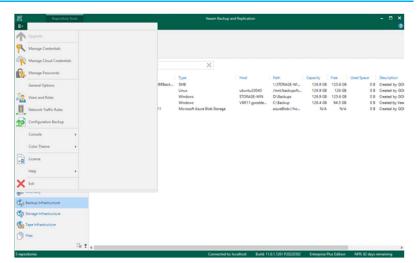
23. Verify the new container created.



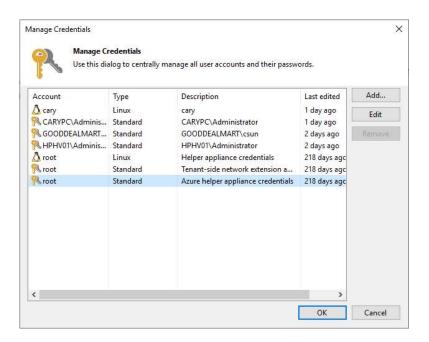
- 24. Log in to the Veeam Backup and replication manager server.
- 25. Open the Veeam Backup & Replication Console, and click Connect.



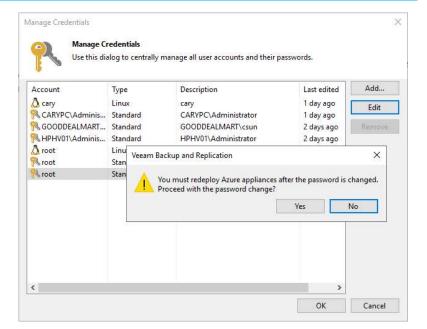
- 26. Veeam Backup &
 Replication uses its builtin credentials record to
 work with all helper
 appliances in Microsoft
 Azure and AzuStack Hub
 by default. I recommend
 changing the password
 for this credentials record
 before you set up helper
 appliances.
- 27. Select Manage
 Credentials from the main menu.



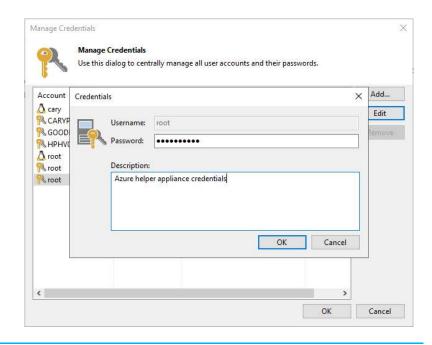
28. Select the built-in credentials record for the Azure helper appliances on the Manage Credentials page and click Edit.



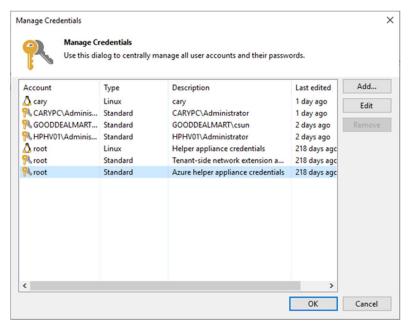
- 29. It would help if you redeployed all existing helper appliances in Microsoft Azure and Azure Stack Hub after you change the password in the built-in credentials record. To redeploy appliances, you must remove all configured appliances and then configure them again.
- 30. Click Yes at the warning message.



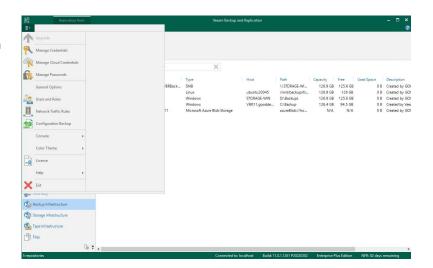
31. On the Credentials page, enter the new password and click OK.



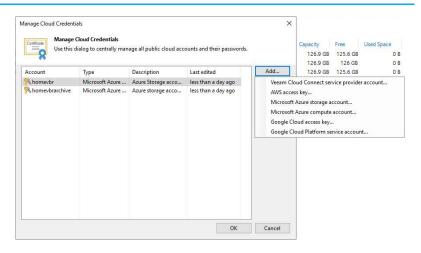
32. On the Manage Credentials page, click OK.



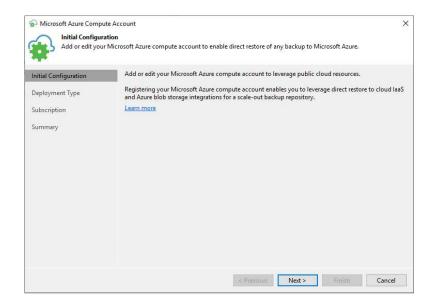
33. Select Manage Cloud Credentials from the main menu.



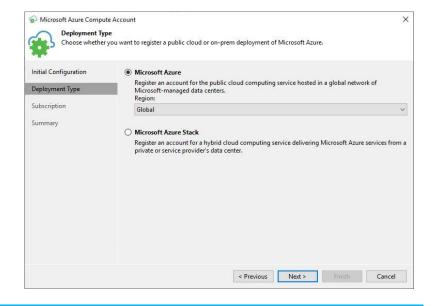
34. Click Add on the Manage Cloud Credentials page and select Microsoft Azure compute account.



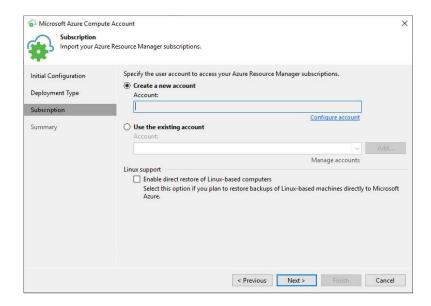
35. On the Initial
Configuration page, click
Next.



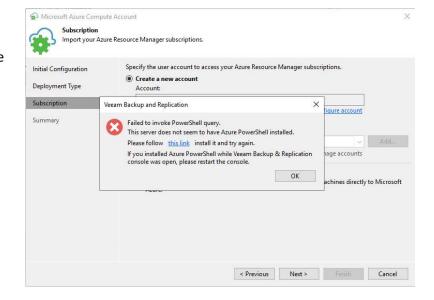
36. On the Deployment Type page, select Microsoft Azure, select a Microsoft Azure region from the Region drop-down list, and click Next.



37. On the Subscription page, select a new account and click Configure account.



- 38. Microsoft Azure
 PowerShell is installed on
 the machine that runs the
 Veeam Backup &
 Replication console. The
 Veeam Backup &
 Replication will display a
 warning if Microsoft
 Azure PowerShell is not
 installed.
- 39. In the warning window, click this link.



40. On the Microsoft Azure PowerShell setup page, select I accept the terms in the License Agreement checkbox and click Install.



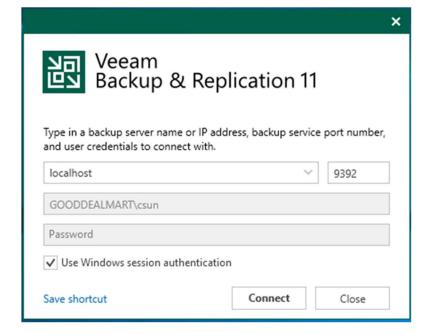
41. On the User Access Control page, click Yes.



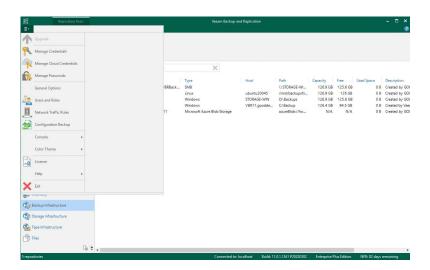
- 42. On the Microsoft Azure PowerShell setup page, click Finish.
- 43. Close the Veeam Backup & Replication console after completing the installation.
- 44. Reboot the machine to allow the Veeam Backup & Replication console to detect the newly installed version of Microsoft Azure PowerShell.



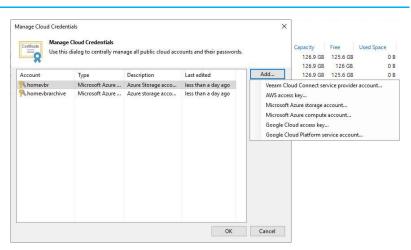
- 45. Log in to the Veeam Backup and replication manager server.
- 46. Open the Veeam Backup & Replication Console, and click Connect.



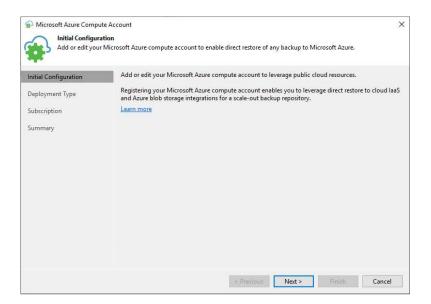
47. Select Manage Cloud Credentials from the main menu.



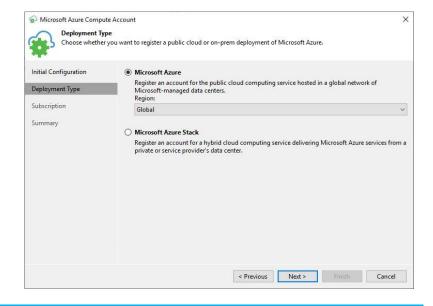
48. Click Add on the Manage Cloud Credentials page and select Microsoft Azure compute account.



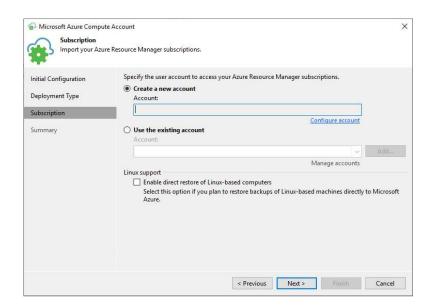
49. On the Initial
Configuration page, click
Next.



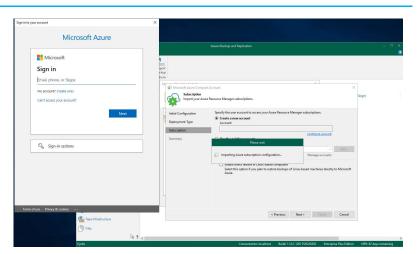
50. On the Deployment Type page, select Microsoft Azure, select a Microsoft Azure region from the Region drop-down list, and click Next.



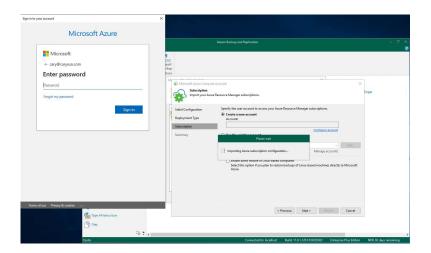
51. On the Subscription page, select a new account and click Configure account.



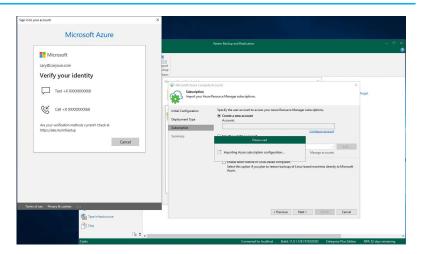
52. Enter the user account on the Microsoft Azure Signin page and click Next.



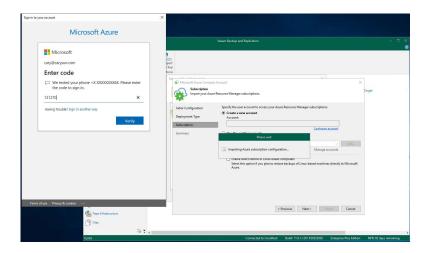
53. Type the password on the Enter password page and click Sign in.



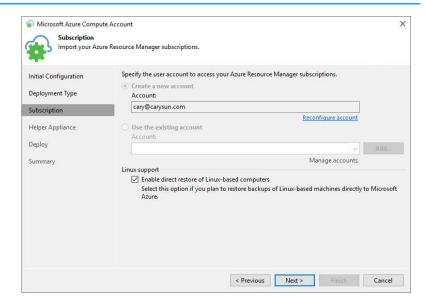
54. Select the verification method on the Verify your identity page.



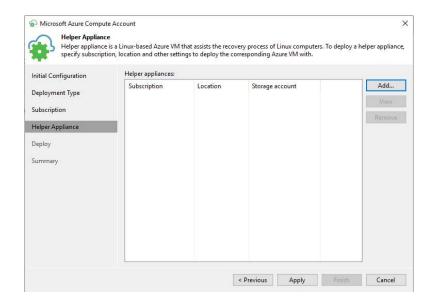
55. On the Enter code page, enter the texted code and click Verify.



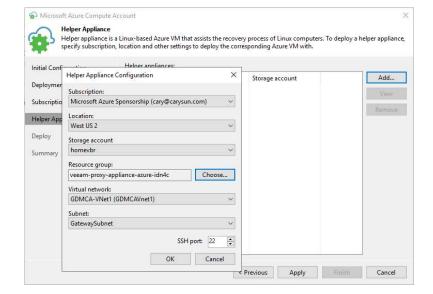
56. On the Subscription page, select the Enable restore of Linux-based computers checkbox to restore Linux-base computers, and click Next.



57. On the Helper Appliance page, click Add.

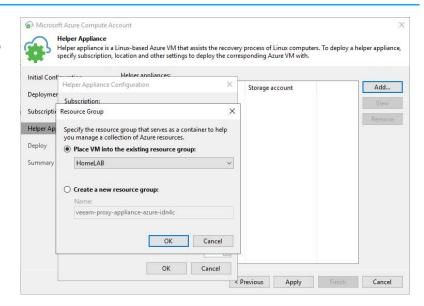


- 58. On the Helper Appliance Configuration page, specify the information for the new helper appliance.
- 59. Select an Azure subscription to configure the helper appliance.
- 60. Select a location where you want to configure a helper appliance.
- 61. Select a storage account whose resources you want to use to store the disks of the helper appliance.
- 62. Click Choose and choose the resource if you do not

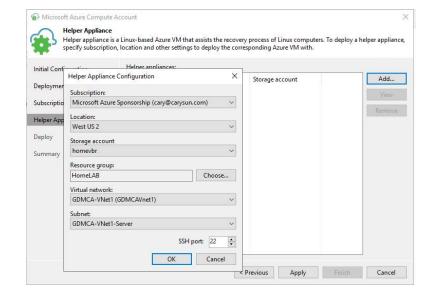


want Veeam Backup & Replication to create a new resource group.

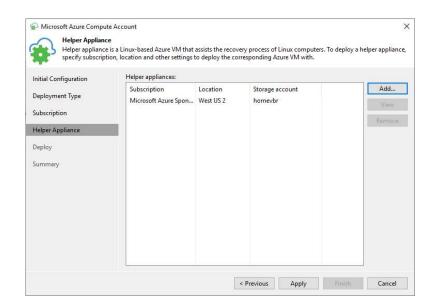
63. On the Resouce Group page, select Place VM into the existing resource group, choose the current resource group from the drop-down list and click OK.



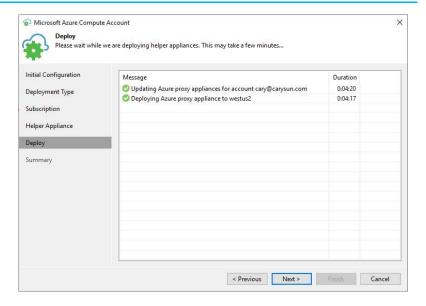
- 64. Select a network for a helper appliance that must be connected.
- 65. Select a subnet for the helper appliance.
- 66. Specify a port over which Veeam Backup & Replication will communicate with the helper appliance.
- 67. Click OK.



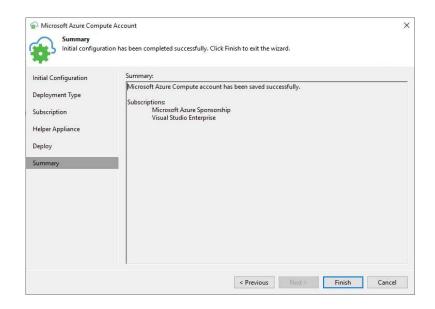
68. On the Helper Appliance page, click Apply.



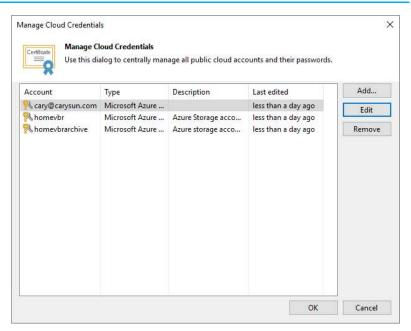
69. On the Deploy page, ensure all deployments are completed without issues, ad click Next.



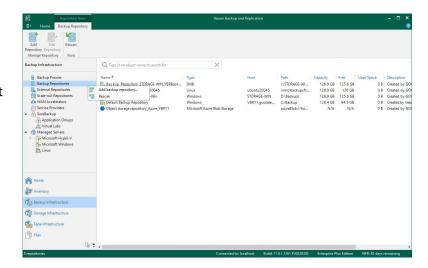
70. On the Summary page, click Finish.



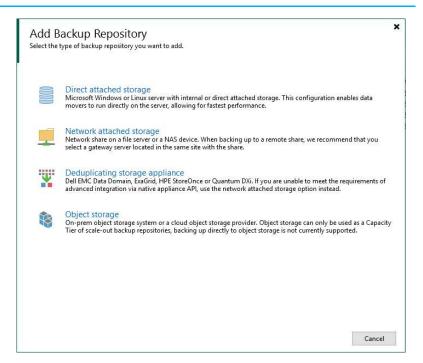
71. On the Manage Cloud Credentials page, click OK.



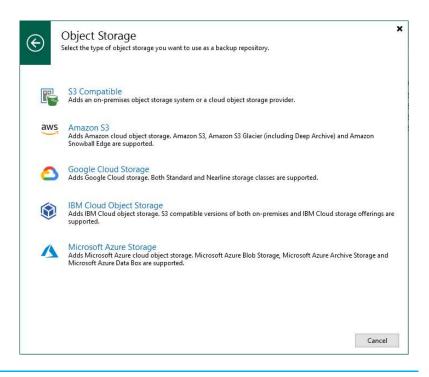
- 72. On the Home page, select Backup Infrastructure.
- 73. On the Backup
 Infrastructure page, select
 Backup Repositories,
 right-click Backup
 Repositories, and select
 Add backup repository.



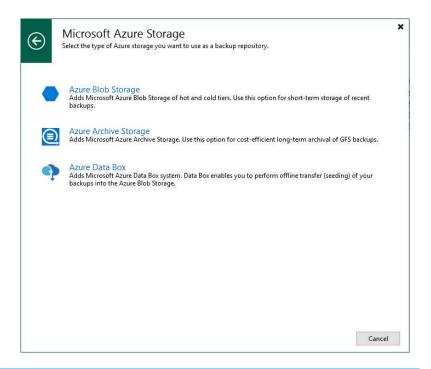
74. Select the Object storage on the Add backup Repository page.



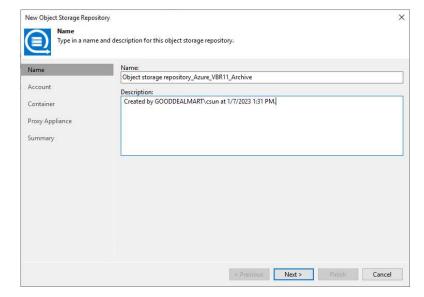
75. On the Object Storage page, select Microsoft Azure Storage.



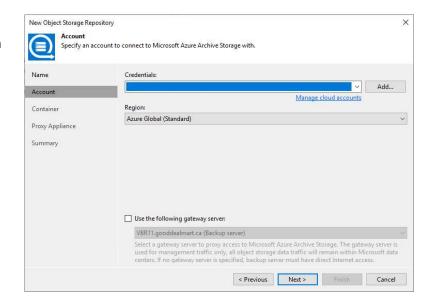
76. On the Microsoft Azure Storage page, select Azure Blob Storage.



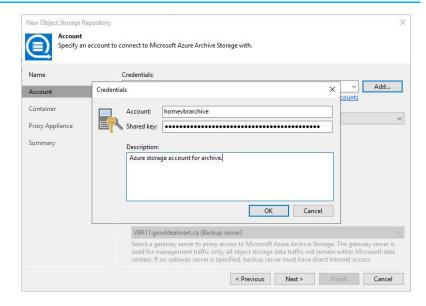
- 77. On the Name page, specify in the Name field.
- 78. In the Description field, describe future references.
- 79. Click Next.



80. Click Add to create a manage cloud account on the Account page.



81. On the Credentials page, paste the Azure storage account name as Account and paste the key1 as Shared key (you can find them from the Access key session of the Azure storage account), and click OK.



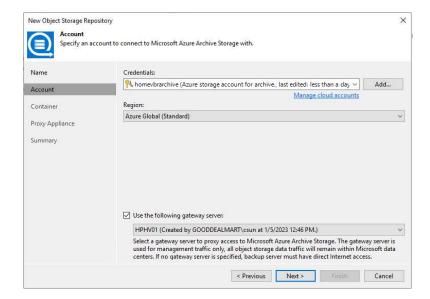
- 82. Select Azure Global (Standard) as Region on the Account page.
- 83. Select Use the following gateway server check box and choose a server from the list.

Note:

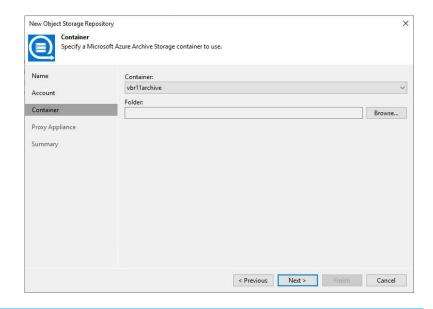
You can select any Microsoft Windows or Linux server added to your backup infrastructure with an internet connection, except the Linux server with the hardened repository role.

If you do not select the Use the following gateway server check box, you must make sure that the backup server has direct internet access

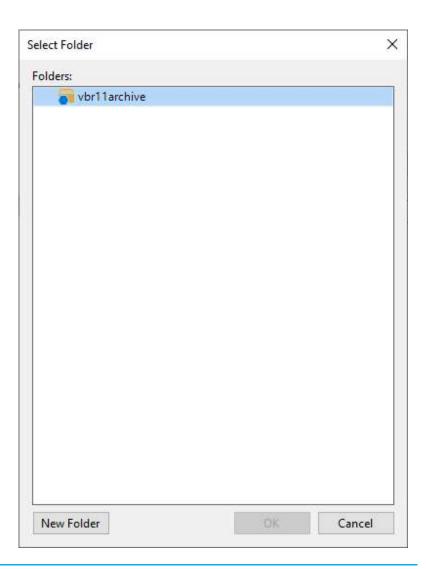
84. Click Next.



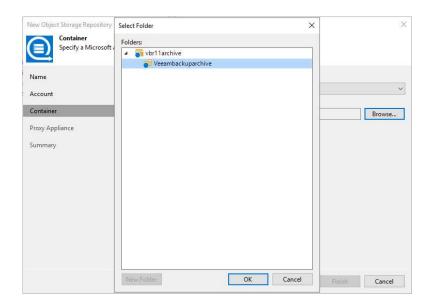
- 85. Select a container from the Container drop-down list on the Container page.
- 86. In the Folder field, click Browse.



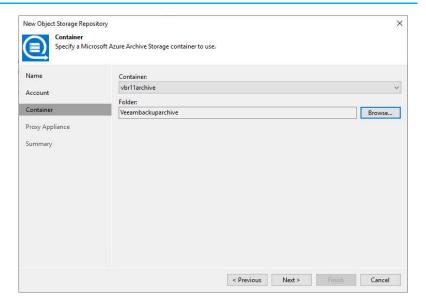
87. On the Folders page, select the container and click New Folder.



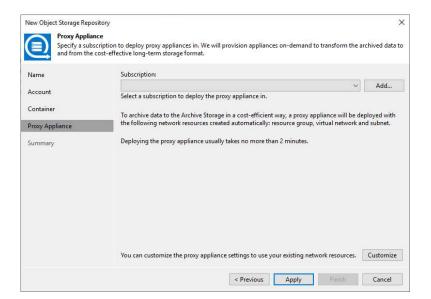
88. Enter the name for the new folder, select the folder and click OK.



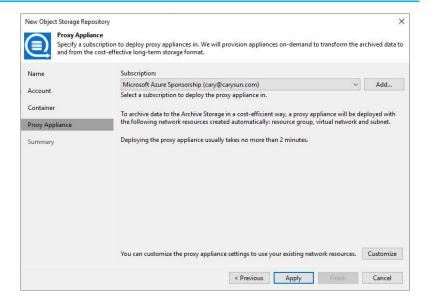
89. On the Container page, click Next.



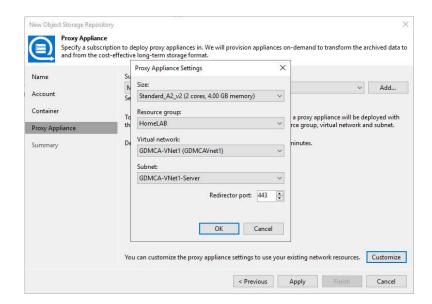
90. Select your Microsoft
Azure subscription
credentials from the
Subscription drop-down
list on the Proxy
Appliance page.



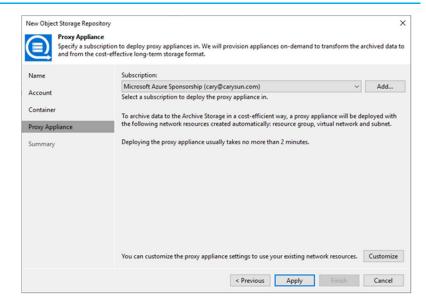
91. On the Proxy Appliance page, click Customise to customize the proxy appliance.



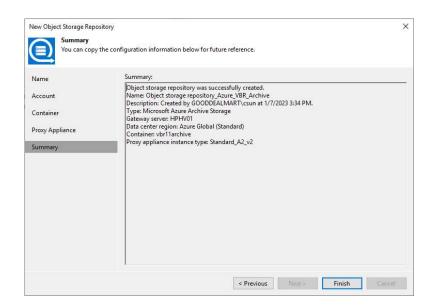
92. On the Proxy Appliance Settings page, specify the information for the proxy appliance and click OK.



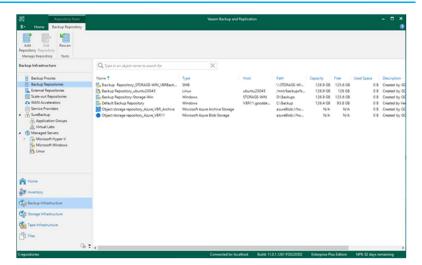
93. On the Proxy Appliance page, click Apply.



94. On the Summary page, click Finish.



95. Verify that the Backup Reposihasy has been added.



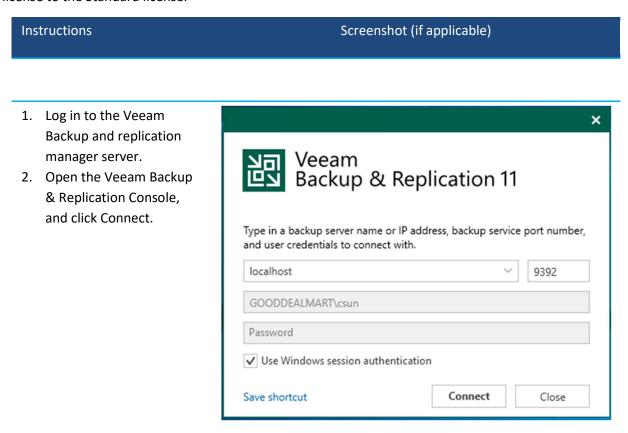
Adding Local Directory and Azure Blob Object Storage as Scaleout Repository without Archive Tier

A scale-out backup repository allows for horizontal scaling for multi-tier data storage.

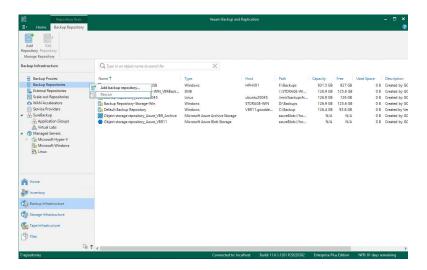
The scale-out backup repository comprises one or more backup repositories known as performance tiers, which object storage repositories can supplement for long-term and archive storage, known as capacity and archive tiers. All storage devices and systems within the scale-out backup repository are linked into a system, and their capacities are summarised.

A scale-out backup repository is included in the Veeam Universal License. An Enterprise or higher edition is required when using a legacy socket-based license.

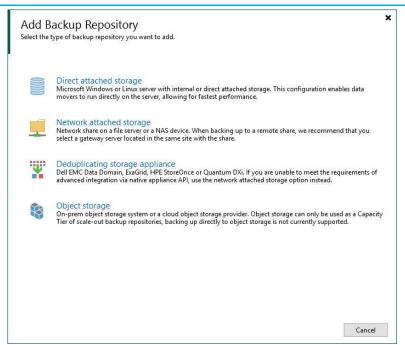
However, you can perform a restore from the scale-out backup repository after downgrading the license to the Standard license.



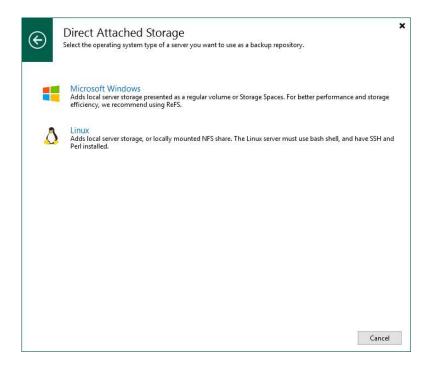
- 3. On the Home page, select Backup Infrastructure.
- 4. On the Backup
 Infrastructure page, select
 Backup Repositories,
 right-click Backup
 Repositories, and select
 Add backup repository.



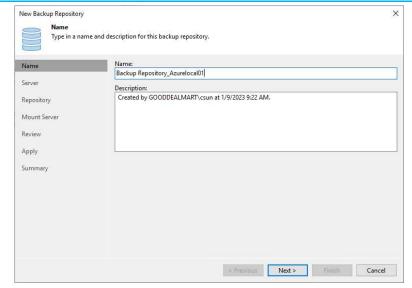
On the Add Backup
 Repository page, select
 Direct attached storage.



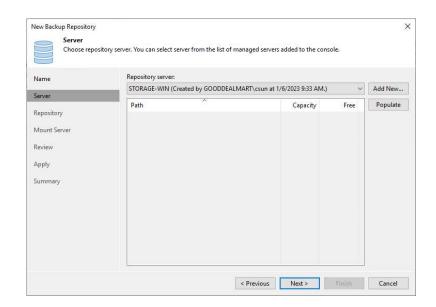
 On the Direct Attached Storage page, select Microsoft Windows.



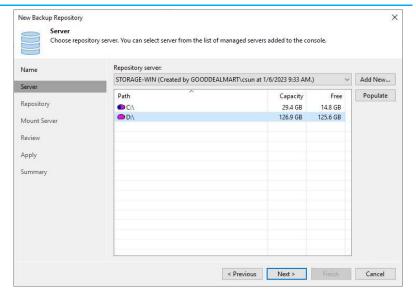
- 7. Specify the repository name in the Name field on the Name page.
- 8. In the Description field, describe future references.
- 9. Click Next.



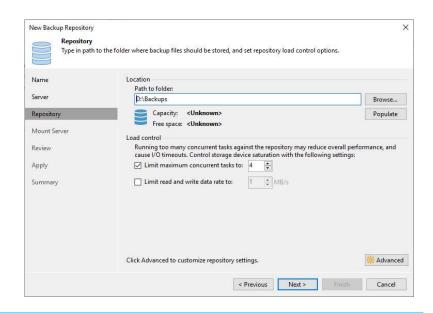
10. Select the repository server from the dropdown list on the Server page and click Populate.



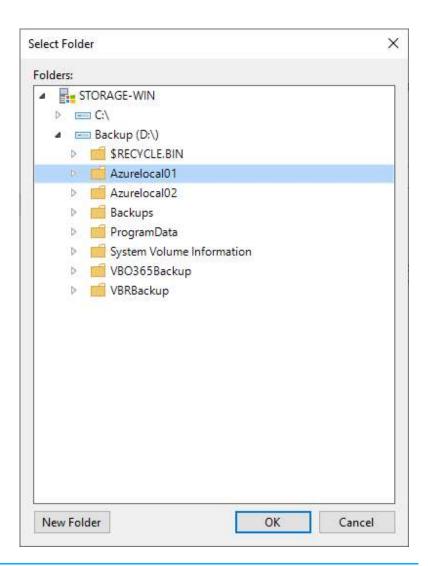
11. Select the disk as a backup repository, and click Next.



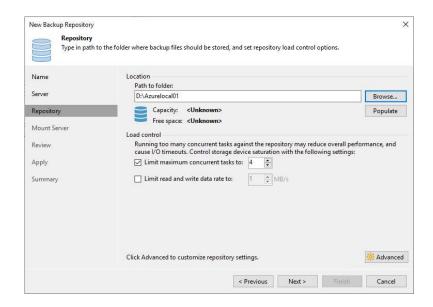
12. In the location session, click Browse.



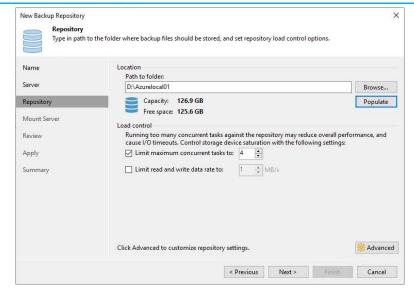
13. On the Select Folder page, select the folder ((or create a New folder) and click OK.



14. On the Repository page, click Populate to review the disk capacity and free space.



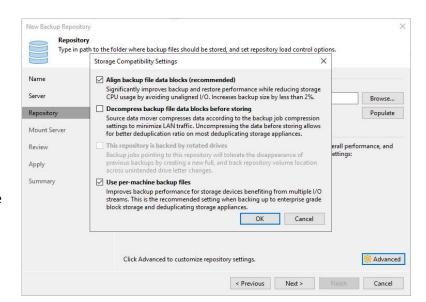
- 15. Use the Load control settings to control the load if you need it.
- 16. Click Advanced.



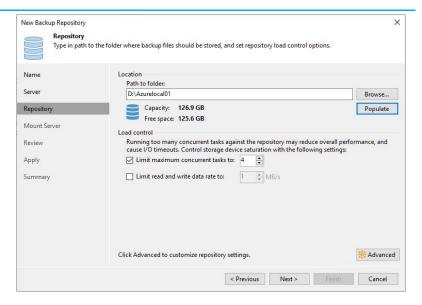
17. On the Storage
Compatibility Settings,
select Align backup file
data blocks, select Use
per-machine backup files,
and click OK.

Note:

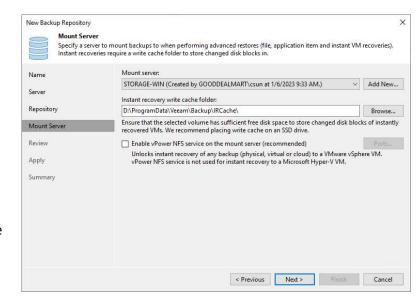
Select Decompress backup file data blocks before storing if you use a deduplicating storage feature or appliance.



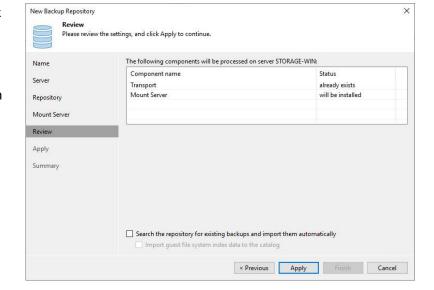
18. On the Repository page, click Next.



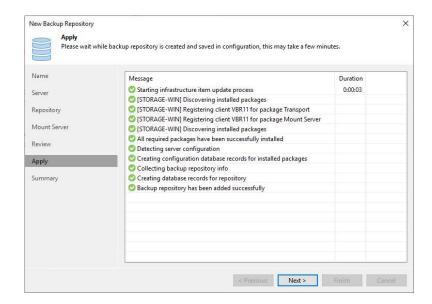
- 19. On the Mount Server page, select a server.
- 20. In the Instant recovery write cache folder field, select a folder used for writing cache during mount operations.
- 21. Unselect Enable vPower NFS service on the mount server, vPower NFS settings are not applicable in Microsoft Hyper-V environments.
- 22. Click Next.



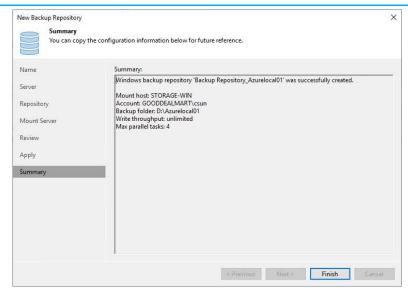
- 23. On the Review page, click Apply.
- 24. Select the Search the repository for existing backups and import them automatically if the backup repository contains backups previously created with Veeam Backup & Replication.
- 25. Select the Import guest file system index checkbox.



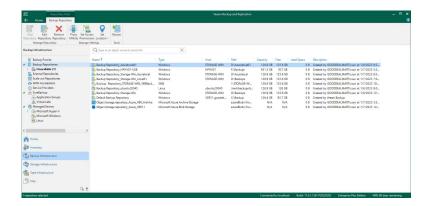
26. On the Apply page, complete adding the backup repository without error, and click Next.



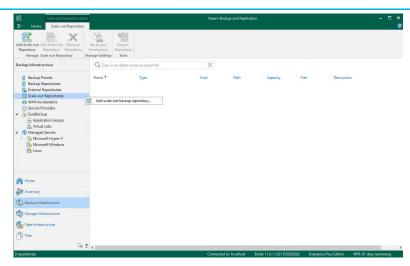
27. On the Summary page, click Finish.



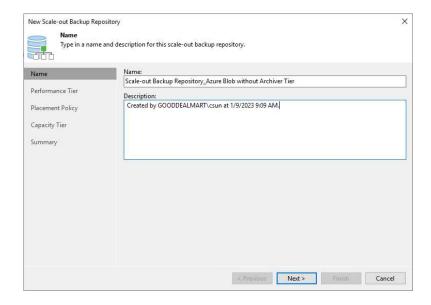
28. Verify that the Backup Repository has been added.



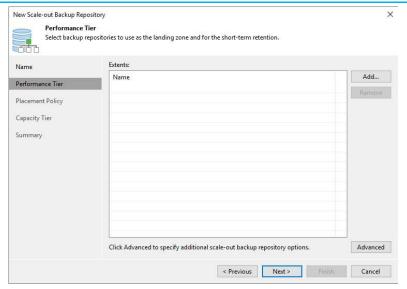
- 29. On the Backup Infrastructure, select Scale-out Repositories.
- 30. Right-click Scale-out
 Repositories and select
 Add scale-out backup
 repository.



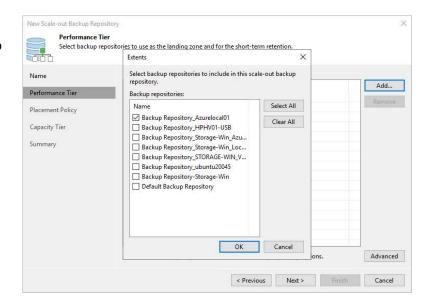
- 31. Specify the scale-out repository name in the Name field on the Name page.
- 32. In the Description field, describe future references.
- 33. Click Next.



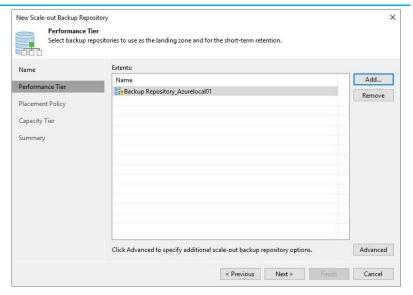
34. On the Performance Tier page, click Add.



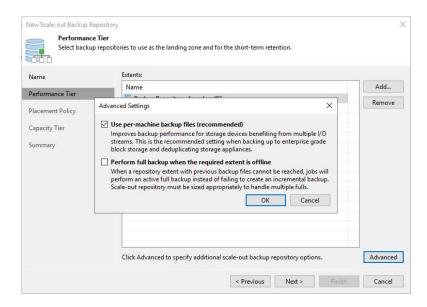
35. On the Extents page, select check boxes next to the backup repositories you want to add as performance extents and click OK.



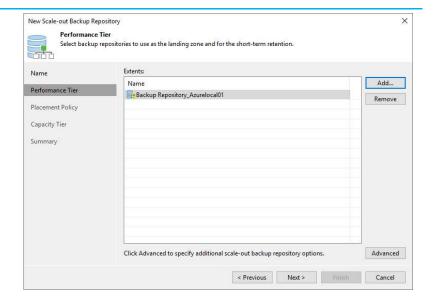
36. On the Performance Tier page, click Advanced.



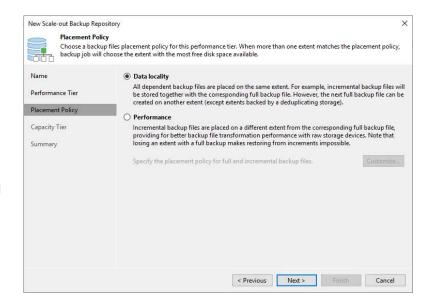
- 37. On the Advanced Settings page, Specify options for the scale-out backup repository and click OK.
- 38. Select the Use permachine backup files checkbox to create a separate backup chain for every machine in the job.
- 39. Select the Perform full backup when necessary as an offline check box to preserve the consistency of backup chains in the scale-out backup repository.



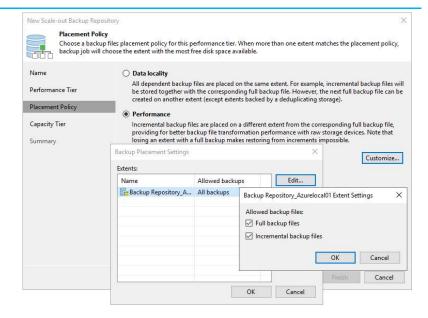
40. On the Performance Tier page, click Next.



- 41. On the Placement Policy page, select one of policy and click Next.
- 42. Select Data locality to store backup files that belong to the same backup chain.
- 43. Select Performance to store full and incremental backup files to different performance extents.



- 44. You can select
 Performance policy if you
 have two extents at
 Performance Tier.
- 45. Select Customize, and click Edit to choose the type of backup files.

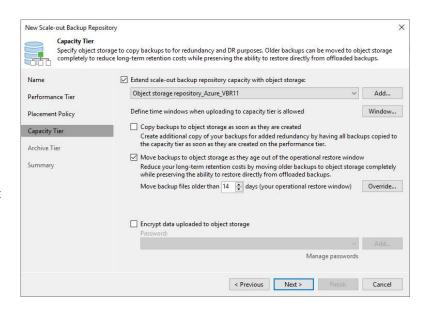


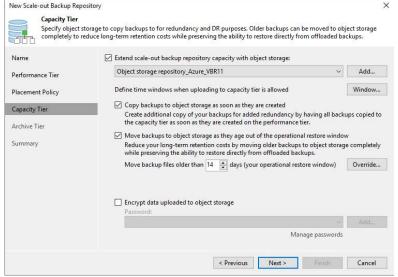
- 46. Select the Extend scaleout backup repository capacity on the Capacity Tier page with the object storage check box.
- 47. Select an object storage repository from the dropdown list. If the object storage repository has not been configured, click Add to configure it.
- 48. Click Windows to define time windows when uploading to capacity tier is allowed if your internet bandwidth is not good enough.
- bandwidth is not good enough.

 49. Select Copy backups to object storage as soon as they are created; checkbox to copy new backups as soon as they are made. It will create an additional copy of your
- 50. Select Move backups to object storage as they age out of the operational restores window checkbox to move the inactive backup chains to the capacity extent.

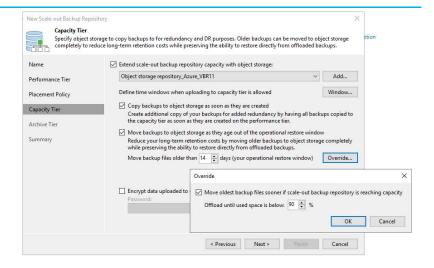
backups for added

redundancy.

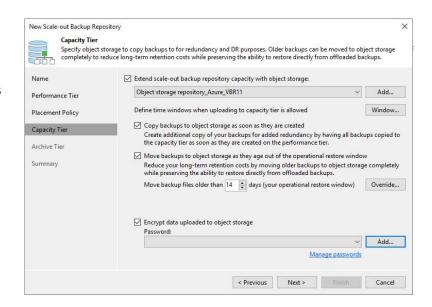




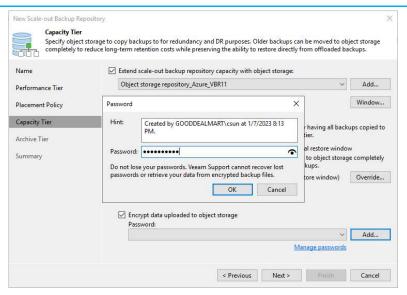
- 51. In Move backup files older than the X days field, select the operational restore window to specify the time after which inactive backup chains on your performance extents will be considered outdated and should be moved to the capacity extent. Consider "0" as an acceptable value for offloading the inactive backup chains on the same day they are created.
- 52. Click Override to override the behaviour of moving old backups
- 53. Select Move oldest backup files sooner if the scale-out backup repository reaches the capacity checkbox. Enter a percentage threshold to force data transfer if a scale-out backup repository reaches the specified threshold.



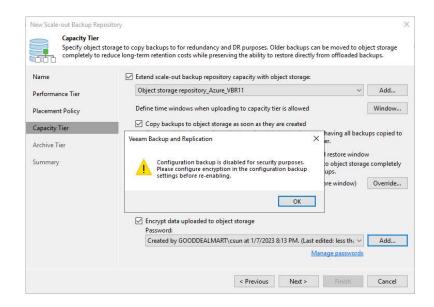
54. Select Encrypt data uploaded to object storage checkbox. The entire collection of blocks and metadata will be encrypted while being offloaded.



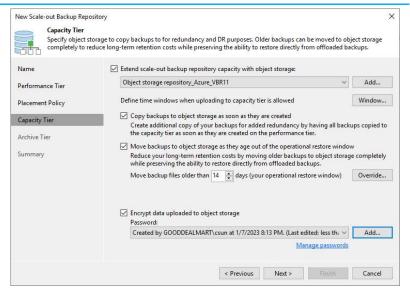
55. Click Add to create a password and click OK.



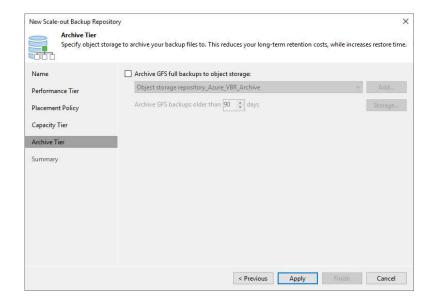
56. You need to enable encryption for your configuration backup.
Click OK at configuration backup is a disabled warning message.



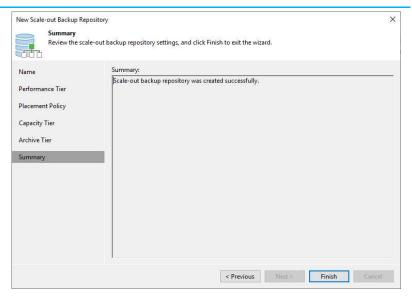
57. On the Capacity Tier page, click Next.



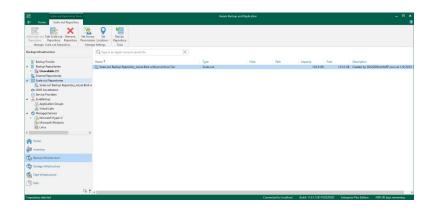
58. On the Archive Tier page, click Apply.



59. On the Summary page, click Finish.



60. Verify that the Scale-out Backup Repository has been added.

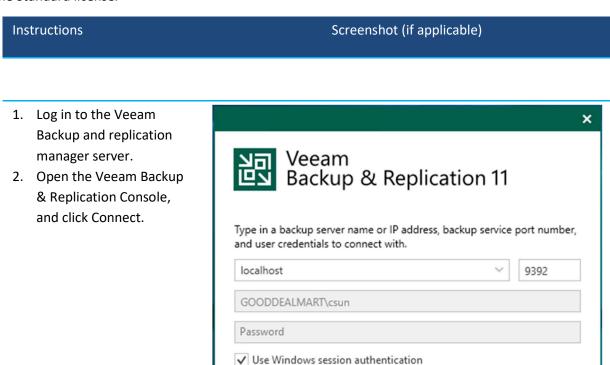


Adding Local Directory and Azure Blob Object Storage as Scaleout Repository with Archive Tier

A scale-out backup repository allows for horizontal scaling for multi-tier data storage.

The scale-out backup repository comprises one or more backup repositories known as performance tiers, which object storage repositories can supplement for long-term and archive storage, known as capacity and archive tiers. All the storage devices and systems within the scale-out backup repository are linked into a system, and their capacities are summarized. A scale-out backup repository is included in the Veeam Universal License. An Enterprise or higher edition is required when using a legacy socket-based license.

However, you can perform a restore from the scale-out backup repository after downgrading to the Standard license.

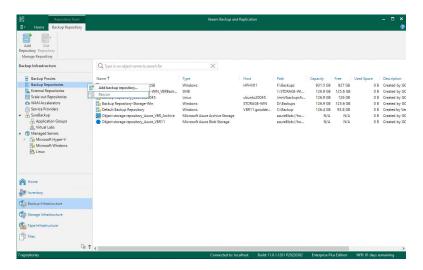


Save shortcut

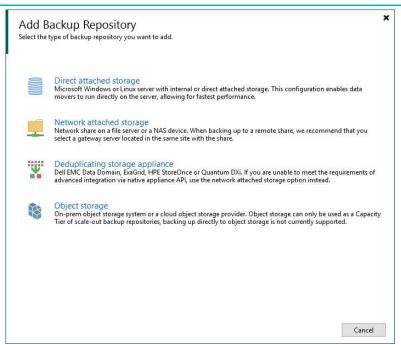
Connect

Close

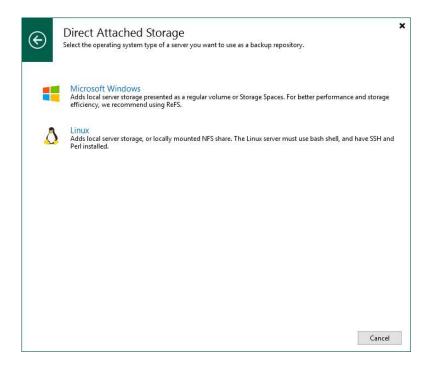
- 3. On the Home page, select Backup Infrastructure.
- 4. On the Backup
 Infrastructure page, select
 Backup Repositories,
 right-click Backup
 Repositories, and select
 Add backup repository.



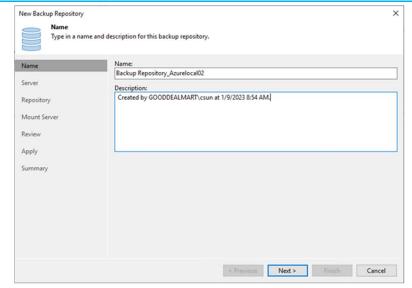
On the Add Backup
 Repository page, select
 Direct attached storage.



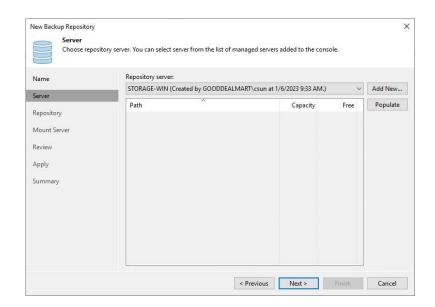
 On the Direct Attached Storage page, select Microsoft Windows.



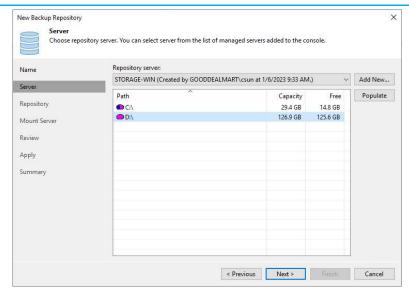
- 7. Specify the repository name in the Name field on the Name page.
- 8. In the Description field, describe future references.
- 9. Click Next.



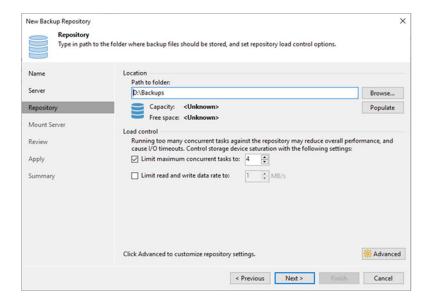
10. Select the repository server from the dropdown list on the Server page and click Populate.



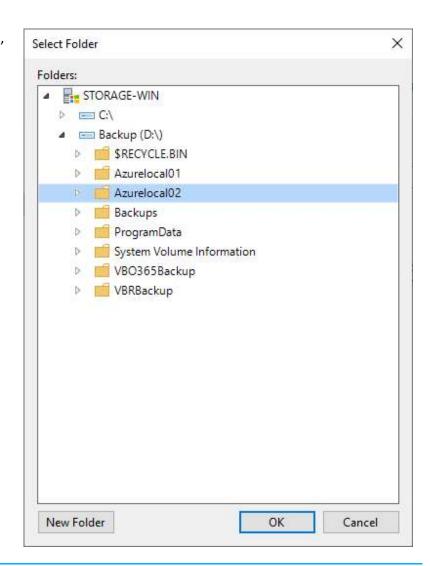
11. Select the disk and click Next.



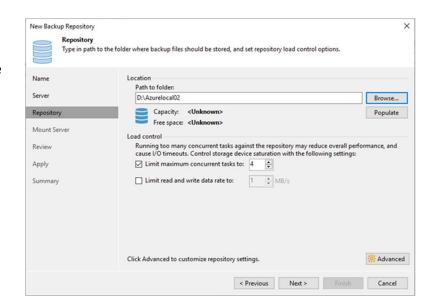
12. In the location session, click Browse.



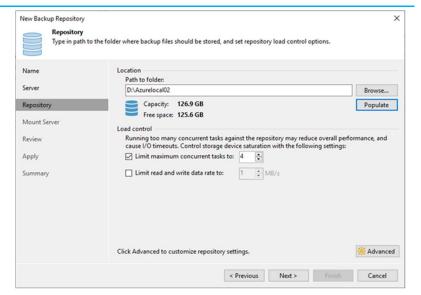
13. On the Select Folder page, select the folder ((or create a New folder) and click OK.



14. On the Repository page, click Populate to review the disk capacity and free space.



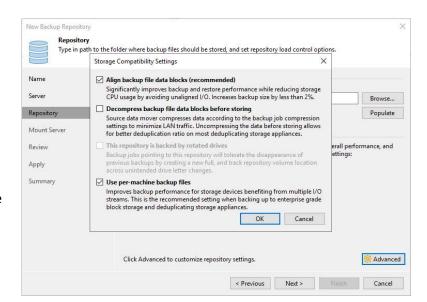
- 15. Use the Load control settings to control the load. If you need it
- 16. Click Advanced.



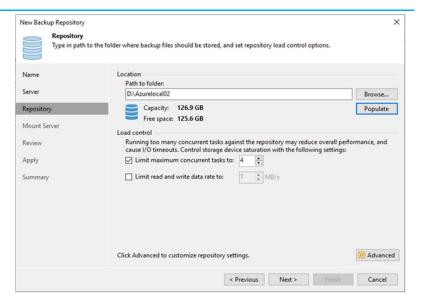
17. On the Storage
Compatibility Settings,
select Align backup file
data blocks, select Use
per-machine backup files,
and click OK.

Note:

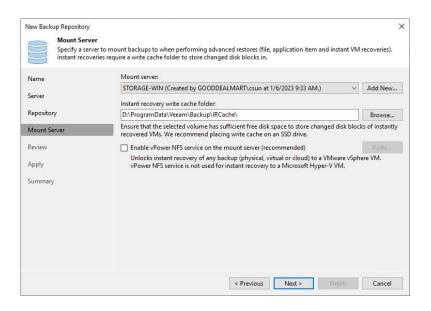
Select Decompress backup file data blocks before storing if you use a deduplicating storage feature or appliance.



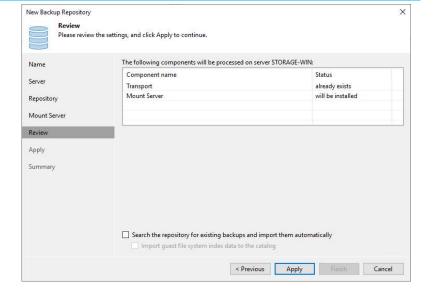
18. On the Repository page, click Next.



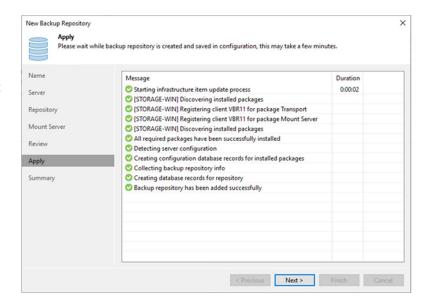
- 19. On the Mount Server page, Specify a server you want to use as a mount server. The mount server is required for file-level and application items restoration.
- 20. In the Instant recovery write cache folder field, specify a folder that will be used for writing cache during mount operations.
- 21. Unselect Enable vPower NFS service on the mount server, vPower NFS settings are not applicable in Microsoft Hyper-V environments.
- 22. Click Next.



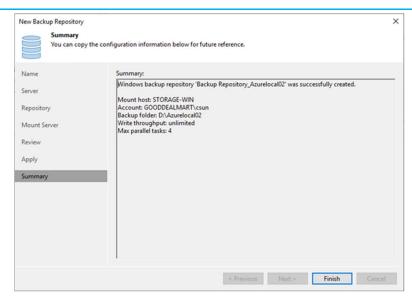
- 23. On the Review page, click Apply.
- 24. Select the Search the repository for existing backups and import them automatically if the backup repository contains backups previously created with Veeam Backup & Replication.
- 25. Select the Import guest file system index.



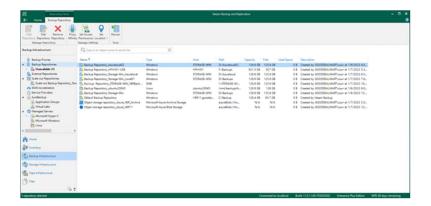
26. On the Apply page, ensure you complete the procedure of the backup repository adding without error, and click Next.



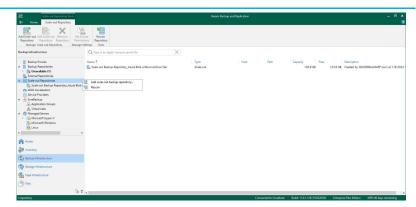
27. On the Summary page, click Finish.



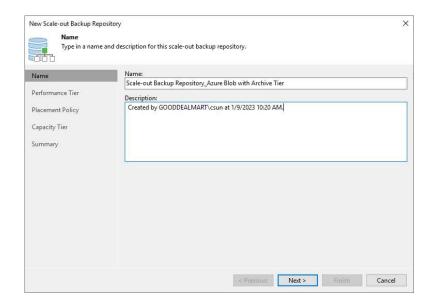
28. Verify that the Backup Repository has been added.



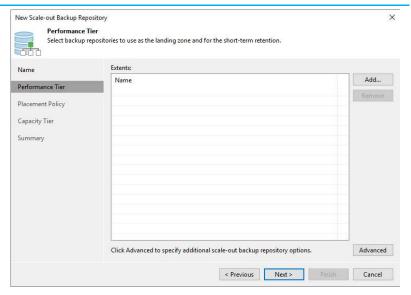
- 29. On the Backup Infrastructure, select Scale-out Repositories.
- 30. Right-click Scale-out Repositories, and select Add scale-out backup repository.



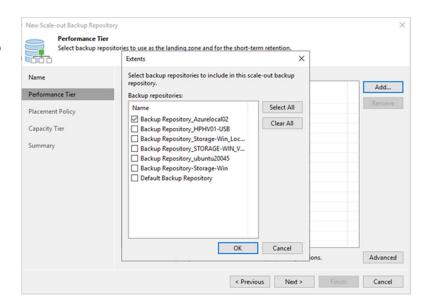
- 31. Specify the scale-out repository name in the Name field on the Name page.
- 32. In the Description field, describe future references.
- 33. Click Next.



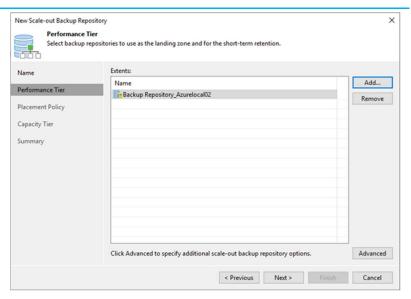
34. On the Performance Tier page, click Add.



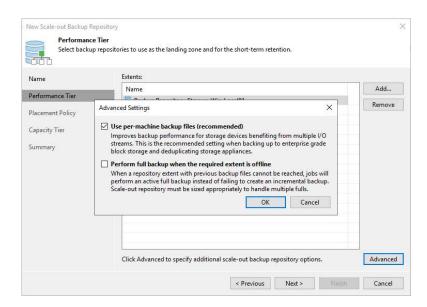
35. On the Extents page, select check boxes next to the backup repositories you want to add as performance extents and click OK.



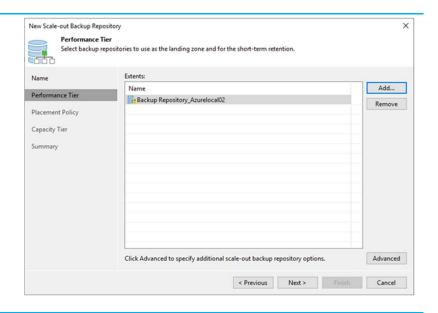
36. On the Performance Tier page, click Advanced.



- 37. On the Advanced Settings page, Specify options for the scale-out backup repository and click OK.
- 38. Select the Use permachine backup files checkbox to create a separate backup chain for every machine in the job.
- 39. Select the Perform full backup when necessary as an offline check box to preserve the consistency of backup chains in the scale-out backup repository.



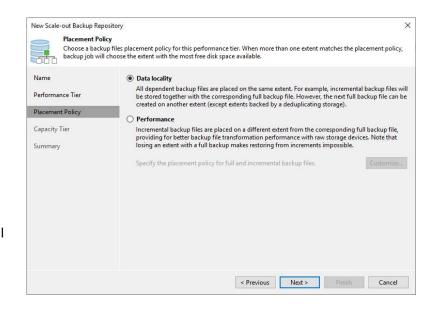
40. On the Performance Tier page, click Next.

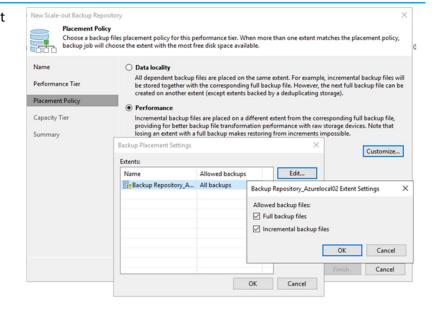


- 41. Select one of the policies for the scale-out backup repository on the Placement Policy page and click Next.
- 42. Select Data locality to store backup files that belong to the same backup chain.
- 43. Select Performance to store full and incremental backup files to different performance extents of the scale-out backup repository.
- repository.

 44. If you have two extents at Performance Tier, select Performance policy, select Customize, and click Edit to choose the type of backup files you want to store on the

extent.

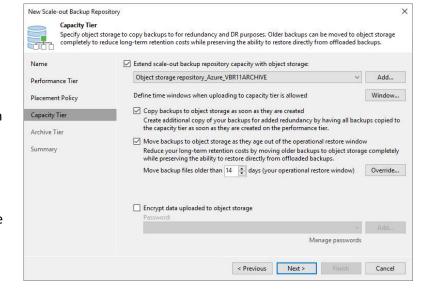




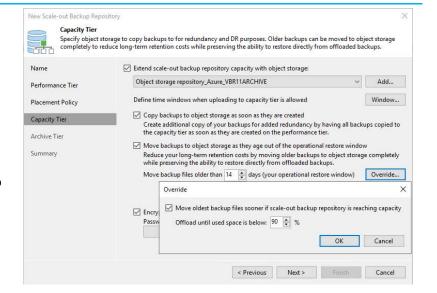
- 45. Select the Extend scaleout backup repository capacity on the Capacity Tier page with the object storage check box.
- 46. Select an object storage repository from the drop-down list. If the object storage repository has not been configured, click Add to configure it.
- 47. Click Windows to define time windows when uploading to capacity tier is allowed if your internet bandwidth is not good enough.
- Capacity Tier Specify object storage to copy backups to for redundancy and DR purposes. Older backups can be moved to object storage completely to reduce long-term retention costs while preserving the ability to restore directly from offloaded backups. Name Extend scale-out backup repository capacity with object storage Object storage repository_Azure_VBR11ARCHIVE ∨ Add... Performance Tier Define time windows when uploading to capacity tier is allowed Window... Placement Policy $\hfill \Box$ Copy backups to object storage as soon as they are created Capacity Tier Create additional copy of your backups for added redundancy by having all backups copied to the capacity tier as soon as they are created on the performance tier. Archive Tier Move backups to object storage as they age out of the operational restore window Summary Reduce your long-term retention costs by moving older backups to object storage completely while preserving the ability to restore directly from offloaded backups. Move backup files older than 14 💂 days (your operational restore window) Override... ☐ Encrypt data uploaded to object storage Manage passwords < Previous Next > Finish Cancel

New Scale-out Backup Repository

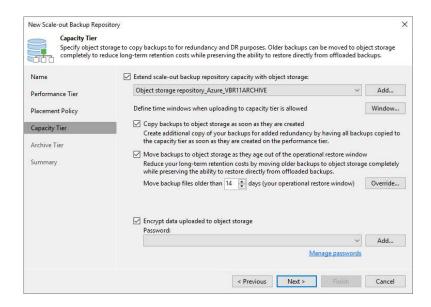
- 48. Select Copy backups to object storage as soon as they are created; checkbox to copy new backups as soon as they are made. It will create an additional copy of your backups for added redundancy.
- 49. Select Move backups to object storage as they age out of the operational restores window checkbox to move the inactive backup chains to the capacity extent.



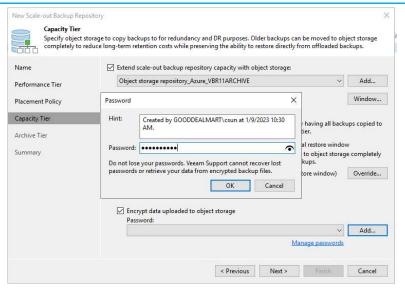
- 50. In Move backup files older than X days field, set the operational restore window to specify the time after inactive backup chains on your performance extents will be considered outdated and should be moved to the capacity extent. Consider "0" as an acceptable value for offloading the inactive backup chains on the same day they are created.
- 51. Click Override to override the behaviour of moving old backups
- 52. Select Move oldest backup files sooner if the scale-out backup repository reaches the capacity checkbox. Enter a percentage threshold to force data transfer if a scale-out backup repository reaches the specified threshold.



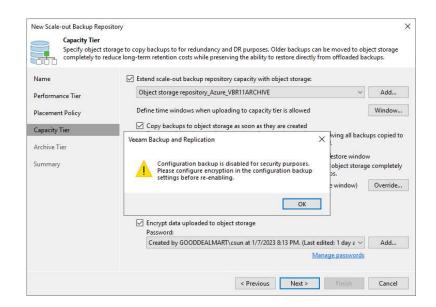
53. Select Encrypt data uploaded to object storage checkbox. The entire collection of blocks with the metadata will be encrypted while being offloaded.



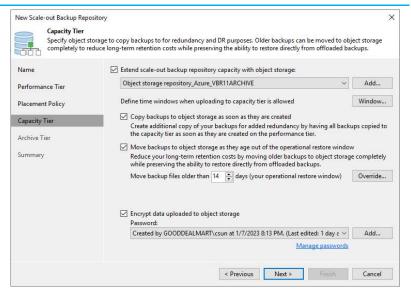
54. Click Add to create a password and click OK.



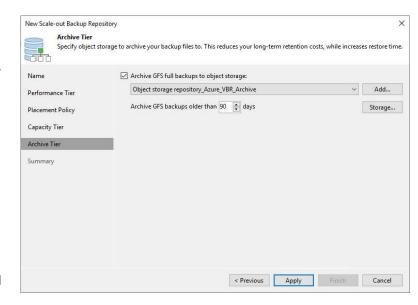
55. You need to enable encryption for your configuration backup.
Click OK at configuration backup is a disabled warning message.



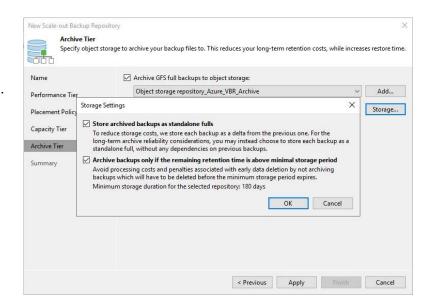
56. On the Capacity Tier page, click Next.



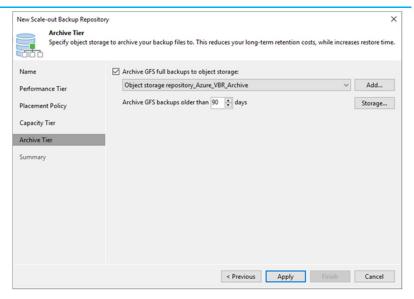
- 57. On the Archive Tier page, select Archive GFS full backups to object storage, choose one of the available archive object storage repositories or click Add to add a new one.
- 58. In the Archive GFS backups older than the number days field, specify the operational restore window to define a period after which inactive backup chains on your capacity extent will be considered outdated and should be moved to the archive extent. Consider the value "0" acceptable for archiving inactive backup chains on the same day they are created.
- 59. You can use the default storage settings or click Storage to specify them manually.



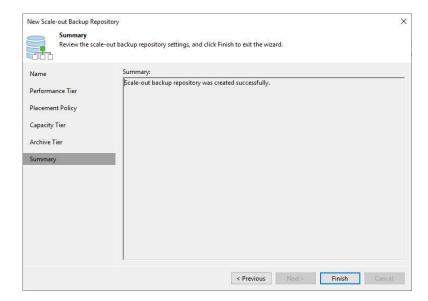
- 60. Select the Store archived backups as a standalone fulls checkbox to forbid the reusing of data blocks.
- 61. Select Archive backups only if the remaining retention time is above the minimum storage period checkbox to transport data blocks to the archive tier.



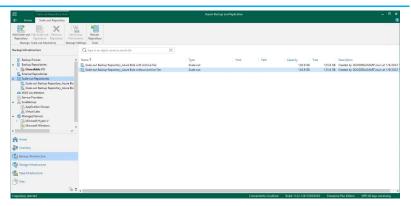
62. On the Archive Tier page, click Apply.



63. On the Summary page, click Finish.



64. Verify that the Scale-out Backup Repository has been added.



Chapter 2

Cloud Backup and Backup Copy

Cloud Backup and Backup Copy are two types of data backup solutions organizations can use to protect their data from loss due to unexpected events such as hardware failure, cyber-attacks, and natural disasters.

Cloud Backup is a data backup solution that stores a copy of your data on remote servers located in the cloud, owned and operated by a third-party service provider. Cloud Backup solutions provide organizations with a highly scalable and cost-effective way to store and protect their data. Cloud Backup typically uses an internet connection to transfer data to the cloud, which can take time, depending on the amount of data being transferred.

Backup Copy, on the other hand, is a data backup solution that creates a second copy of your data on a different storage device. Backup Copy solutions are typically used to create an additional layer of protection for critical data in case the primary backup solution fails. Backup Copy can be set up to store data locally or remotely, depending on the organization's needs.

Both Cloud Backup and Backup Copy have their advantages and disadvantages. Cloud Backup provides a highly scalable and cost-effective way to store and protect data, but it can be slow to transfer large amounts of data to the cloud. Backup Copy, however, provides additional protection for critical data, but it can be expensive to maintain and manage.

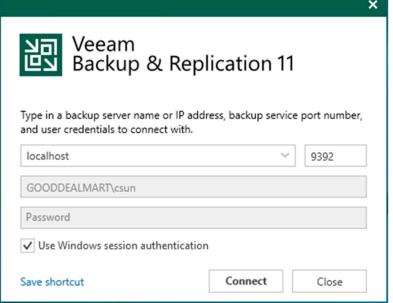
In general, it is recommended that organizations use a combination of both Cloud Backup and Backup Copy to ensure that their data is protected against loss due to unexpected events. This approach provides the benefits of both solutions while mitigating their respective disadvantages.

Creating a Backup job using Azure Blob repositories as Cloud Redundant Data

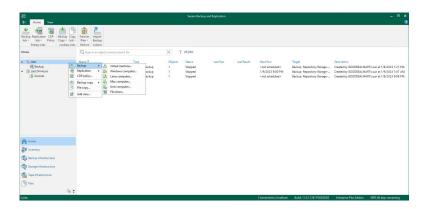
This procedure immediately creates a backup job to sync backup files with Azure cloud and off-loads Azure blob after performing a full backup. It would be best to have a scale-out repository ready before beginning this backup job.



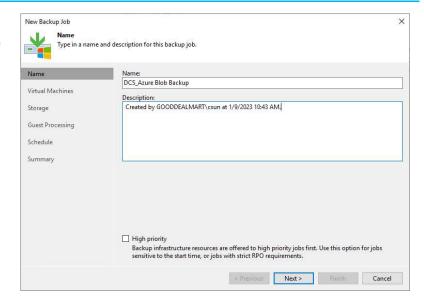
 Open the Veeam Backup & Replication Console, and click Connect.



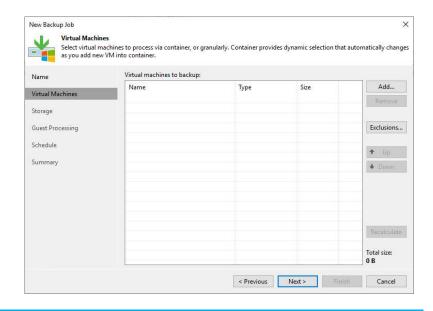
3. On the Home page, select Jobs, right-click Jobs, select Backup and click Virtual machine.



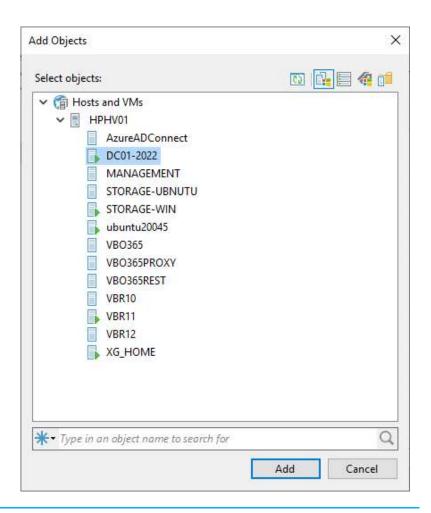
- 4. On the Name page, enter a name for the backup job in the Name field.
- 5. Describe the Description field.
- Select the High priority check box if you want the resource scheduler of Veeam Backup & Replication to prioritize this job higher than similar jobs and allocate resources to it in the first place.
- 7. Click Next.



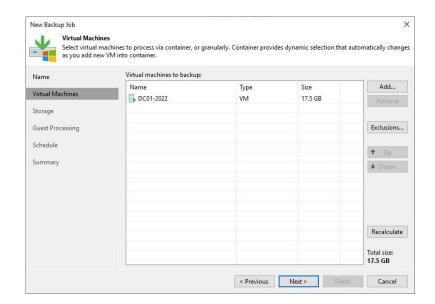
8. On the Virtual Machines page, click Add.



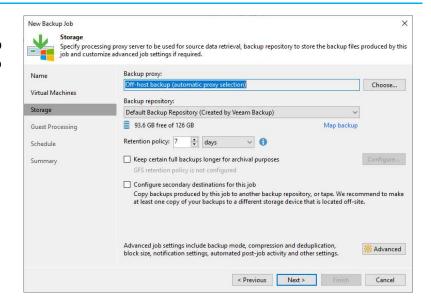
- 9. Select the VM in the list on the Add Objects page and click Add.
- 10. If you have multiple VMS that needs to back up in the same backup job, you can repeat the step to add them.



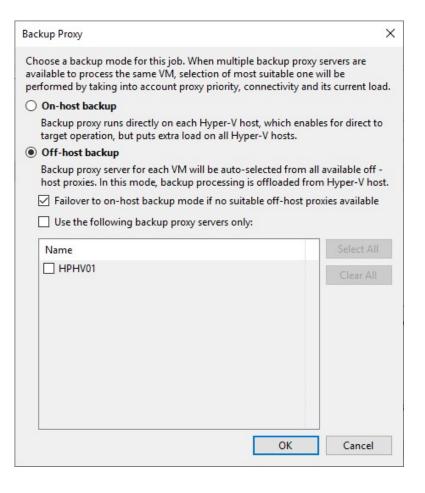
11. On the Virtual Machines page, click Next.



12. On the Storage, click
Choose to select a backup
proxy if you don't want to
use the default Off-host
backup (automatic proxy
selection) setting.

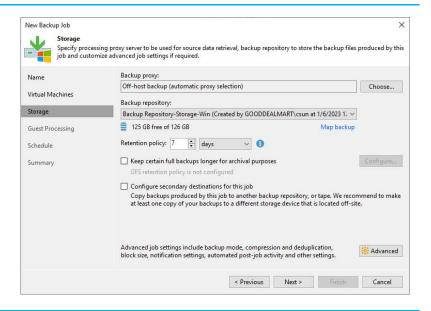


- page, if you select Onhost backup mode, the source Microsoft HyperV host will serve as both the source host and the backup proxy. In this mode, Veeam Data Mover runs directly on the source host, which speeds up data retrieval but places additional strain on the host.
- 14. If you select Off-host backup mode, Veeam Data Mover will run on a dedicated off-host backup proxy. All backup processing operations from the source host are routed to the off-host backup proxy in this mode.
- 15. If the off-host backup mode is selected for the job, but there are no off-host backup proxies available when the job begins, Veeam Backup & Replication will transition to on-host backup mode.
- 16. You unselect the Failover to on-host backup mode if no suitable off-host

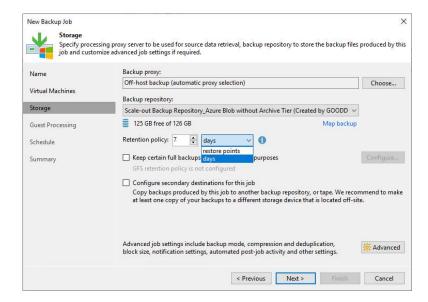


proxies are available checkbox. Still, the job will only start if off-host backup proxies are available or configured properly.

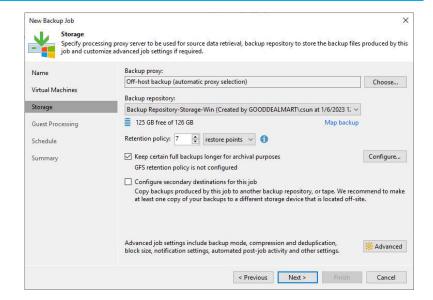
- 17. Click OK.
- 18. Select the Scale-out backup repository from the Backup repository drop-down list.



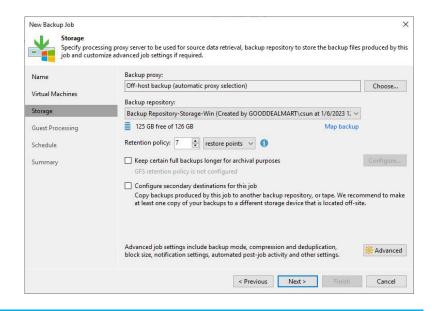
- Set the retention policy settings for restore points in the Retention Policy field.
- 20. Select days or restore points from the drop-down list.



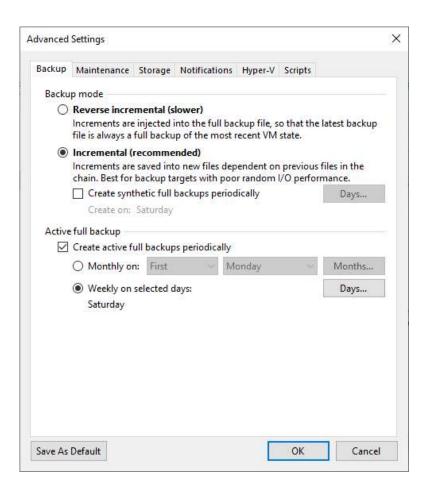
21. You can configure GFS retention policy settings for the backup job to ignore the short-term retention policy for some full backups and store them for long-term archiving.



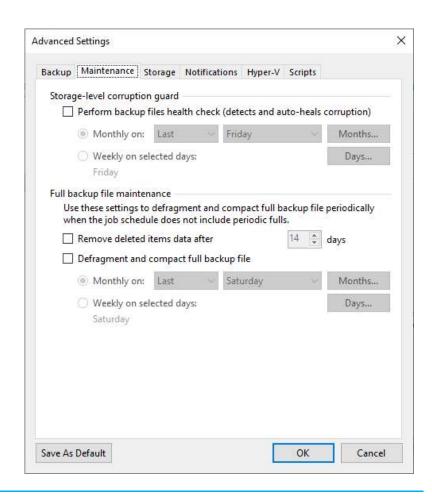
22. On the Storage page, click Advanced.



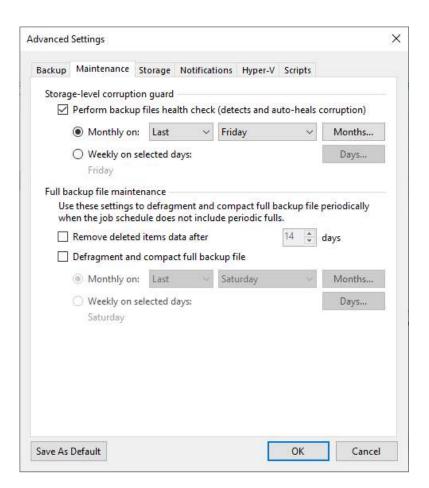
- 23. On the Backup page, Select Incrementally and disable synthetic full.
- 24. In the Active full backup session, select Create active full backups periodically check box.
- 25. Select the Monthly on or Weekly on selected days options to define scheduling settings.
- 26. After creating a full backup file, all backup files start to upload from the local directory to the Azure blob.



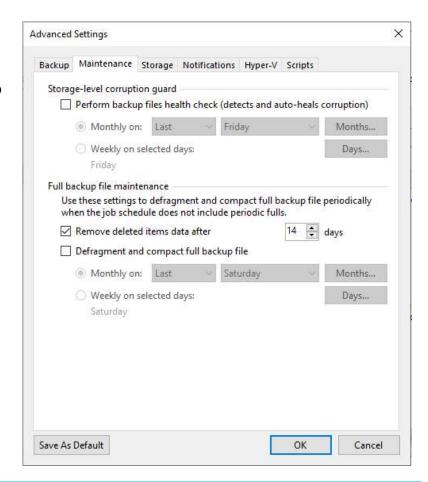
27. On the Advanced Settings, Maintenance.



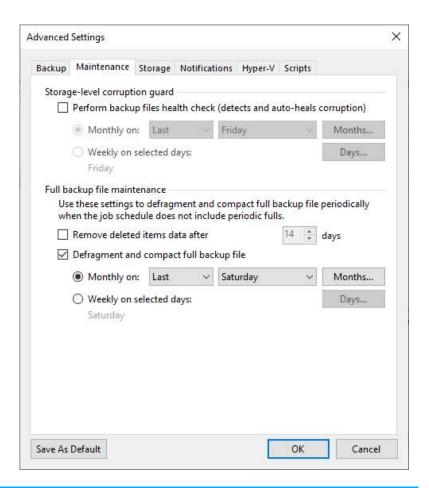
28. To regularly perform a health check on the backup chain's most recent restore point, select the Perform backup files health check (detects and auto-heals corruption) checkbox in the Storage-level corruption guard section and set a timetable for the health check.



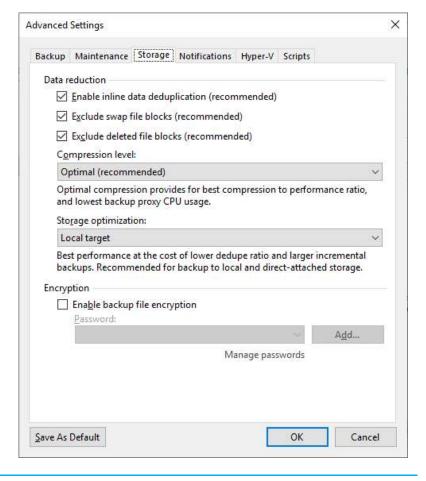
29. Select the Remove deleted items data after the check box and enter the days you want backup data for deleted VMs to be kept.



30. Select the Defragment and compact full backup file check box and specify the schedule for the compact operation to compact a full backup periodically.



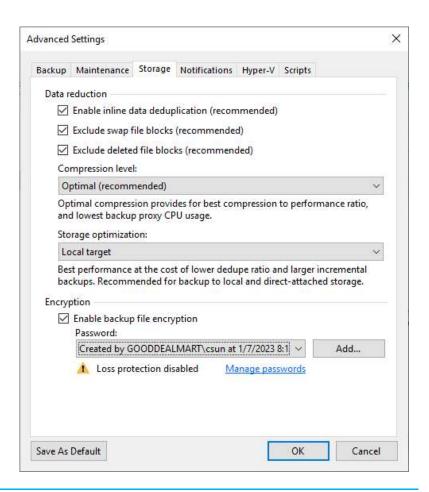
- 31. On Advanced Settings, click Storage.
- 32. Select the Enable inline data deduplication check box.
- 33. Select the Exclude swap file blocks checkbox.
- 34. Select the Exclude deleted file blocks check box.
- 35. Select the compression level for the backup from the drop-down list.



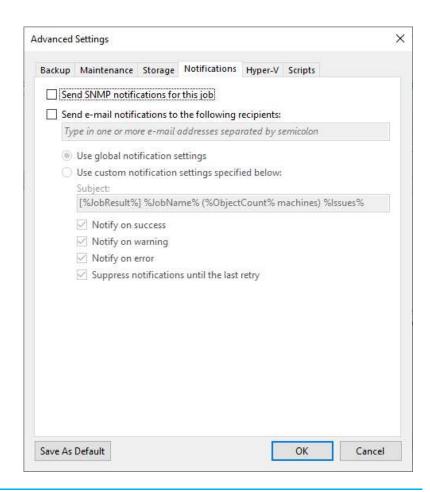
36. Select Storage optimization from the drop-down list.

Storage optimization option	Block size	Description
Local target (large blocks)	4096 KB	Recommended for files that are larger than 16 TB. This option will provide the lowest deduplication ratio and the largest size of incremental files.
Local target	1024 KB	Recommended for backup and replication to SAN, DAS or local storage. This option provides the fastest job performance but reduces the deduplication ratio, because with larger data blocks it is less likely to find identical blocks.
LAN target	512 KB	Recommended for backup and replication to NAS, and on-site backup and replication. This option provides a better deduplication ratio and reduces the size of a file because of reduced data block sizes.
WAN target	256 KB	Recommended if you are planning to use WAN for off-site backup and replication. This option provides the maximum deduplication ratio and the smallest size of files that allows you to reduce the amount of traffic over WAN.

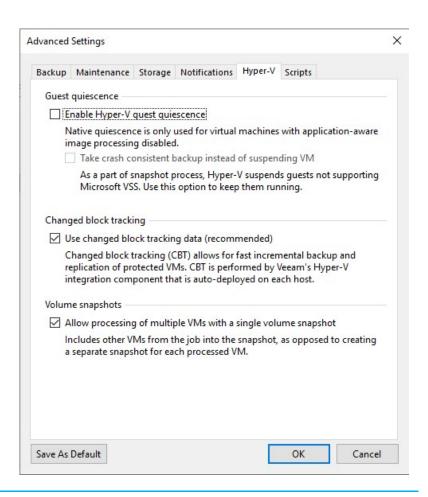
- 37. Select the Enable backup file encryption checkbox to encrypt the content of backup files.
- 38. Select a password from the drop-down list. Then, if you still need to do, click Add or use the Manage passwords link to create a new password.



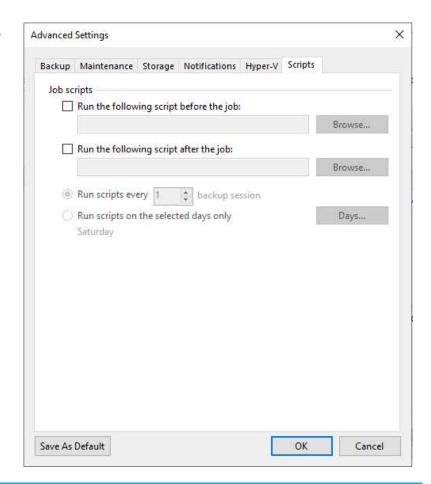
- 39. On the Advanced Settings, select Notifications.
- 40. Keep the default settings.



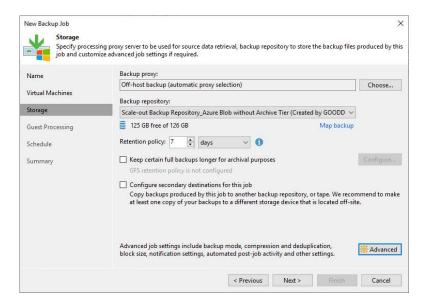
- 41. On the Advanced Settings, select Hyper-V.
- 42. Select the Enable Hyper-V guest quiescence check box if application-aware processing cannot be used for some reason.
- 43. Select the Take crash consistent backup instead of suspending VM check box if you do not want to suspend VMs in the job.
- 44. Select the Use changed block tracking data (recommended) check box.
- 45. Select the Allow processing of multiple VMs with a single volume snapshot check box.



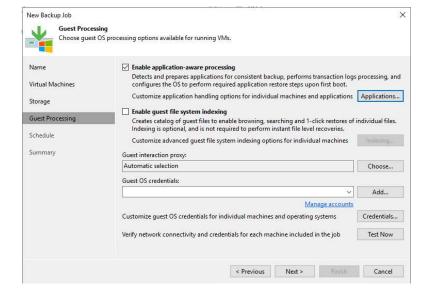
- 46. On the Advanced Settings page, click Scripts.
- 47. Keep the default settings.
- 48. Click OK.



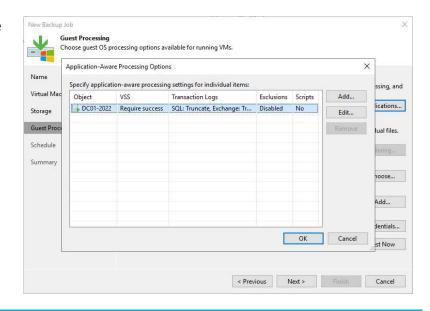
49. On the Storage page, click Next.



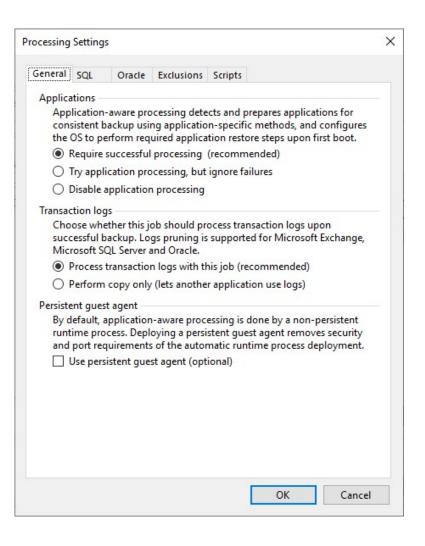
- 50. When you add VMs running VSS-aware applications to the backup job, you can enable application-aware processing to create a transactionally consistent backup. The transactionally consistent backup ensures that applications on VMs can be recovered without data loss.
- 51. Select the Enable application-aware processing check box on the Guest Processing page and click Applications.



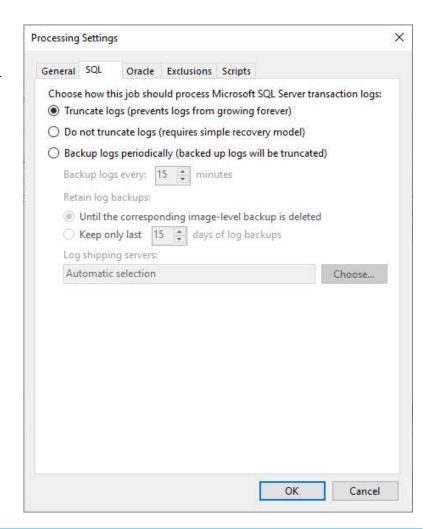
52. On the Application-Aware Processing Options page, select the VM and click Edit.



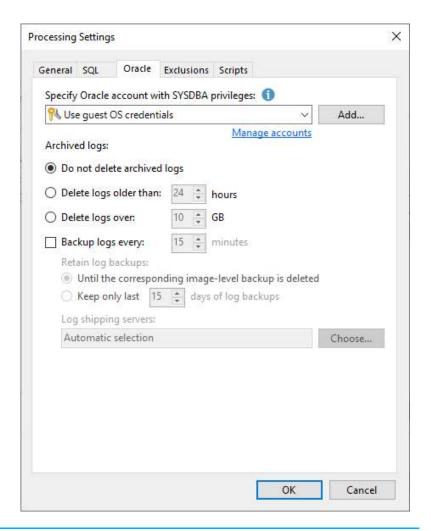
- 53. On the Processing Settings, click General.
- 54. Suppose you need Veeam Backup & Replication to stop the backup process if any error occurs during application-aware processing. Select Require successful processing (recommend).
- 55. Suppose you must continue the backup process even if there is an error during applicationaware processing. Select Try application processing but ignore failures.
- 56. Select Disable application processing to disable application-aware processing for the VM.
- 57. Select Process transaction logs with this job (recommend).
- 58. Select Perform copy only to let another application use
- 59. Select the Use persistent guest agent (optional) checkbox to enable persistent agent.



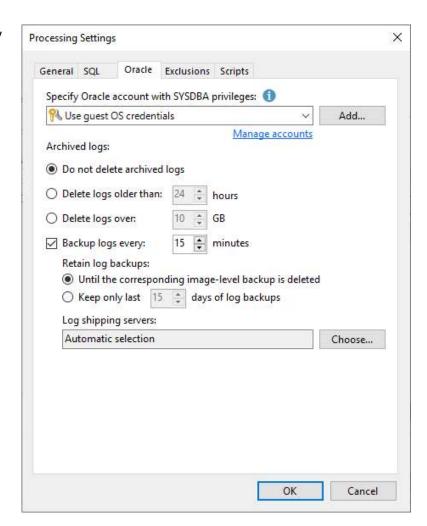
- 60. On the Processing
 Settings page, click SQL if
 the VM is a Microsoft SQL
 Server VM.
- 61. Select Truncate logs
 (Prevents logs from
 growing forever) to
 truncate transaction logs
 after a successful backup.



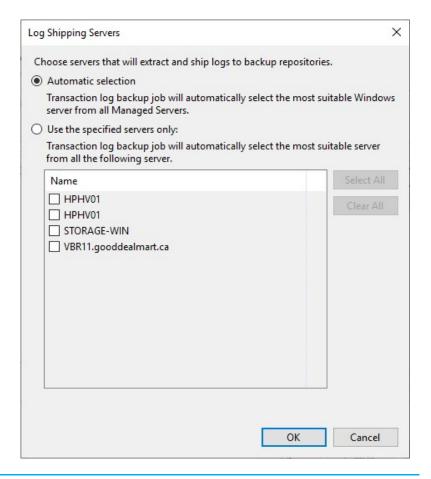
- 62. On the Processing
 Settings page, click Oracle
 if the VM is an Oracle
 Server.
- 63. Select a user account from the drop-down list.
- 64. Select Do not delete archived logs.



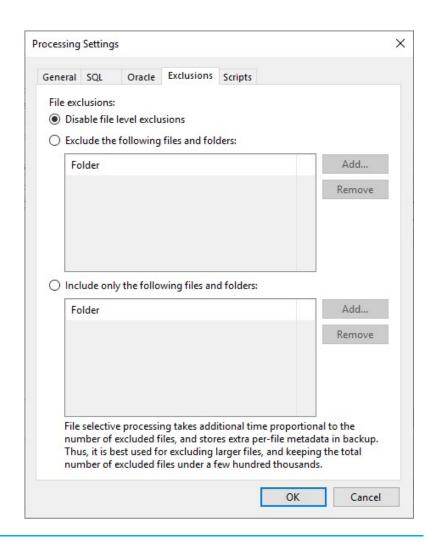
- 65. Select the retention policy settings for archived logs in the Retain log backups section.
- 66. Click Choose In the Log shipping servers.



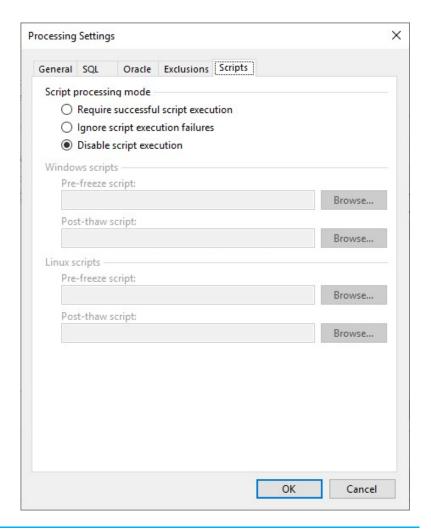
- 67. On the Log Shipping
 Servers page, Select
 Automatic selection if you
 need Veeam Backup &
 Replication to choose an
 optimal log shipping
 server automatically.
- 68. Select Use the specified servers only and then select check boxes next to those you want to use as log shipping servers.
- 69. Click OK.



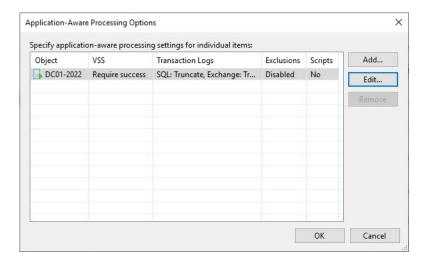
70. On the Processing
Settings page, click
Exclusions and keep the
default settings.



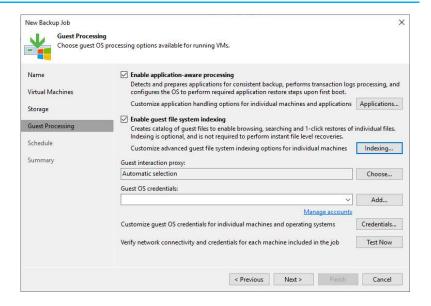
- 71. On the Processing
 Settings page, click Scripts
 and keep the default
 settings.
- 72. Click OK.



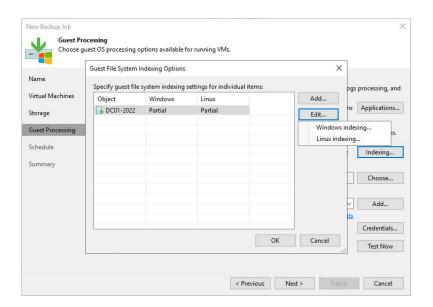
73. On the Application-Aware Processing Options page, click OK.



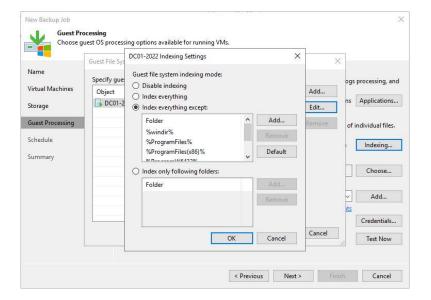
74. Select the Enable guest file system indexing checkbox and click Indexing.



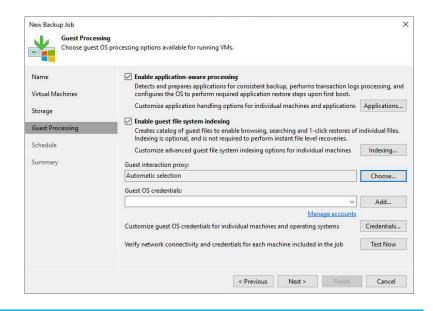
75. On the Guest File System Indexing Options page, select the VM, click Edit and select Windows indexing.



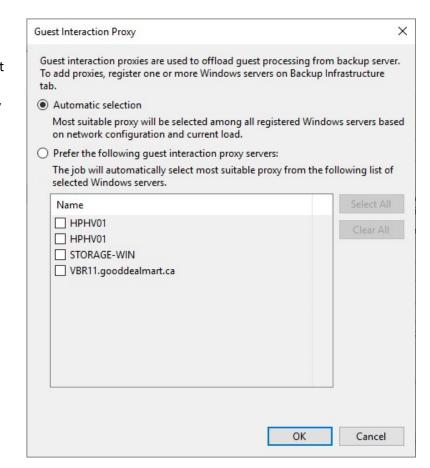
- 76. On the Guest file system indexing mode page, keep the default settings.
- 77. Click OK.



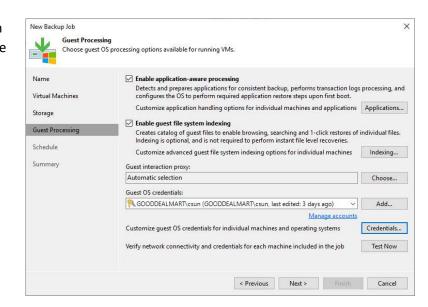
78. Click Choose in the Guest interaction proxy field on the Guest Processing page.



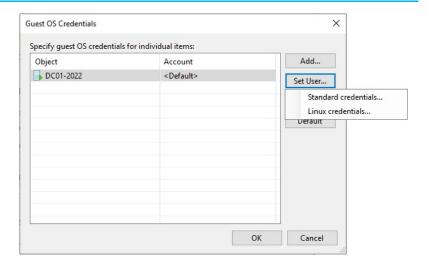
- 79. On the Guest Interaction Proxy page, select Automatic selection to let Veeam Backup & Replication automatically select the guest interaction proxy.
- 80. Select Prefer the following guest interaction proxy servers to explicitly define which servers will perform the guest interaction proxy role.
- 81. Click OK.



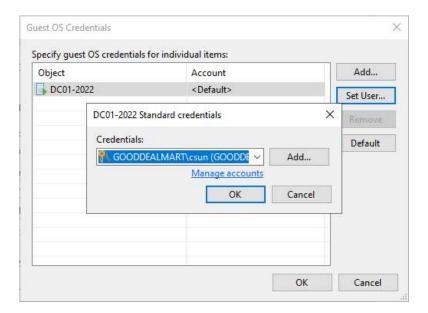
- 82. Choose a user account on the Guest Processing page with sufficient permissions from the Guest OS credentials drop-down list.
- 83. Click Credentials to Customize guest OS credentials for individual machines and operating systems.



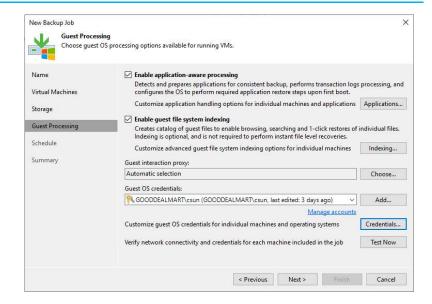
- 84. On the Guest OS
 Credentials page, select
 the VM and click Set User.
- 85. Select Standard credentials.



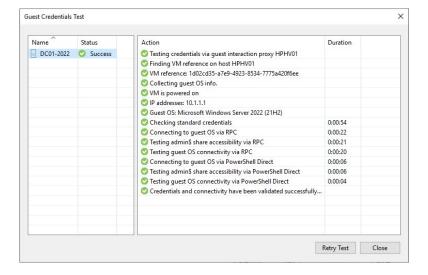
- 86. Choose a user from the Credentials drop-down list, and click OK.
- 87. Repeat the steps for each VM.



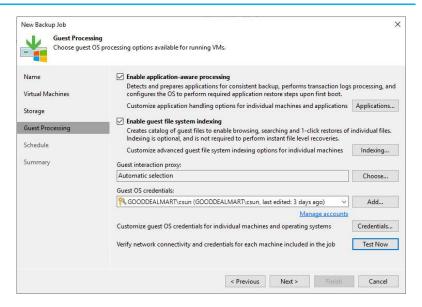
88. On the Guest Processing page, click Test Now to verify network connectivity and credentials for each machine included in the job.



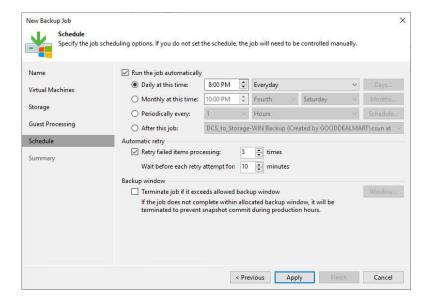
- 89. On the Guest Credentials Test page, verify each machine's success.
- 90. Click Close.



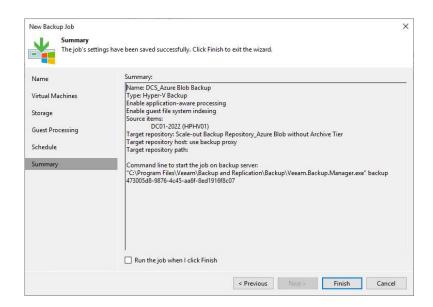
91. On the Guest Processing page, click Next.



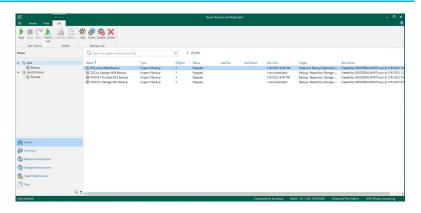
- 92. Select Run the job automatically on the Schedule page and select your specified schedule.
- 93. Define whether Veeam Backup & Replication should retry the backup job if it fails in the Automatic retry section.
- 94. Define the time interval the backup job must complete in the Backup window section. The backup window ensures that the job does not overlap with production hours and that there is no unnecessary overhead on the production environment.
- 95. Click Apply.



96. On the Summary page, click Finish.

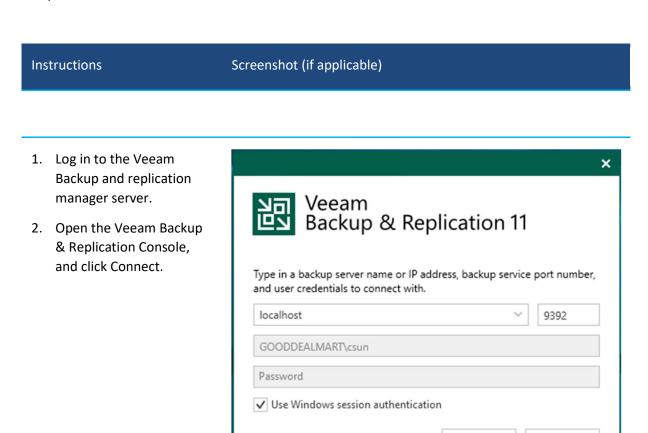


97. Verify that the backup job has been added



Creating a Backup Copy Jobs offload backups to Azure Blob without using the Archive Tier

This procedure creates a backup copy job that off-loads to the Azure blob after performing a full backup. Therefore, it will not use the Azure Blob archive tier.

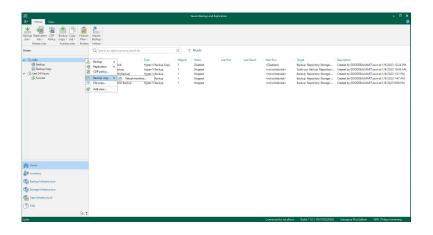


Save shortcut

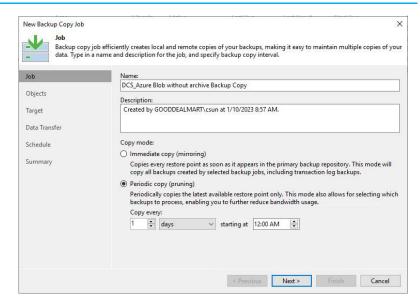
Connect

Close

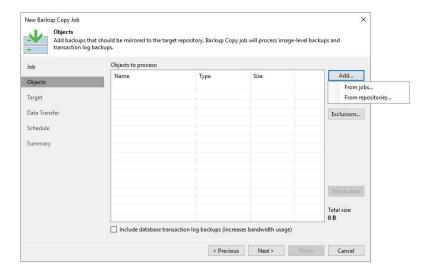
3. On the Home page, select Jobs, right-click Jobs, select Backup copy and click Virtual machine.



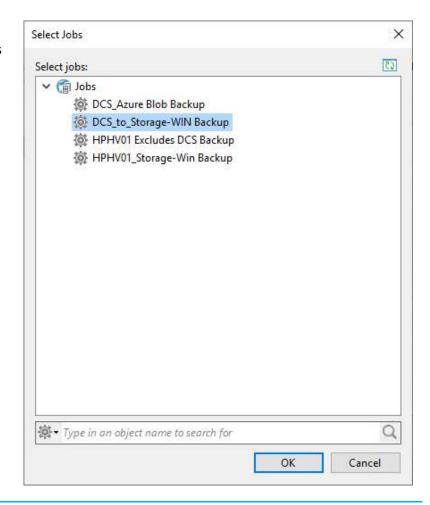
- 4. On the Name page, enter a name for the backup job in the Name field.
- 5. Describe the Description field.
- In the copy mode session, select a backup copy mode. You cannot change the set mode after configuring the backup copy job.
- 7. Click Next.



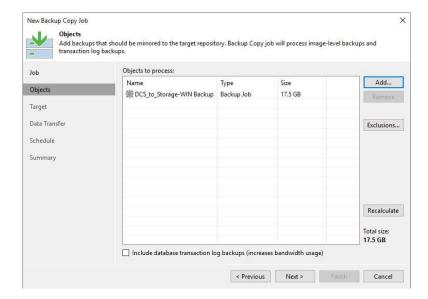
8. On the Objects page, click Add and select From jobs.



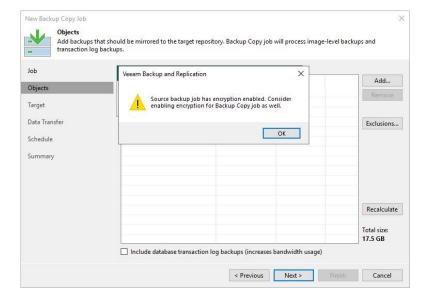
9. Select the job from the jobs list on the Select jobs page and click OK.



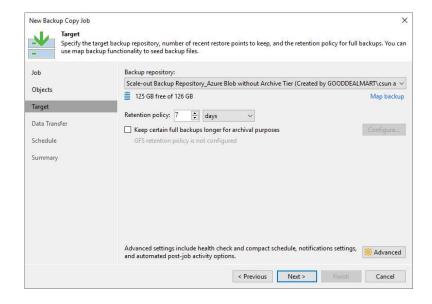
- 10. On the Objects page, click Next.
- 11. Select Include database transacting log backups (increases bandwidth usage) If required.



12. Click OK in the encryption-enabled warning message if the source backup job has encryption enabled.

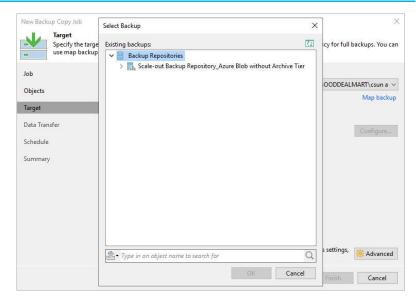


13. Select the Scale-out backup repository_Azure Blob on the Target page without Archive Tier (We created it at Configuring Backup Infrastructure session) from the dropdown list.

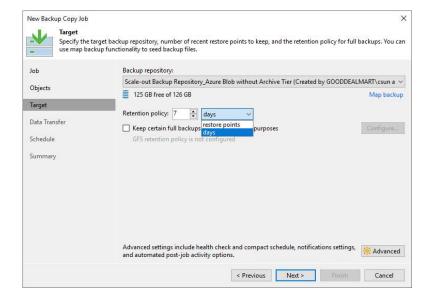


14. Click Map backup if required. It is helpful if you have relocated backup copy files to a new backup repository and want to point the job to existing backups in this new backup repository.

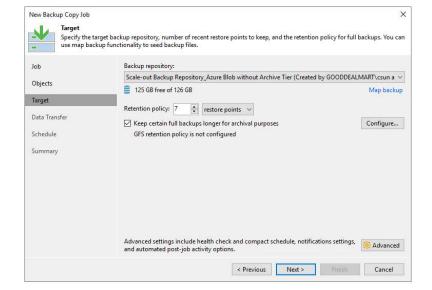
Backup copy job mapping can also be used if the configuration database becomes corrupt and you need to reconfigure backup jobs.



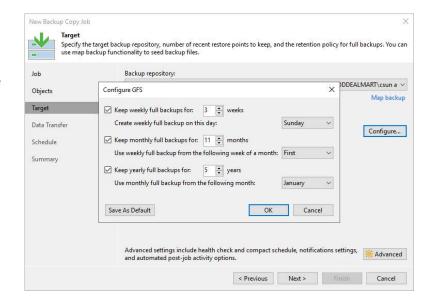
- Set the retention policy settings for restore points in the Retention Policy field.
- 16. Select days or restore points from the drop-down list.



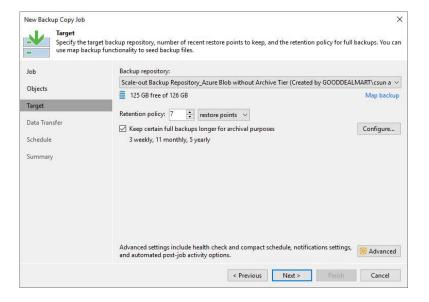
- 17. You can configure GFS retention policy settings for the backup copy job for long-term archiving.
- 18. Select Keep specific full backups for longer for archival purposes, and click Configure.



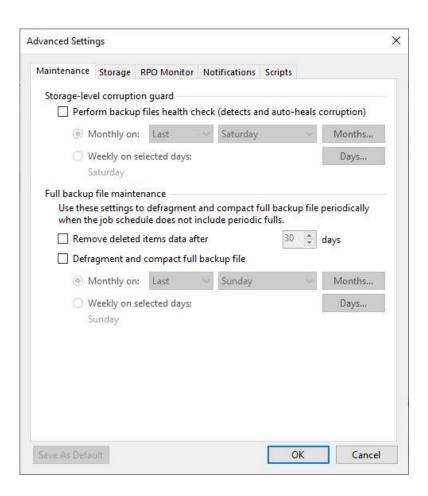
- 19. On the Configure GFS page, select the Keep weekly full backups for check box and specify the number of weeks you want to prevent restore points from being modified and deleted.
- 20. Select the Keep monthly full backups for check box, and specify the months you want to prevent restore points from being modified and deleted.
- 21. Select the Keep yearly full backups for check box, and specify the years you want to prevent restore points from being modified and deleted.
- 22. Click OK.



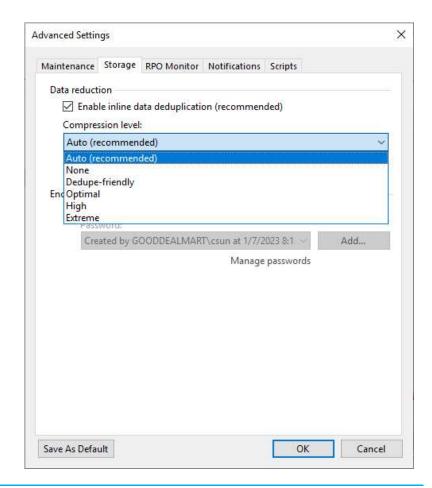
23. On the Target page, click Advanced.



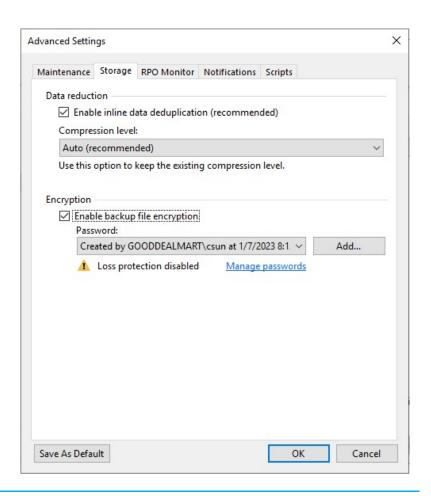
- 24. On the Advanced Settings, click Maintenance.
- 25. Select the Perform backup files health check (detects and auto-heals corruption) checkbox and specify the schedule for the health check if required.
- 26. Select the Remove deleted items data after checkbox and specify the retention days settings for deleted workloads if required.
- 27. Select the Defragment and compact full backup file checkbox and specify the schedule for the compacting operation if required.



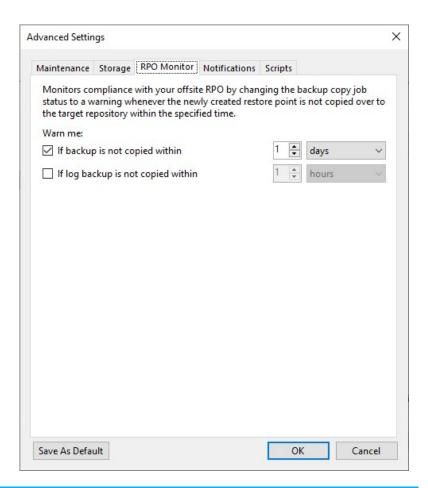
- 28. On Advanced Settings, click Storage.
- 29. Select the Enable inline data deduplication check box.
- 30. Select the compression level for the backup copy from the drop-down list.



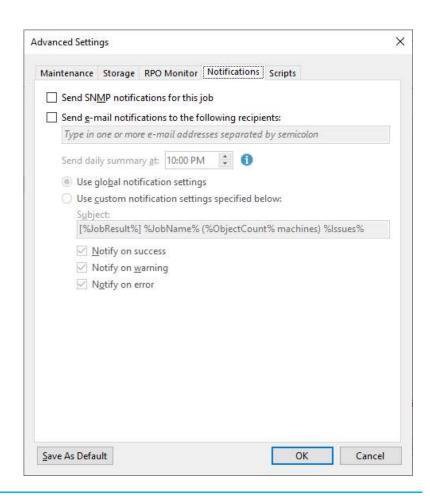
- 31. Select the Enable backup file encryption checkbox to encrypt the content of backup files.
- 32. Select a password from the drop-down list. If you haven't done so, click Add or use the Manage passwords link to create a new password.



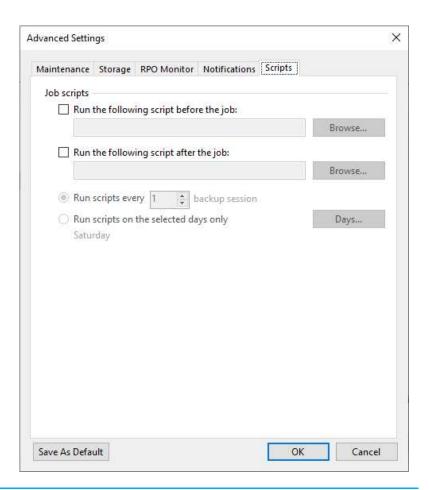
- 33. On the Advanced Settings page, select RPO Monitor.
- 34. Select Warn me if the backup is not copied within the checkbox, and specify the desired RPO in minutes, hours or days.
- 35. Select Warn me if log backup is not copied within the checkbox. If you have enabled copying of log backups, specify the desired RPO in minutes, hours or days.



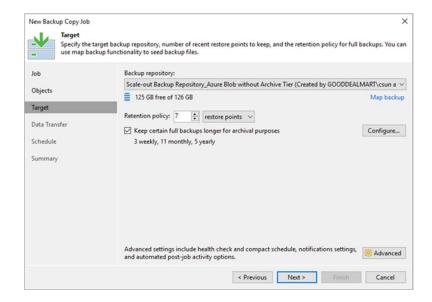
- 36. On the Advanced Settings, select Notifications.
- 37. Keep the default settings.



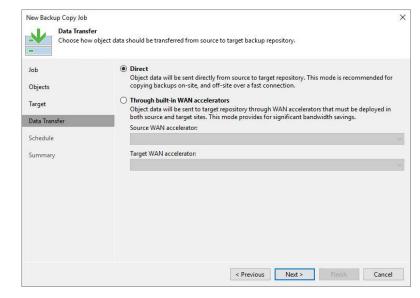
- 38. On the Advanced Settings page, click Scripts.
- 39. Keep the default settings and click OK.



40. On the Target page, click Next.

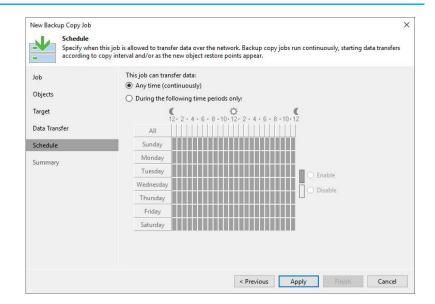


- 41. On the Data Transfer page, Select Direct if you plan to copy backup files over high-speed connections.
- 42. Select the Through builtin WAN accelerators if you copy backup files over WAN or slow connections.
- 43. Select a WAN accelerator configured in the source site from the Source WAN accelerator drop-down list.
- 44. Select a WAN accelerator configured in the target site from the Target WAN

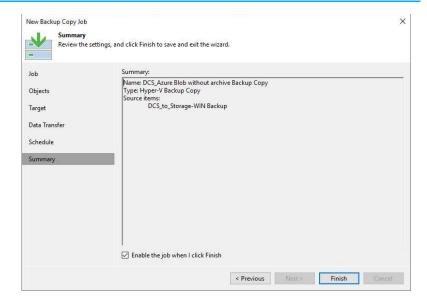


accelerator drop-down list.

- 45. Click Next.
- 46. On the Schedule page, select Any time (continuously) if this job can transfer data at any time.
- 47. Select During the following periods only if required.
- 48. Click Apply.

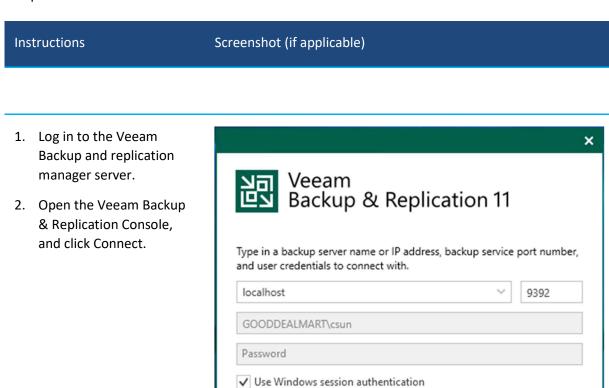


49. Select Enable the job on the Summary page when I click the Finish check box. If you want to start the job after creating it, click Finish.



Creating a Backup Copy Jobs offload backups to Azure Blob using the Archive Tier

This procedure creates a backup copy job that off-loads to the Azure blob after performing a full backup. It will use the Azure Blob archive tier.

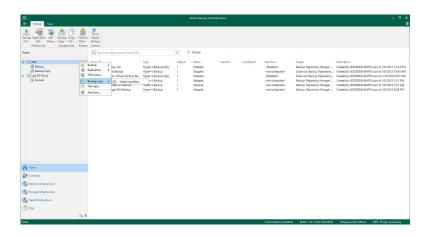


Save shortcut

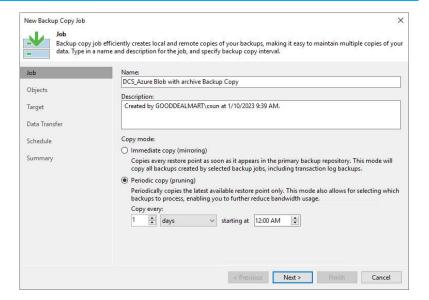
Close

Connect

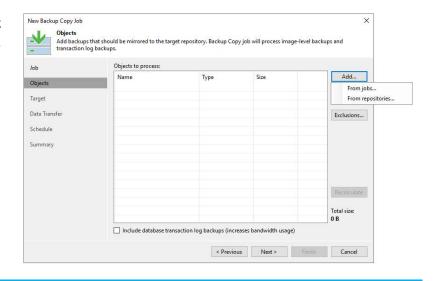
3. On the Home page, select Jobs, right-click Jobs, select Backup copy and click Virtual machine.



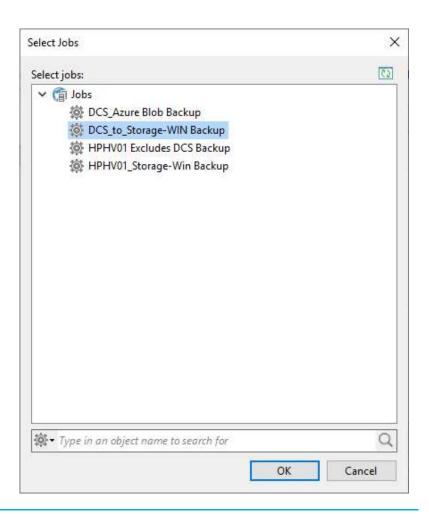
- 4. On the Name page, enter a name in the Name field.
- 5. Describe the Description field.
- In the copy mode session, select a backup copy mode. You cannot change the set mode after configuring the backup copy job.
- 7. Click Next.



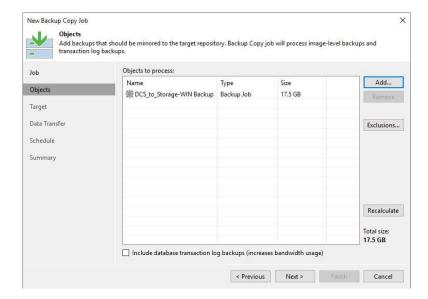
8. On the Objects page, click Add and select From jobs.



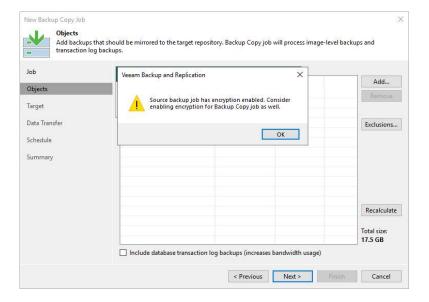
9. Select the job from the jobs list on the Select jobs page and click OK.



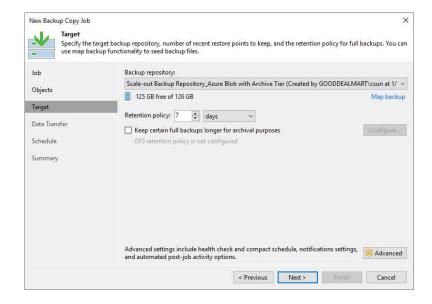
- 10. On the Objects page, click Next.
- 11. Select Include database transacting log backups (increases bandwidth usage) If required.



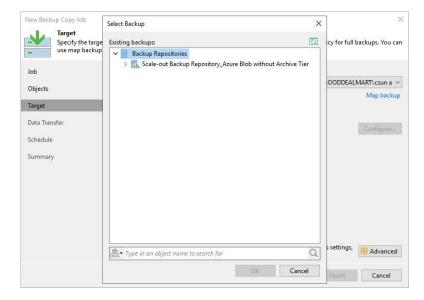
12. Click OK in the encryption-enabled warning message if the source backup job has encryption enabled.



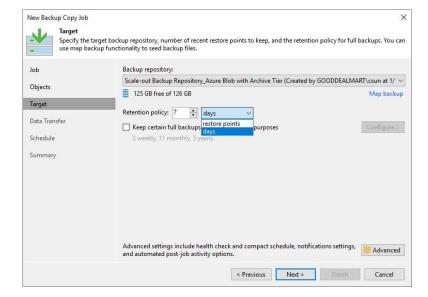
13. Select the Scale-out backup repository_Azure Blob on the Target page with Archive Tier (We created it at Configuring Backup Infrastructure session) from the dropdown list.



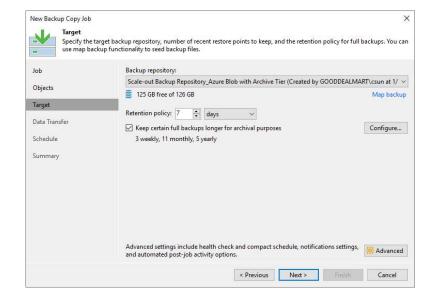
14. Click Map backup if required. It is helpful if you have relocated backup copy files to a new backup repository and want to point the job to existing backups in this new backup repository. Backup copy job mapping can also be used if the configuration database becomes corrupt and you need to reconfigure backup jobs.



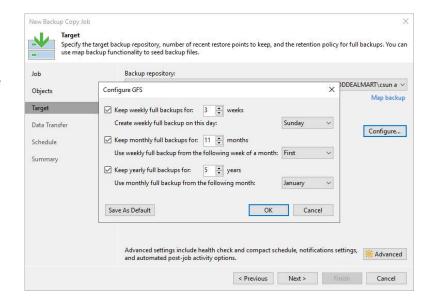
- 15. Set the retention policy settings for restore points in the Retention Policy field.
- 16. Select days or restore points from the drop-down list.



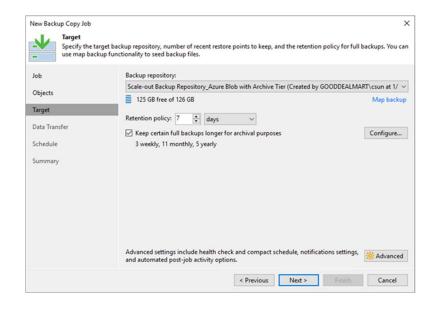
- 17. You can configure GFS retention policy settings for the backup copy job for long-term archiving.
- 18. Select Keep specific full backups for longer for archival purposes, and click Configure.



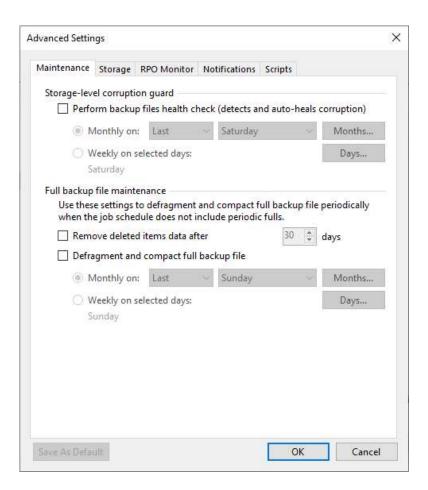
- 19. On the Configure GFS page, select the Keep weekly full backups for check box and specify the number of weeks you want to prevent restore points from being modified and deleted.
- 20. Select the Keep monthly full backups for check box, and specify the months you want to prevent restore points from being modified and deleted.
- 21. Select the Keep yearly full backups for check box, and specify the years you want to prevent restore points from being modified and deleted.
- 22. Click OK.



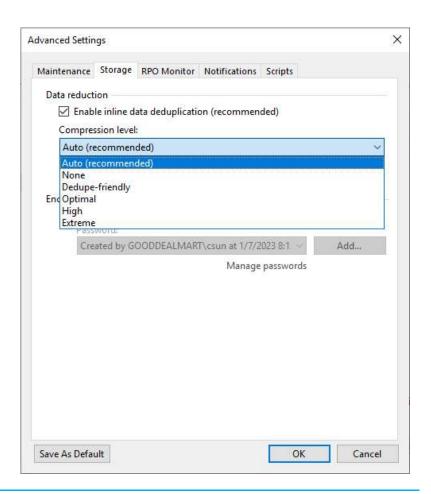
23. On the Target page, click Advanced.



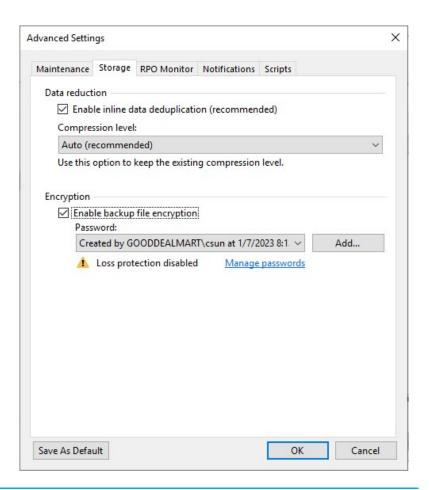
- 24. On the Advanced Settings, click Maintenance.
- 25. Select the Perform backup files health check (detects and auto-heals corruption) checkbox and specify the schedule for the health check if required.
- 26. Select the Remove deleted items data after checkbox and specify the retention days settings for deleted workloads if required.
- 27. Select the Defragment and compact full backup file checkbox and specify the schedule for the compacting operation if required.



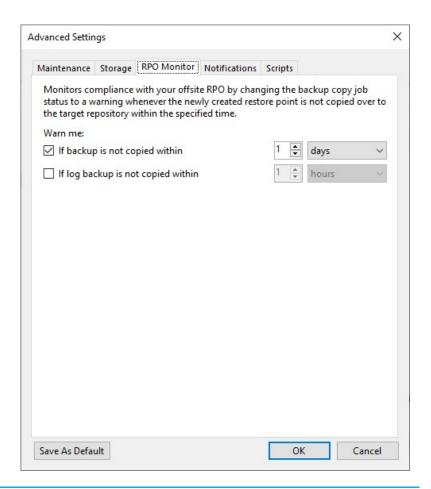
- 28. On Advanced Settings, click Storage.
- 29. Select the Enable inline data deduplication check box.
- 30. Select the compression level for the backup copy from the drop-down list.



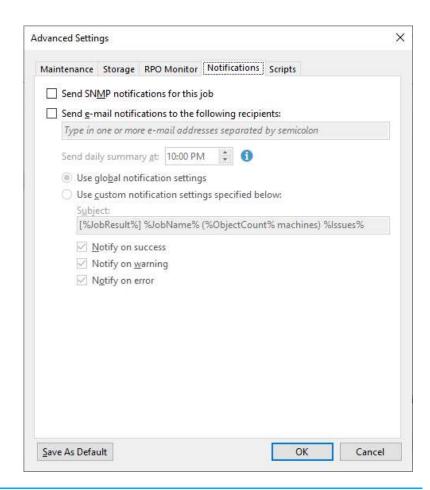
- 31. Select the Enable backup file encryption checkbox to encrypt the content of backup files.
- 32. Select a password from the drop-down list. If you haven't done so, click Add or use the Manage passwords link to create a new password.



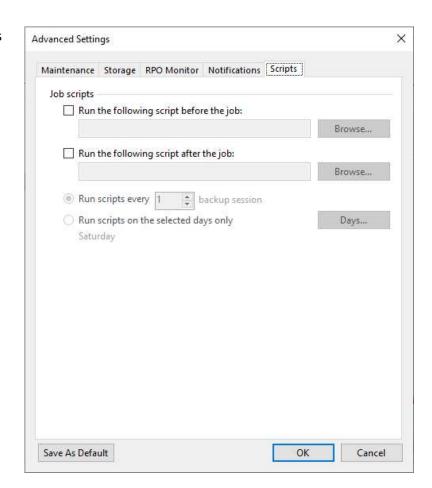
- 33. On the Advanced Settings page, select RPO Monitor.
- 34. Select Warn me if backup is not copied within the checkbox, and specify the desired RPO in minutes, hours or days.
- 35. Select Warn me if log backup is not copied within the checkbox. If you have enabled copying of log backups, specify the desired RPO in minutes, hours or days.



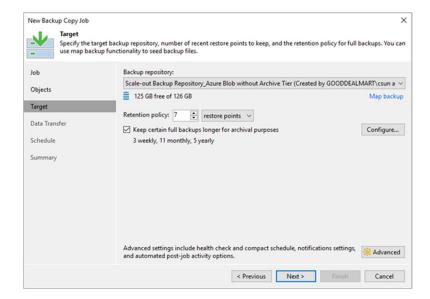
- 36. On the Advanced Settings, select Notifications.
- 37. Keep the default settings.



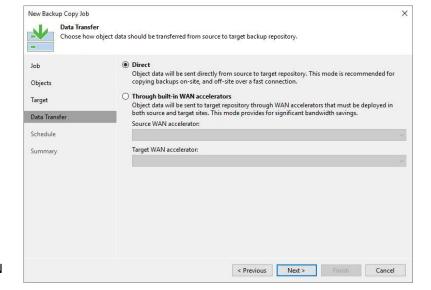
- 38. On the Advanced Settings page, click Scripts.
- 39. Keep the default settings.
- 40. Click OK.



41. On the Target page, click Next.

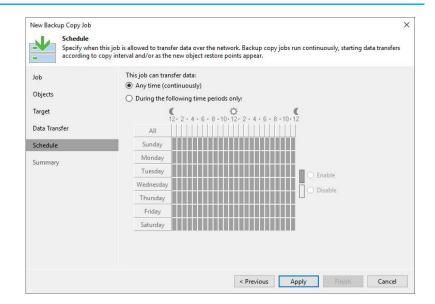


- 42. On the Data Transfer page, Select Direct if you plan to copy backup files over high-speed connections.
- 43. Select the Through builtin WAN accelerators if you copy backup files over WAN or slow connections.
- 44. Select a WAN accelerator configured in the source site from the Source WAN accelerator drop-down list.
- 45. Select a WAN accelerator configured in the target site from the Target WAN

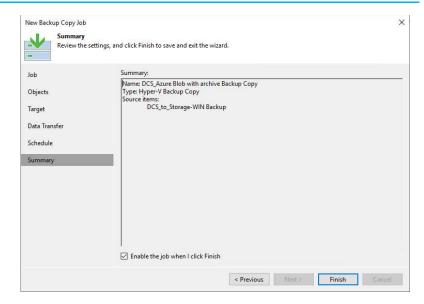


accelerator drop-down list.

- 46. Click Next.
- 47. On the Schedule page, select Any time (continuously) if this job can transfer data at any time.
- 48. Select During the following periods only if required.
- 49. Click Apply.



50. Select Enable the job on the Summary page when I click the Finish check box. If you want to start the job after creating it, click Finish.



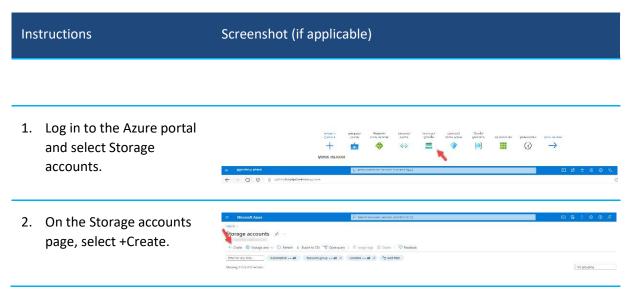
Chapter 3

Restore to Microsoft Azure

Veeam Backup & Replication allows you to restore various workloads (VMs, Google VM instances, physical servers, and so on) to Microsoft Azure from backups.

Restore On-premises VM to Microsoft Azure

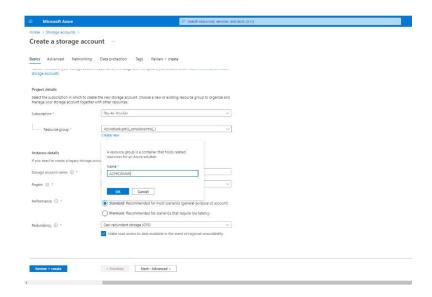
This procedure recovers (or moves) a backup VM to Microsoft Azure.



 On the Basics tab, configure Project details and Instance details settings for the new Storage account and click Next: Advanced.

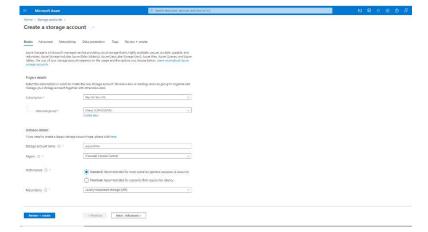
Project details

- Subscription: Select your azure subscription, in my case, is Pay-As-You-Go.
- Resource group: Select Create new, type
 AZPRODVMS in the name field, and click OK.

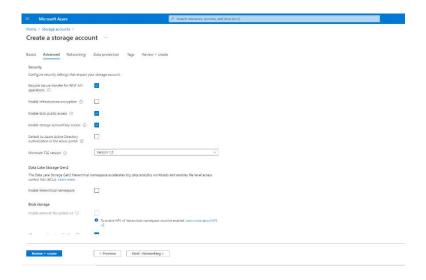


Instance details

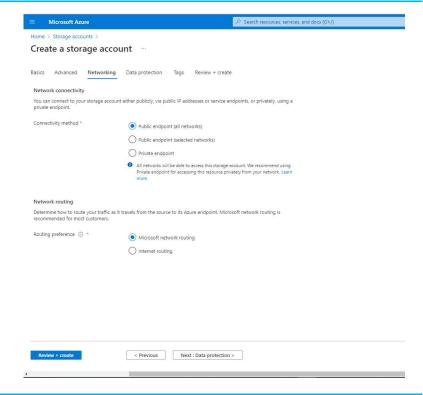
- Storage account name: Type approves.
- Region: Select (Canada)
 Canada Central.
- Performance: Select Standard: Recommend for most scenarios (generalpurpose v2 account).
- Redundancy: Select Locallyredundant storage (LRS).



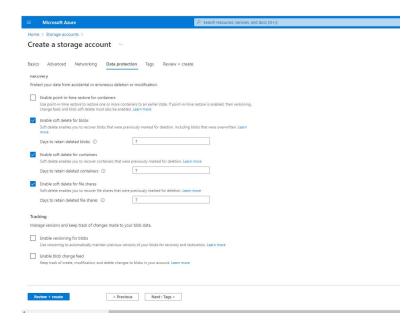
 On the Advanced tab, keep the default settings and click Next: Networking. You can change settings based on your requirements.



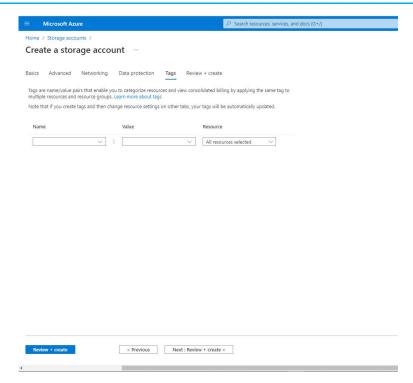
5. On the Networking tab, keep the default settings and click Next: Data protection. You can change settings based on your requirements.



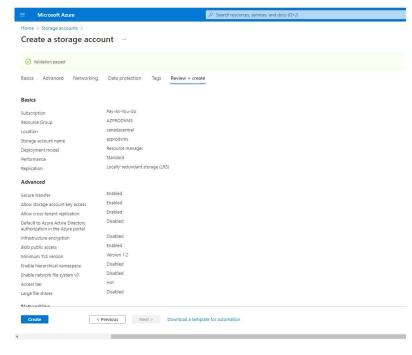
 Keep the default settings on the Data Protection tab, and click Next: Tags. You can change settings based on your requirements.



 On the Tags tab, keep the default settings and click Next: Review + create.
 You can change settings based on your requirements.



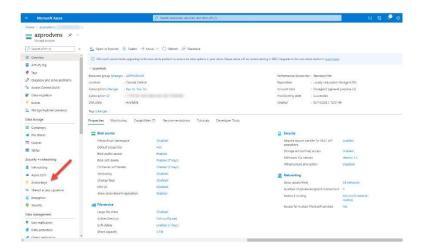
8. On the Review + create tab, review settings, and click Create.



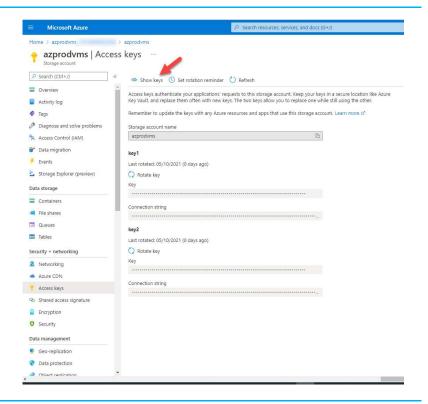
9. Click Go to resource after creating a new storage account completed.



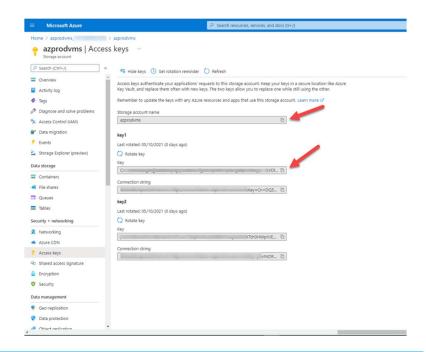
10. On the new storage account page, select Access keys.



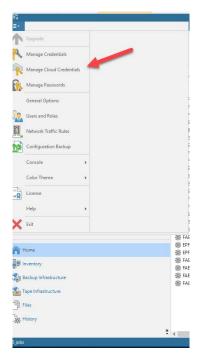
11. On the Access keys page, select Show keys.



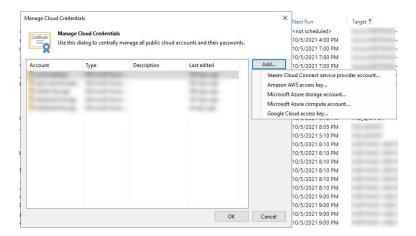
12. Copy the storage account name and Key to a safe place. We will need them at laser settings.



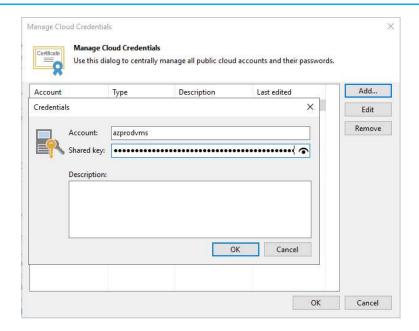
13. On the Veeam management console, from the main menu, select Manage Cloud Credentials.



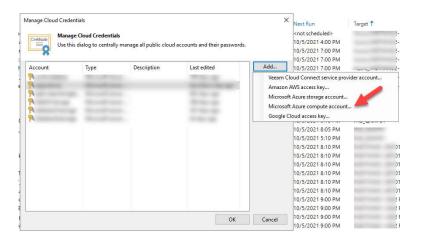
14. Click Add on the Manage Cloud Credentials page and select Microsoft Azure storage account.



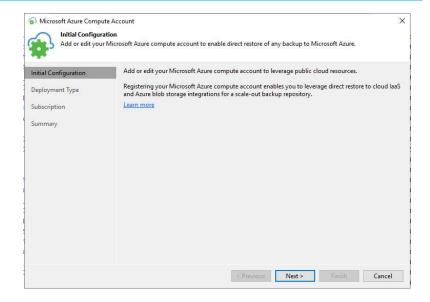
15. Paste the storage account name to Account, paste the key to Shared key, and click OK.



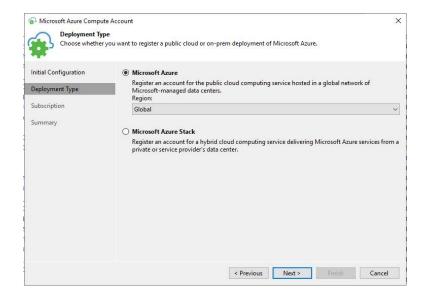
16. Click Add on the Manage Cloud Credentials page and select Microsoft Azure compute account.



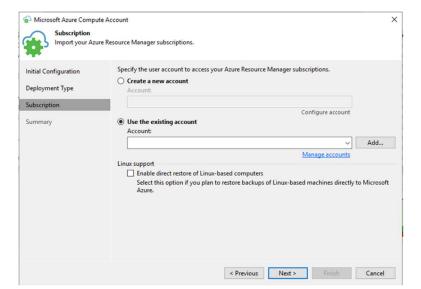
17. On the Initial Configuration page, click Next.



- 18. On the Deployment Type page, select Microsoft Azure.
- 19. Select Global from the region list and click Next.

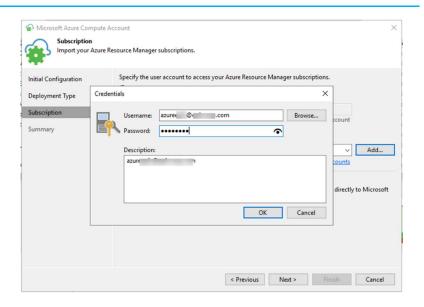


- 20. Select the method of importing your Azure Resource Manager subscription on the Subscription page. You have two options: Use the existing account and Create a new account.
- The Azure account must have the Contributor role privileges for the required subscription. However, you can create a custom role with minimal permissions if the Contributor role cannot be used.
- Only subscriptions that belong to the selected

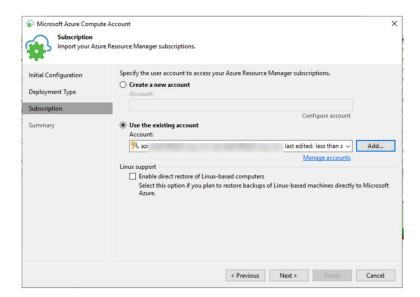


account's directory will be added.

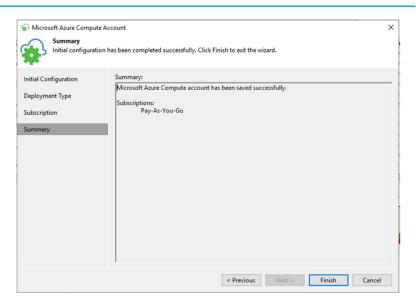
- In this scenario, you cannot add an account with enabled MFA. You must disable MFA for the required account. App passwords are not supported.
- If you have more than one Azure Active Directory tenant associated with your account, specify which tenant to use.
- 21. On the Subscription page, select Use the existing account and click Add.
- 22. Enter a username in the Username field.
- 23. Enter a password in the Password field and click OK.



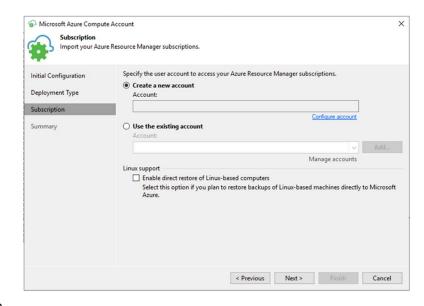
24. On the Subscription page, click Next.



25. On the Summary page, click Finish.

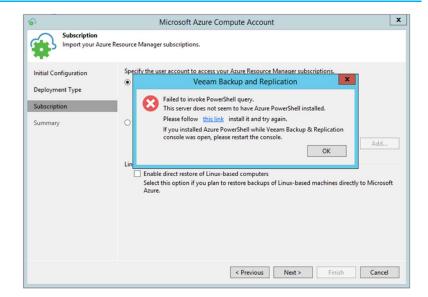


- 26. If you Select Use the
 Create a new account,
 Veeam Backup &
 Replication will register a
 special application on
 Azure. Veeam Backup &
 Replication will use this
 application to
 communicate with Azure.
 Mind the following
 prerequisites.
 - A Microsoft Azure account that you plan to add to Veeam Backup & Replication must have the Owner role privileges for the subscription that will be used for restoring to Microsoft Azure. Owner role privileges are required to provide access to a subscription for the created application.
 - The user must have privileges to register applications: Global Administrator privileges or the enabled Users can register applications option in the Azure portal.
- 27. Select Create a new account on the

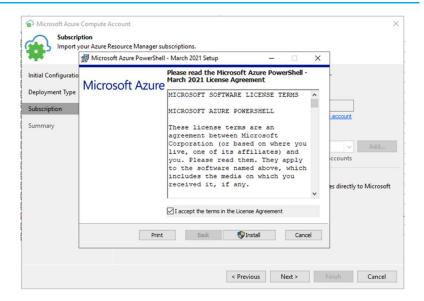


Subscription page and click Configure account.

28. On the Failed to invoke PowerShell query error page, click this link to install Azure PowerShell.



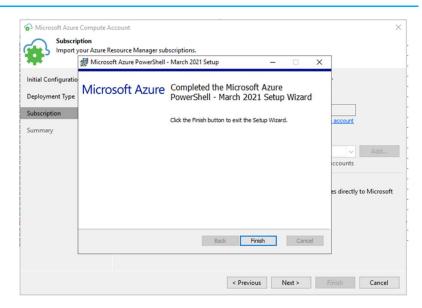
29. On the Microsoft Azure PowerShell Setup page, select I accept the terms in the License Agreement checkbox and click Install.



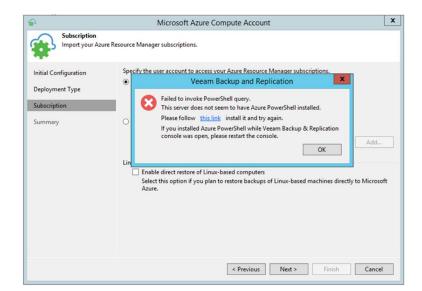
30. On the User Account Control page, click Yes.



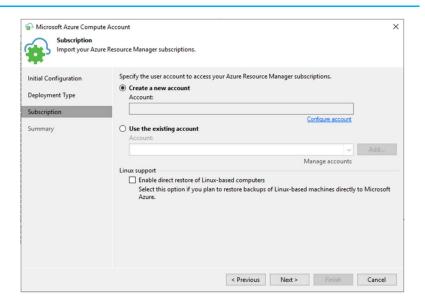
31. On the Microsoft Azure PowerShell Setup page, click Finish.



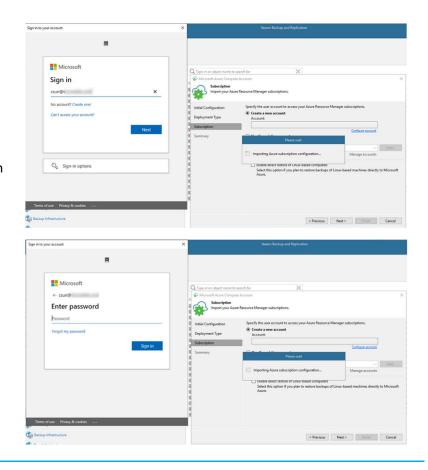
32. Click OK on the Failed to invoke PowerShell query error page.



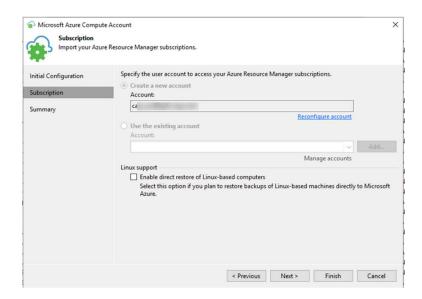
33. Re-click the Configure account again after installing Azure PowerShell.



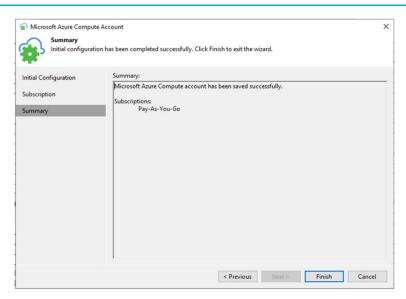
34. Enter the credentials of an existing Microsoft Azure account in the browser window. Veeam Backup & Replication will retrieve information about subscriptions and resources associated with this account.



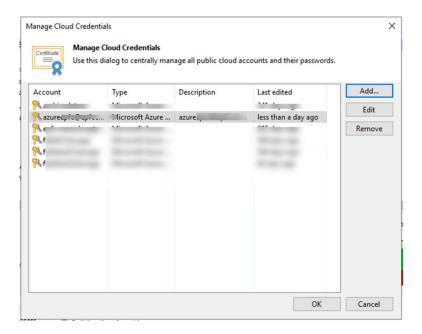
35. On the Subscription page, click Next.



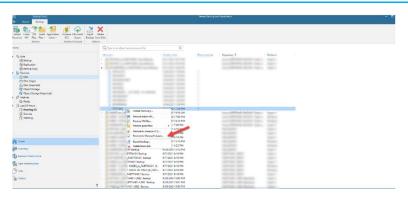
36. On the Summary page, click Finish.



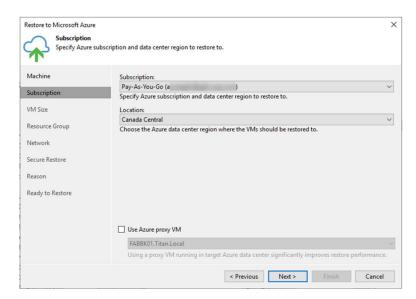
37. On the Manage Cloud Credentials page, click OK.



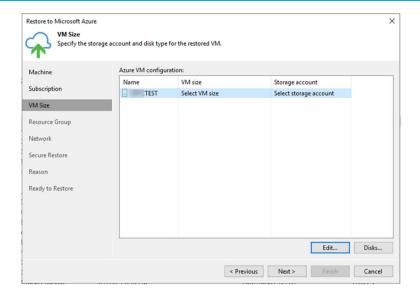
38. Right-click the backup VM you want to restore to Azure and select Restore to Microsoft Azure.



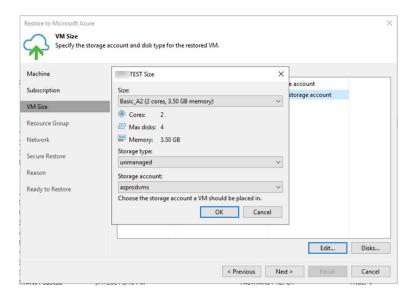
- 39. On the Subscription page, select a subscription from the Subscription dropdown list, whose resources you want to use from the Locations list and select a geographic region where you want to place the restored machine.
- 40. Select a geographic region with at least one storage account associated with the subscriptions and click Next.



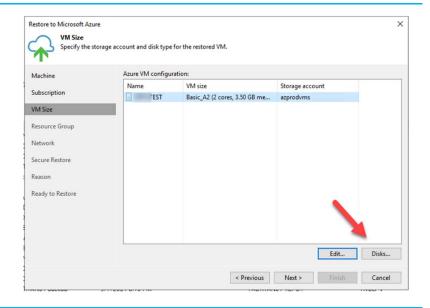
41. In the Azure VM
Configuration list, select
the machine and click
Edit.



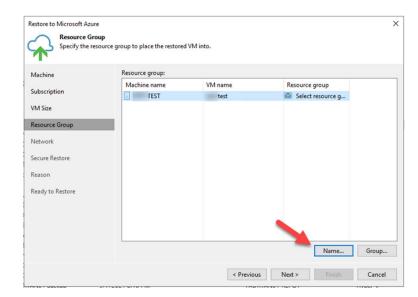
- 42. Select a size for the restored VM from the Size list. Next, select the right VM size corresponding to the initial machine configuration.
- 43. From the Storage account list, select a storage account whose resources you want to use to store the disks of the restored machine. The storage account must be compatible with the VM size you select. The list of storage accounts contains only general-purpose storage accounts. Click OK.



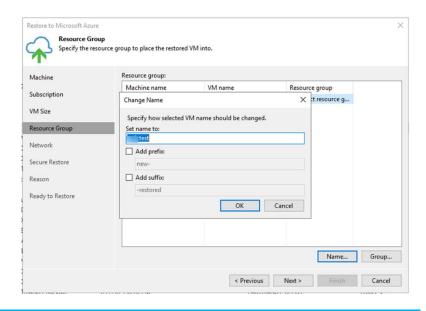
44. In the Azure VM
Configuration list, select
the machine and click
Disks.



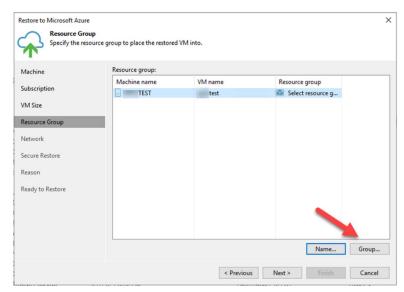
- 45. Select a disk and click the Disk Type button if you have selected the managed storage type. Then, in the Select Azure VM Disk Type window, select one of the following types: Standard HDD, Standard SSD or Premium SSD.
- 46. On the Resource page, by default, Veeam Backup & Replication restores a machine with its original name. However, if necessary, you can define a new name for the restored machine. Select the machine In the Resource group list and click the Name button.



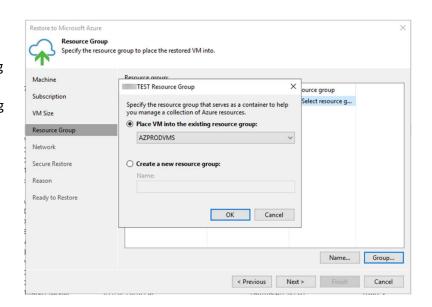
47. In the Change Name window, specify a change name rule — add a prefix and suffix to the original machine name — and click OK.



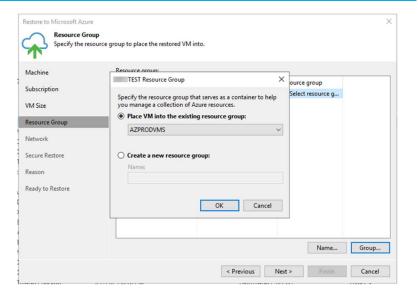
48. By default, Veeam Backup & Replication creates a new resource group for the restored machine and places the machine in it. However, if necessary, you can place the machine in an existing resource group. In the Resource group list, select the machine and click Group.



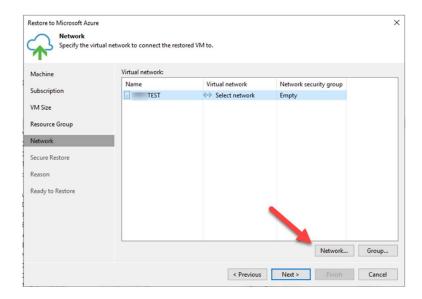
- 49. In the VM Resource
 Group window, select
 Place VM into the existing
 resource group to place
 the machine in an existing
 resource group or select
 Create a new resource
 group. For example,
 suppose you want to
 create a dedicated
 resource group for the
 restored machine.
- 50. Click OK.



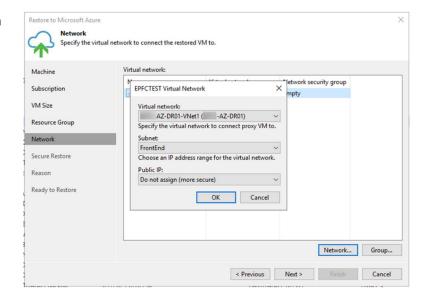
51. At the Resource Group page, click Next.



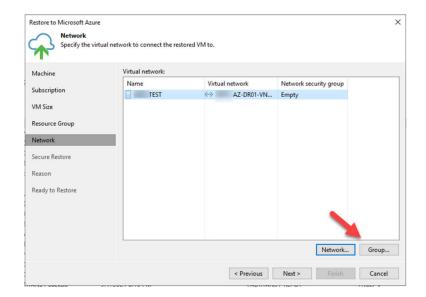
52. Select the machine in the Virtual network list and click Network.



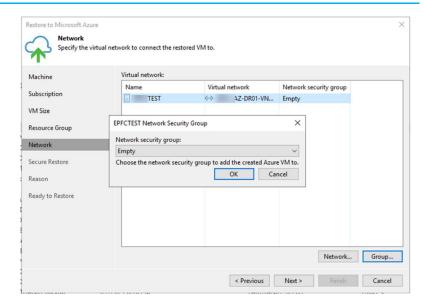
- 53. Select a network to which the machine must be connected from the Virtual network list.
- 54. From the Subnet list, select a subnet for the machine.
- 55. Veeam Backup &
 Replication can assign a
 public IP for the restored
 VM, which can be used
 for communications over
 the internet. By default, a
 public IP is not assigned.



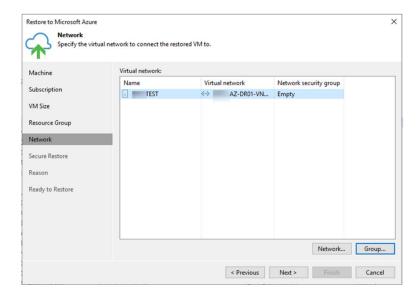
56. From the Virtual network list, select the machine and click Group.



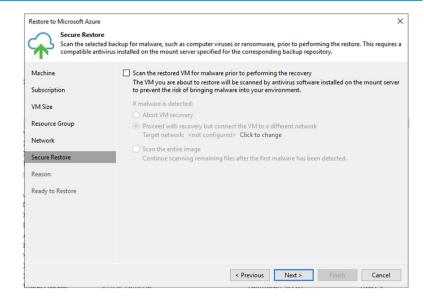
57. Select the network security group from the list and click OK. The restored machine will be added to the selected network security group. If you leave the field empty, Veeam Backup & Replication will create a new network security group.



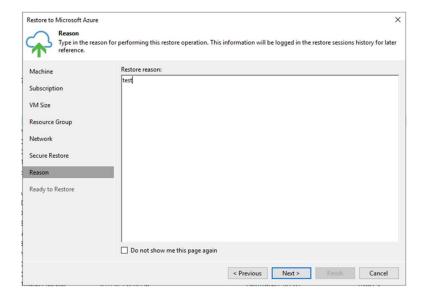
58. On the Network page, click Next.



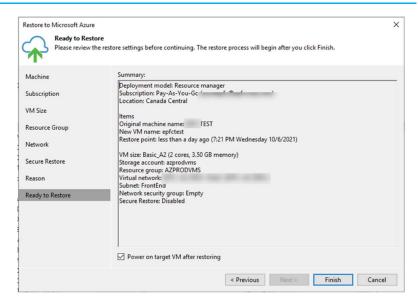
- 59. On the Secure Restore page, you can scan machine data with antivirus software before restoring the machine to Microsoft Azure
- 60. Click Next.



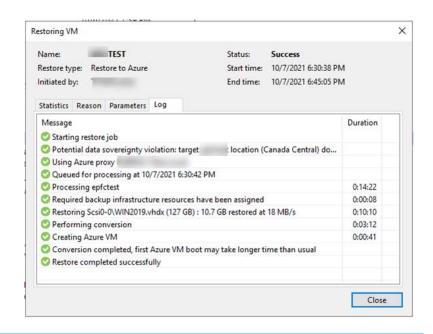
61. Enter a reason for restoring the machine on the Reason page and click Next.



62. At the Summary page, click Finish.



63. Make sure the restore status is Success.



Chapter 5

SureBackup

SureBackup is a Veeam technology that allows you to test VM backups and see if you can recover data from them. In addition, you can check any restore point of a backed-up VM.

During a SureBackup job, Veeam Backup & Replication performs "live" verification, which includes scanning the backed-up data for malware, booting the VM from the backup in an isolated environment, running tests for the VM, powering the VM off, and creating a report on recovery verification results.

Creating Application Group

The application group contains one or more VMs dependent on the verified VM. These VMs run applications and services that must be started before the verified VM can function correctly. For example, the application group usually includes a domain controller, DNS server, and DHCP server.

When you create an application group, you specify the role of each VM, as well as its boot priority and delay. You can also specify which tests must be run to validate the VMs in the application group.

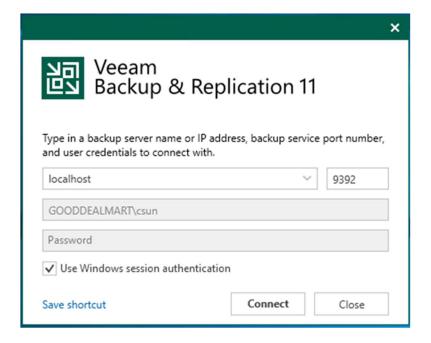
When a SureBackup job starts, Veeam Backup & Replication starts in the virtual lab VMs from the application group in the correct order and runs the necessary tests against them. Veeam Backup & Replication creates the required environment for the verified VM in this manner. Veeam Backup & Replication starts the verified VM in the virtual lab only after all VMs in the application group have been started and tested.

Note:

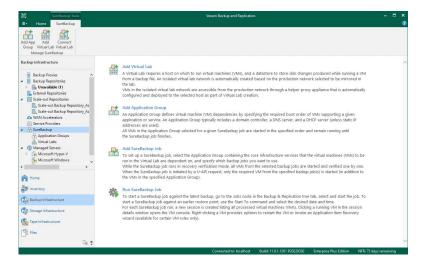
- A valid Enterprise edition Veeam Backup & Replication license must be installed on the backup server.
- If you intend to ping test VMs, the firewall on the tested VMs must allow ping requests.
- If you intend to use a heartbeat test to validate VMs, Hyper-V Integration Services must be installed in the tested VMs.
- To access the console of a verified VM, the backup server must have RDP client version 7.0 or later installed. The RDP client comes standard with Microsoft Windows 7 and later.

Instructions Screenshot (if applicable)

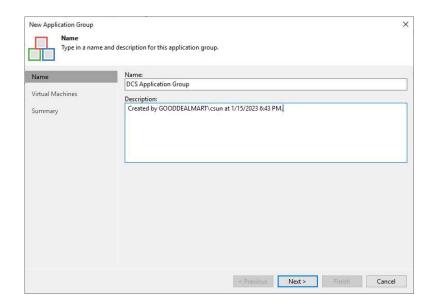
- Log in to the Veeam
 Backup and replication
 manager server.
- Open the Veeam Backup & Replication Console, and click Connect.



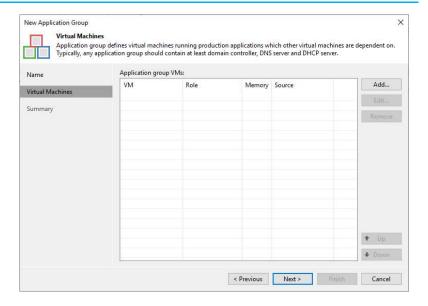
- 3. On the management console, Select Backup Infrastructure.
- On the Backup Infrastructure page, select SureBack and click Add Application Group.



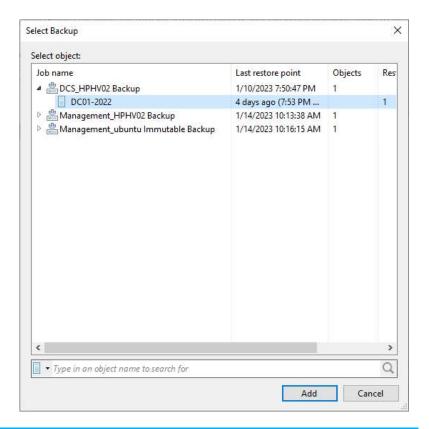
- 5. Enter the specified application group name on the Name page in the name field.
- Describe the group information in the Description, and click Next.



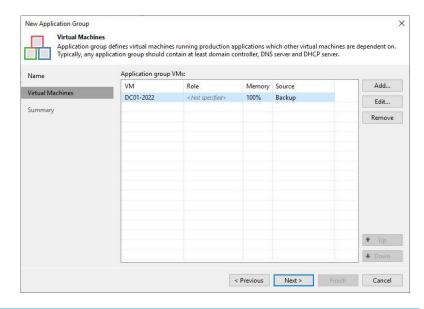
7. On the Virtual Machines page, click Add.



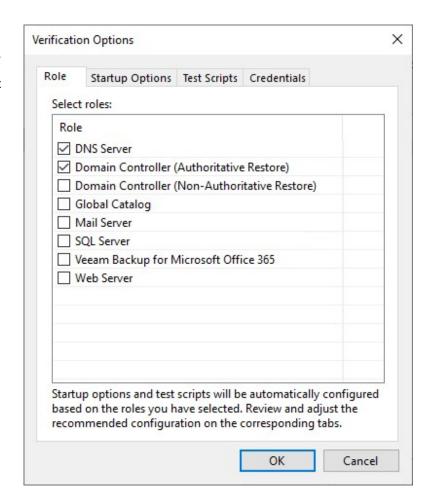
- 8. Expand the jobs name on the Select Backup page and select the machine.
- 9. Click Add.



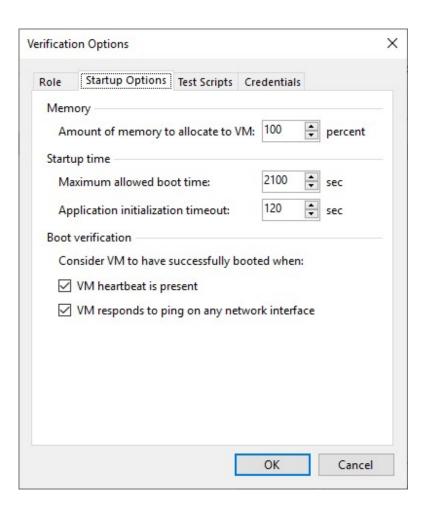
 Select the machine on the Virtual Machines page and click Edit.



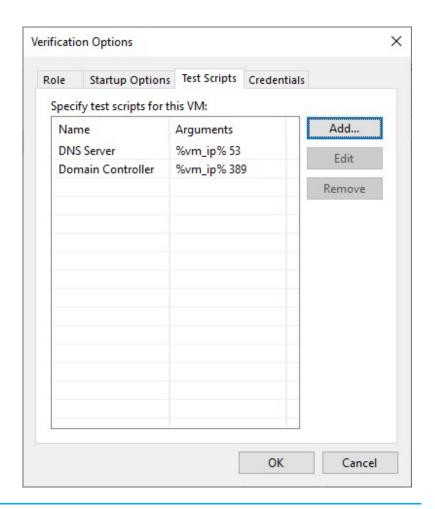
- 11. On the Verification Options page, select Role.
- 12. Select the roles check box in the Role.



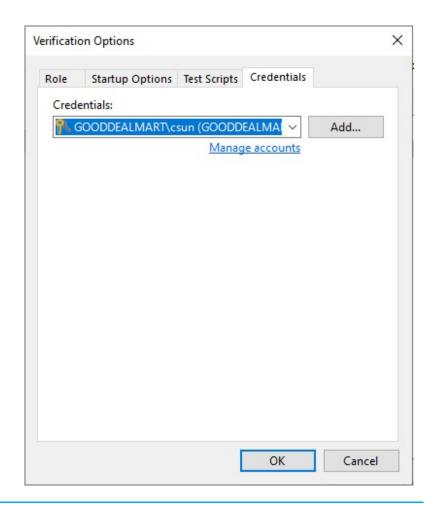
- 13. On the Verification Options page, select Startup Options.
- 14. On the Start Options page, In the Memory section, specify how much memory you want to pre-allocate to the VM when the system boots.
- 15. In the Startup time section, enter the maximum boot time for the VM and the timeout for the VM to initialize applications.
- In the Boot verification section, select the VM heartbeat is present checkbox.
- 17. Select VM responds to ping on any network interface checkbox.



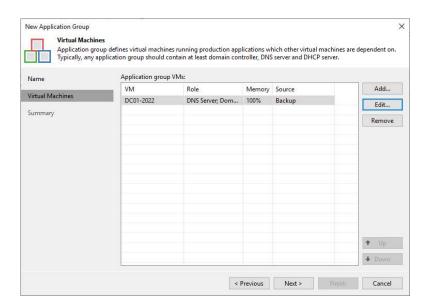
- 18. On the Verification
 Options page, select Test
 Scripts.
- 19. You can add or edit the script if needed.



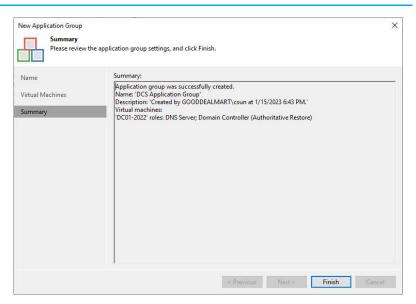
- 20. On the Verification Options page, select Credentials.
- 21. Select a user account from the Credentials drop-down list.
- 22. Click OK.



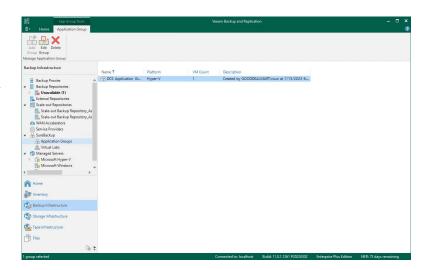
23. On the Virtual Machines page, click Next.



24. On the Summary page, click Finish.



- 25. On the Backup Infrastructure page, expand SureBackup and select Application Groups.
- 26. Ensure the new application group is created.



Configuring Basic Single-Host Virtual Labs

The virtual lab is a self-contained virtual environment. where Veeam Backup & Replication verifies virtual machines (VMs). With Veeam Backup & Replication, it is possible to start VMs from the application group and the verified VM in the virtual lab, allowing for seamless recovery and continuity of operations in case of a disaster or unexpected downtime.

The virtual lab itself does not necessitate the allocation of additional resources. VMs running in the virtual lab, on the other hand, consume CPU and memory resources on the Hyper-V host of the virtual lab. Therefore, all VM changes made during recovery verification are saved to the differencing disc (AVHD/AVHDX file) created by Veeam Backup & Replication for the recovered VM. The changes are discarded once the recovery verification process is completed.

The virtual lab is wholly isolated from the production environment. The virtual lab's network configuration is identical to the production environment's. For example, if verified VMs and application group VMs are located in two logical networks in the production environment, the virtual lab will also have two networks. The virtual lab networks will be mapped to the corresponding production networks.

The IP addresses of VMs in isolated networks are the same as those in the production network. This allows VMs in the virtual lab to function precisely as they were in the production environment.

The introductory single-host virtual lab is ideal if all the VMs that need to be verified, including the VMs from the application group and the backup server, are connected to the same network. This ensures efficient communication between them and simplifies the setup process.

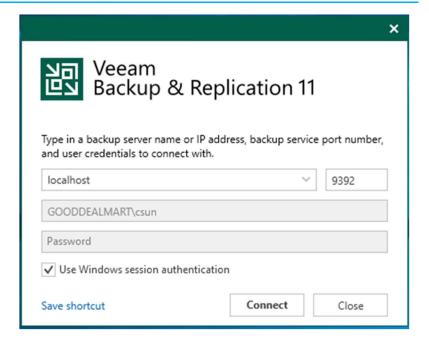
Veeam Backup & Replication creates one virtual network mapped to the production network for the introductory single-host virtual lab. In addition, Veeam Backup & Replication adds a virtual switch for the virtual lab. The new virtual switch is only available to VMs running in the virtual lab. There is no routing to other networks outside of the virtual lab.

Veeam Backup & Replication automatically configure all settings for the introductory single-host virtual lab. In addition, the proxy appliance is also automatically created and configured on the Hyper-V host where the virtual lab is created.

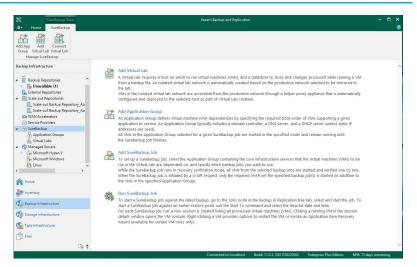
Instructions

Screenshot (if applicable)

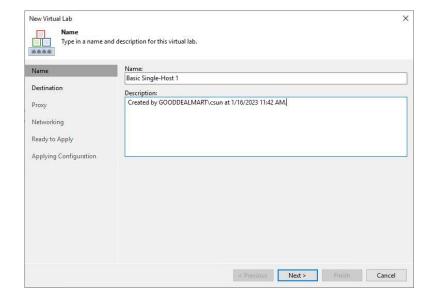
- Log in to the Veeam
 Backup and replication manager server.
- Open the Veeam Backup & Replication Console, and click Connect.



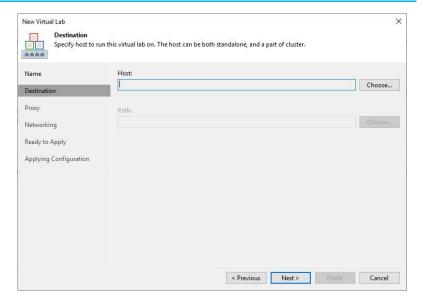
- 3. On the Home page, select Backup Infrastructure.
- 4. Select SureBackup and click Add Virtual Lab on the Backup Infrastructure page.



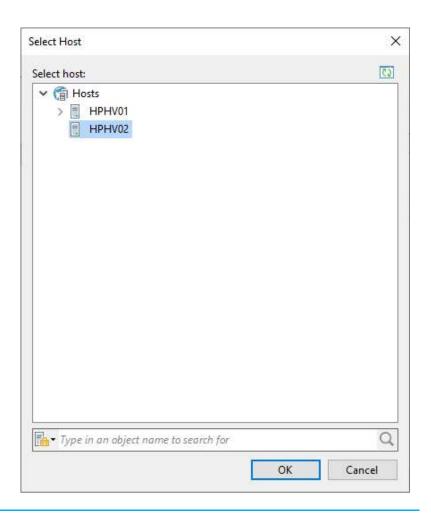
- 5. On the Name page, specify in the Name field.
- 6. In the Description field, describe future references.
- 7. Click Next.



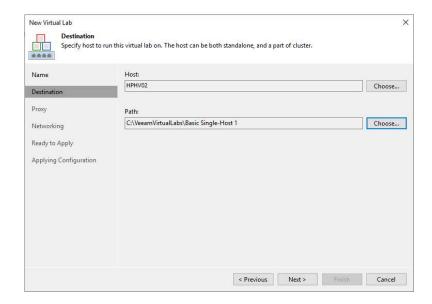
8. On the Destination page, click Choose in the Hose field.



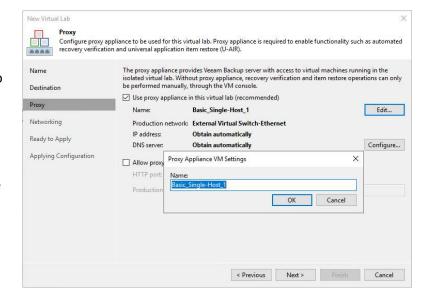
9. Select the host on the Select Host page, and click OK.



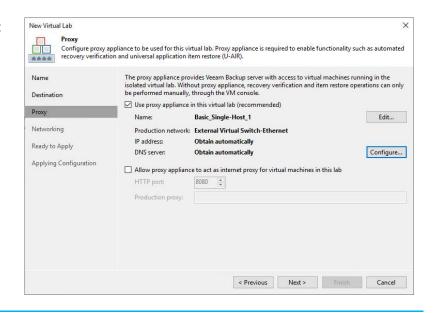
Select the path from Choose in the Path field, and click Next.



- 11. On the Proxy page, select the Use proxy appliance in this virtual lab (recommend) checkbox to enable automatic recovery verification of VMs.
- 12. The proxy appliance uses the same virtual lab name by default. However, you can click Edit to change the Name.



13. Click Configure and select a production network.

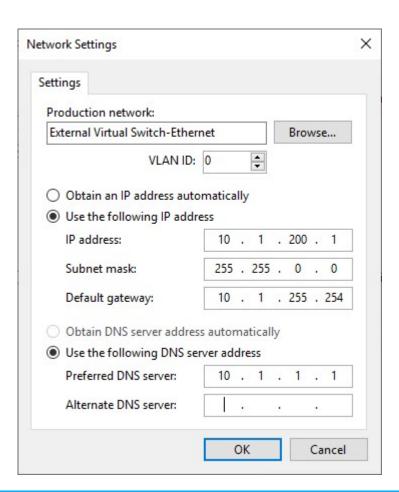


14. On the Network Settings page, enter the IP address of the proxy appliance in the production network and the DNS server settings.

15. Click OK.

Note:

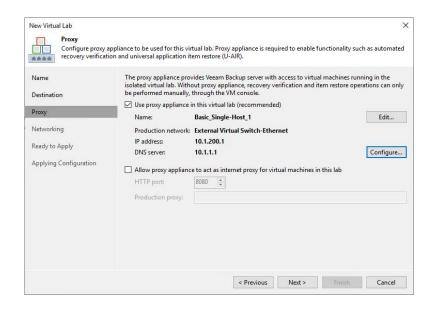
This ensures that the backup server can communicate with the proxy appliance and the VMs running in the virtual lab. Still, it's essential to keep in mind that the routing table of the backup server must be updated accordingly, depending on the IP address assigned to the proxy appliance.



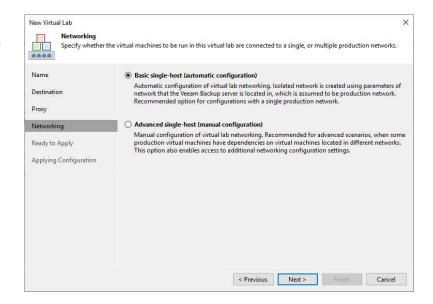
- 16. Select Allow proxy appliance to act as internet proxy for virtual machines in this lab checkbox if VMs in the virtual lab can connect to the internet.
- 17. Enter a port number for HTTP traffic in the Port field.
- 18. If the VMs need to access the Internet, enter the proxy server name in the Production proxy field.
- 19. Click Next.

Note:

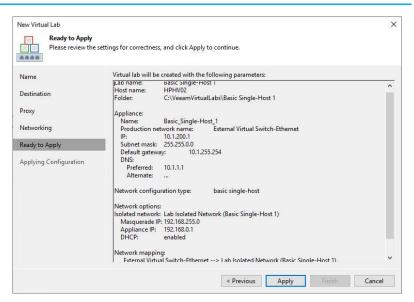
Allowing the proxy appliance to act as an internet proxy enables HTTP(S) internet access for virtual lab VMs. Other protocols (such as the ICMP protocol used for ping tests) are not proxied by the proxy appliance for VMs in the virtual lab.



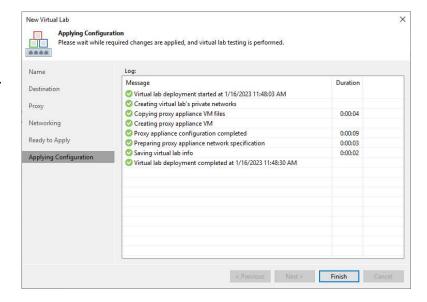
20. Select Basic single-host (automatic configuration) on the Networking page and click Next.



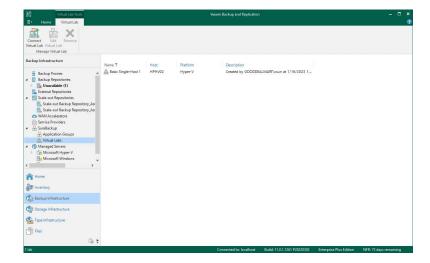
21. On the Ready to Apply page, click Apply.



- 22. On the Applying
 Configuration page,
 ensure the processes are
 completed and successful.
- 23. Click Finish.



- 24. On the Backup
 Infrastructure page,
 expand SureBackup and
 select Virtual Labs.
- 25. Ensure the new Virtual Lab is created.



Configuring Advanced Single-Host Virtual Labs

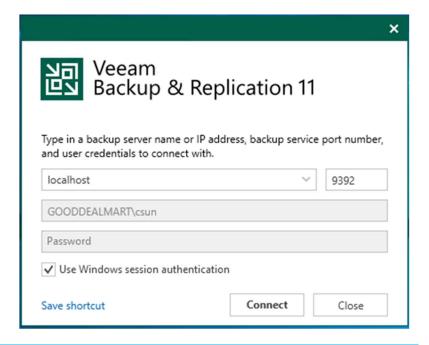
You can use the advanced single-host virtual lab if the VMs you want to test and the VMs in the application group are on different networks.

Veeam Backup & Replication creates several virtual networks for the advanced single-host virtual lab. The number of virtual networks corresponds to the number of verified VMs connected to production networks. The virtual lab networks are mapped to production networks.

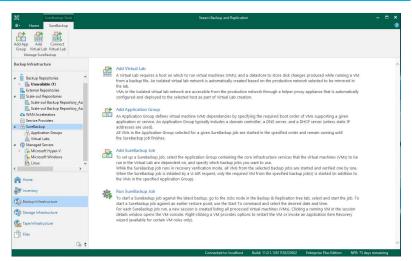
Every network in the virtual lab receives a new virtual switch from Veeam Backup & Replication. For example, if you have two networks in your production environment, Veeam Backup & Replication will create two networks in the virtual lab and two virtual switches for each network on the Hyper-V host. The virtual lab's VMs only uses the additional virtual switches. There is no routing to other networks outside of the virtual lab.

Instructions	Screenshot (if applicable)

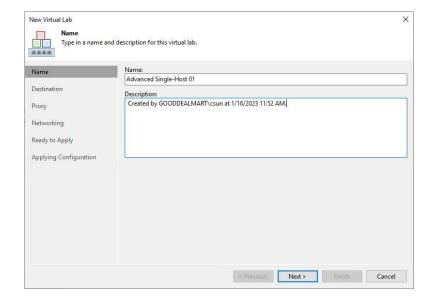
- Log in to the Veeam
 Backup and replication
 manager server.
- Open the Veeam Backup & Replication Console, and click Connect.



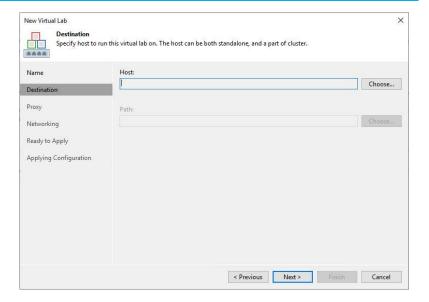
- 3. On the Home page, select Backup Infrastructure.
- 4. Select SureBackup and click Add Virtual Lab on the Backup Infrastructure page.



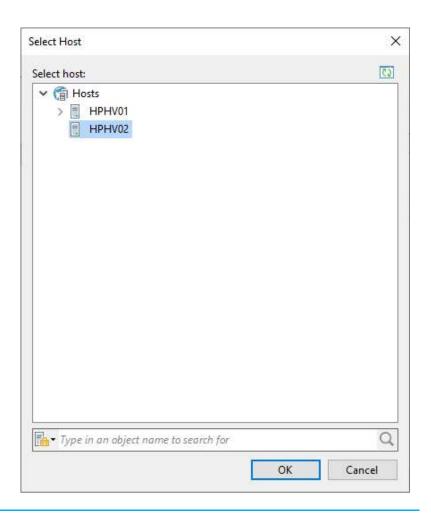
- 5. On the Name page, specify in the Name field.
- 6. In the Description field, describe future references.
- 7. Click Next.



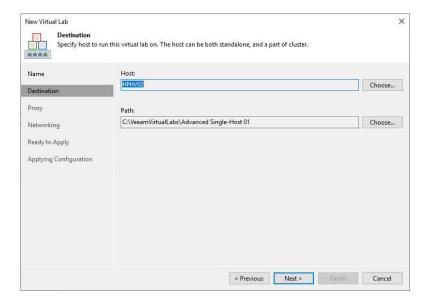
8. On the Destination page, click Choose in the Hose field.



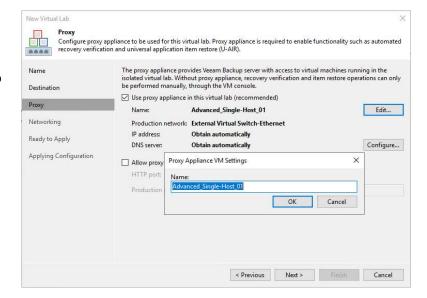
9. Select the host on the Select Host page, and click OK.



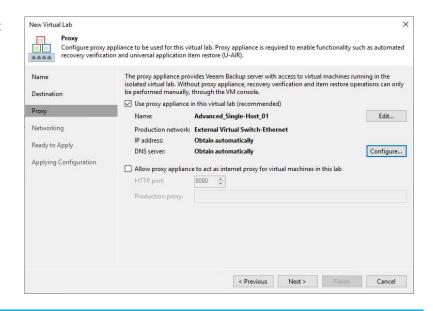
Select the path from Choose in the Path field, and click Next.



- 11. On the Proxy page, select the Use proxy appliance in this virtual lab (recommend) checkbox to enable automatic recovery verification of VMs.
- 12. The proxy appliance uses the same virtual lab name by default. However, you can click Edit to change the Name.



13. Click Configure and select a production network.

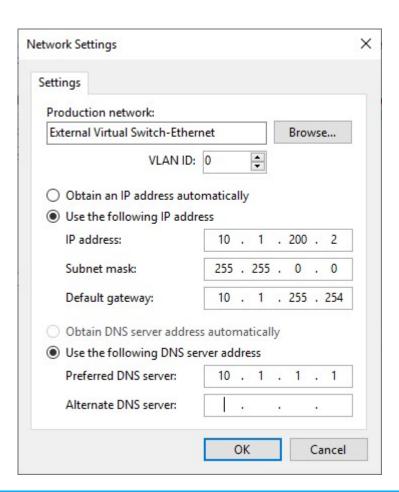


14. On the Network Settings page, enter the IP address of the proxy appliance in the production network and the DNS server settings.

15. Click OK.

Note:

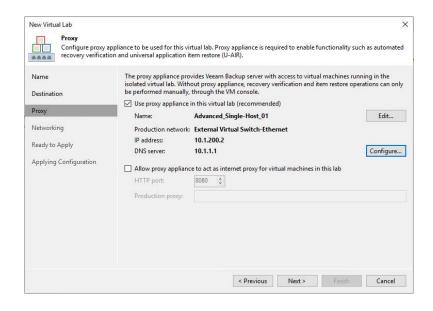
This ensures that the backup server can communicate with the proxy appliance and the VMs running in the virtual lab. Still, it's essential to keep in mind that the routing table of the backup server must be updated accordingly, depending on the IP address assigned to the proxy appliance.



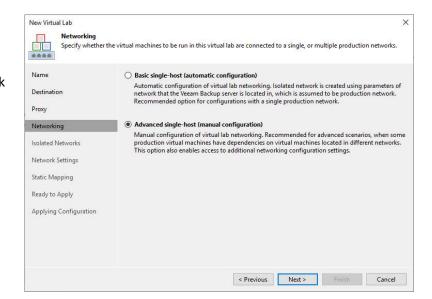
- 16. Select Allow proxy appliance to act as internet proxy for virtual machines in this lab checkbox if VMs in the virtual lab can connect to the internet.
- 17. Enter a port number for HTTP traffic in the Port field.
- 18. If the VMs need to access the Internet, enter the proxy server name in the Production proxy field.
- 19. Click Next.

Note:

Allowing the proxy appliance to act as an internet proxy enables HTTP(S) internet access for virtual lab VMs. Other protocols (such as the ICMP protocol used for ping tests) are not proxied by the proxy appliance for VMs in the virtual lab.

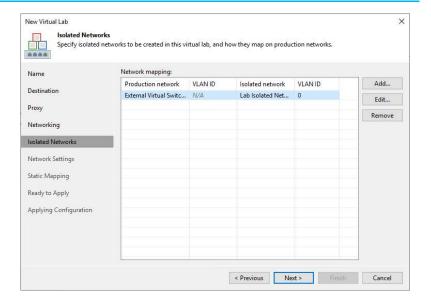


20. Select Advanced singlehost (manual configuration) on the Networking page and click Next.

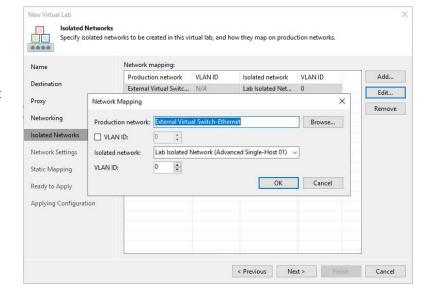


21. Select the existing

Network mapping on the
Isolated Networks page,
and click Edit.



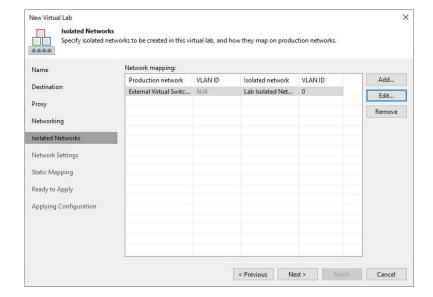
- 22. On the Networking Mapping page, enter a name for isolated and select a production network from the list that contains VMs from the application group and verified VMs.
- 23. Select the VLAN ID check box and enter an ID number if necessary.
- 24. Select an isolated network from the drop-down list.
- 25. Select the VLAN ID check box and enter an ID number if necessary.



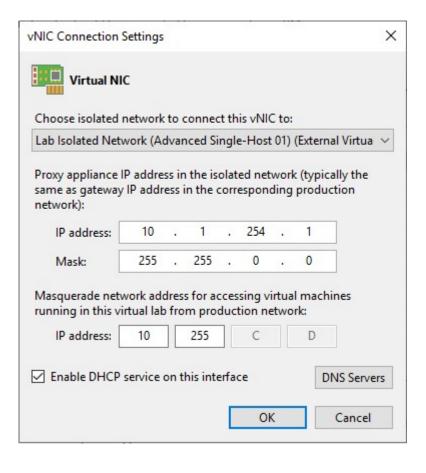
26. On the Isolate Networks page, click Next.

Note:

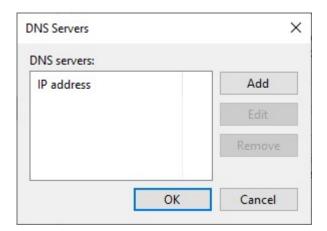
Several production networks can be mapped to the same isolated network. However, the production networks you intend to map must use the same network masks and IP address pools.



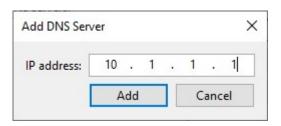
- 27. Choose an isolated network to connect this vNIC from the drop-down list on the Virtual NIC page.
- 28. Assign the IP address and Mask to the proxy appliance in the isolated network.
- 29. This allows for seamless connectivity between the VMs in the virtual lab and the production network, as the masquerade IP address is used to mask the actual IP address of the VMs running in the virtual lab, making them appear as if they are running on the production network. In cases where a specific IP address range needs to be used, the masquerade network IP address can be easily changed in the virtual lab settings.
- Select Enable DHCP service on this interface check box, and click DNS Servers.



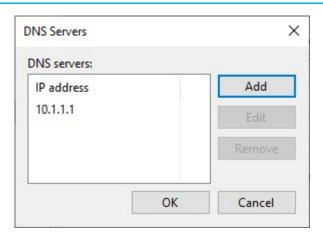
31. On the DNS Servers page, click Add.



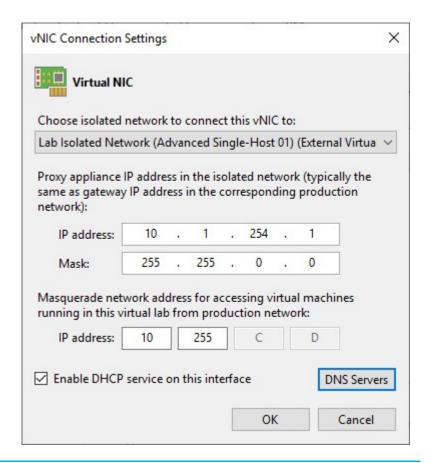
32. Enter the DNS Server IP address on the Add DNS Server page and click Add.



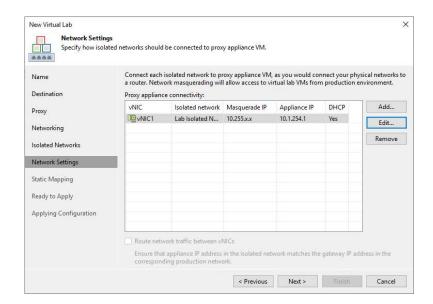
33. On the DNS Servers page, click OK.



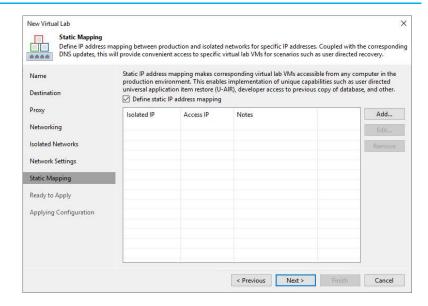
34. On the Virtual NIC page, click OK.



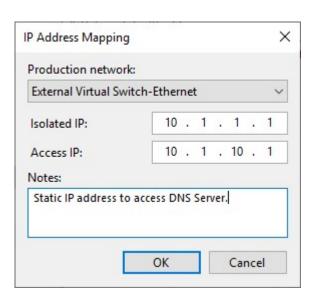
- 35. On the Network Settings page, click Next.
- 36. Select the Route network traffic between vNICs checkbox if you have multiple vNICs.

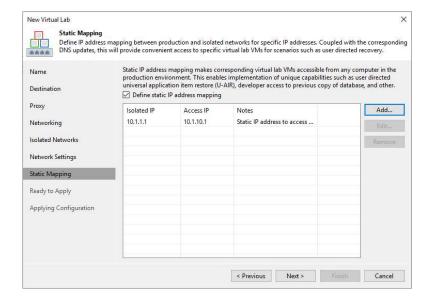


37. Select Define static IP address mapping on the Static page and click Add if necessary.

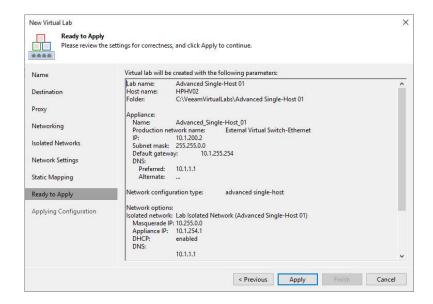


- 38. Select the production network from the drop-down list on the IP address Mapping page.
- 39. In the Isolated IP field, enter an IP address of the VM in the production.
- 40. In the Access IP field, enter the production network's IP address to access the virtual lab VM.
- 41. Describe the future references in the Notes.
- 42. Click OK.
- 43. On the Static Mapping page, click Next.

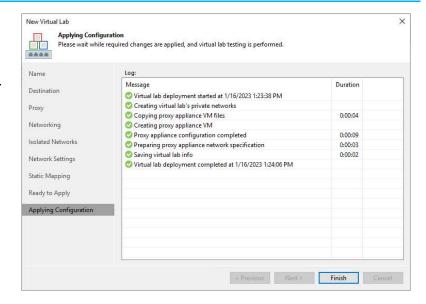




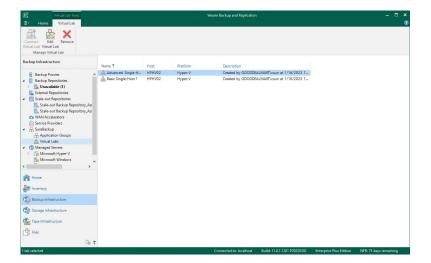
44. Review the parameters on the Ready to Apply page, and click Apply.



- 45. On the Applying
 Configuration page,
 ensure the processes are
 completed and successful.
- 46. Click Finish.



- 47. On the Backup
 Infrastructure page,
 expand SureBackup and
 select Virtual Labs.
- 48. Ensure the new Virtual Lab is created.



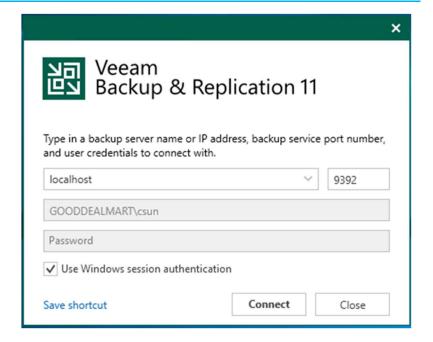
Creating a SureBackup Job

A SureBackup job is a recovery verification task. The SureBackup job collects all recovery verification task settings and policies, such as the application group and virtual lab to be used, VM backups that must be verified in the virtual lab, and so on. SureBackup can be run manually or scheduled to run automatically.

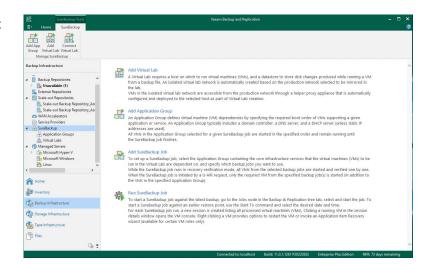
When the verification process is finished, the VMs in the application group are powered down. Optionally leave the application group's VMs running to perform manual testing or enable user-directed application item-level recovery.

Instructions Screenshot (if applicable)

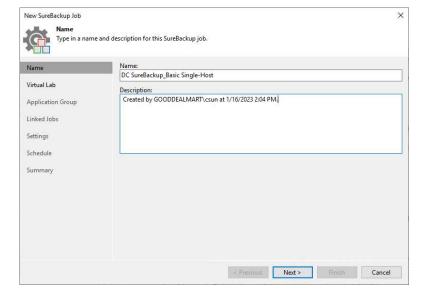
- Log in to the Veeam
 Backup and replication
 manager server.
- Open the Veeam Backup & Replication Console, and click Connect.



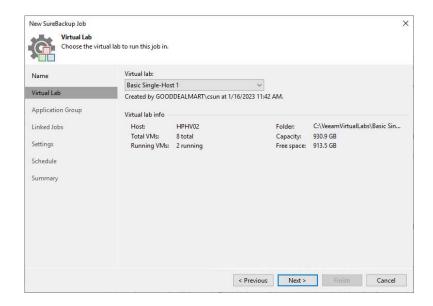
- 3. On the Home page, select Backup Infrastructure.
- 4. Select SureBackup and click Add SureBackup Job on the Backup Infrastructure page.



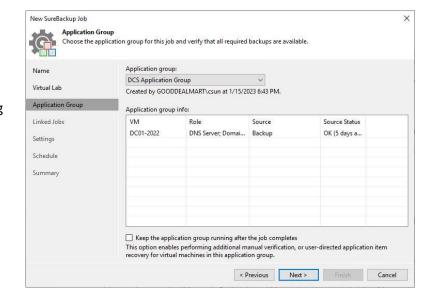
- 5. On the Name page, specify in the Name field.
- 6. In the Description field, describe future references.
- 7. Click Next.



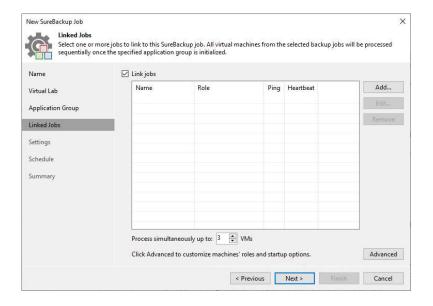
8. Select the Basic Single-Host virtual lab from the drop-down list on the Virtual Lab page and click Next.



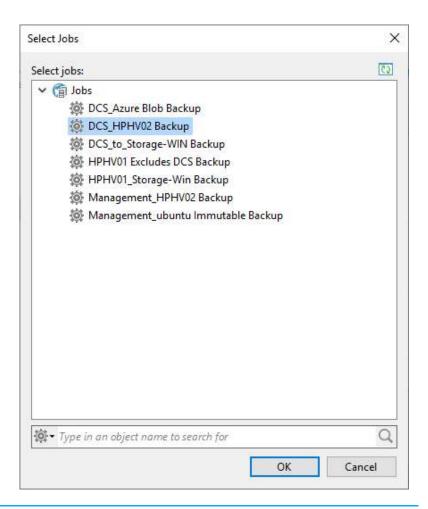
- Select the application group from the dropdown list.
- Select Keep the application group running after the job completes if necessary.
- 11. Click Next.



12. Select the Link jobs checkbox on the Linked jobs page and click Add.



13. On the Select page, select the job and click OK.

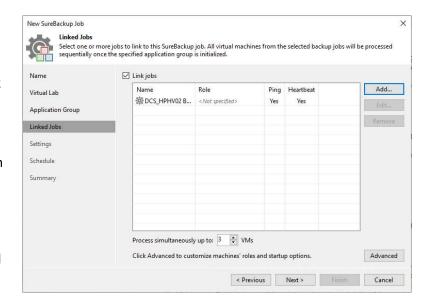


14. On the Link jobs page, change the number of processes simultaneously VMs if necessary and click Advanced.

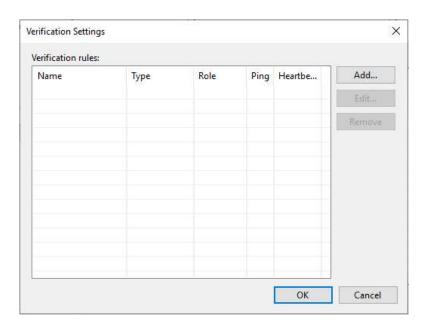
Note:

By default, you can launch and test up to three virtual machines simultaneously. You can also increase the number of VMs started and tested simultaneously. However, if these VMs are resource intensive, the performance of the SureBackup job and the Hyper-V host on which the virtual lab is hosted may suffer.

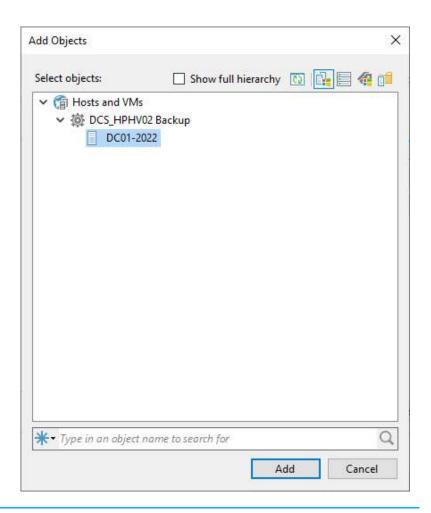
15. Click Advanced.



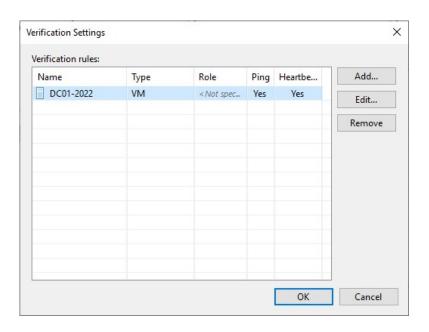
16. On the Verification Settings page, click Add.



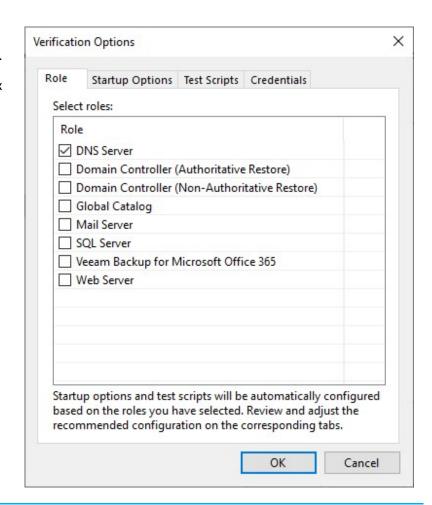
17. On the Add Objects, select the Machine and click Add.



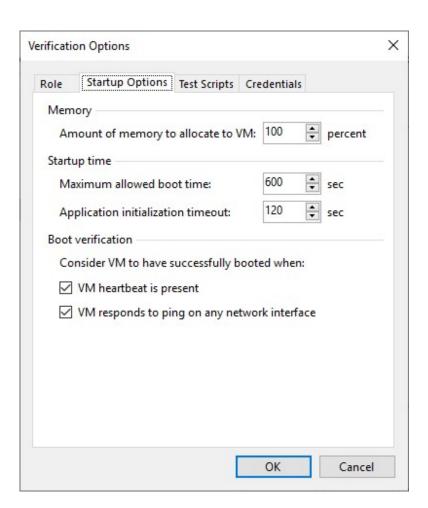
18. Select the machine and click Edit on the Verification Settings page.



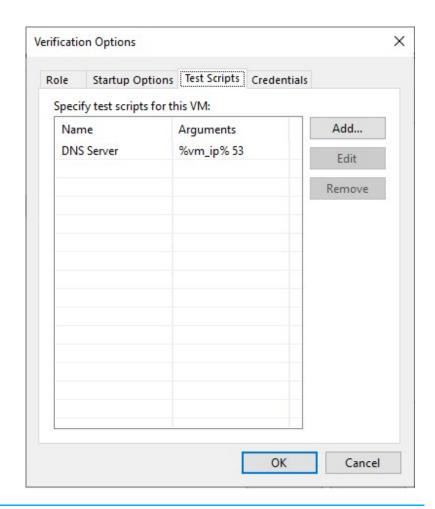
- 19. On the Verification Options page, select Role.
- 20. Select the roles check box in the Role.



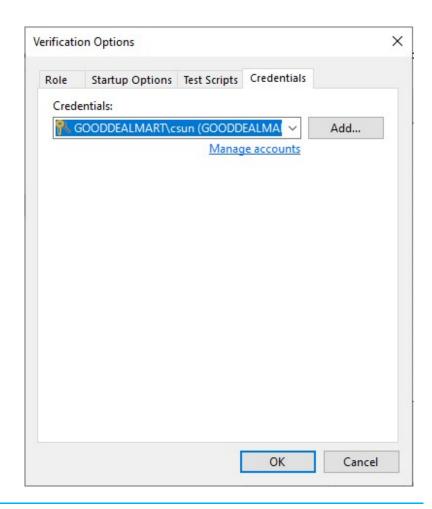
- 21. On the Verification Options page, select Startup Options.
- 22. On the Start Options page, In the Memory section, specify how much memory you want to pre-allocate to the VM when the system boots.
- 23. In the Startup time section, enter the maximum boot time for the VM and the timeout for the VM to initialize applications.
- 24. In the Boot verification section, select the VM heartbeat is present checkbox.
- 25. Select VM responds to ping on any network interface checkbox.



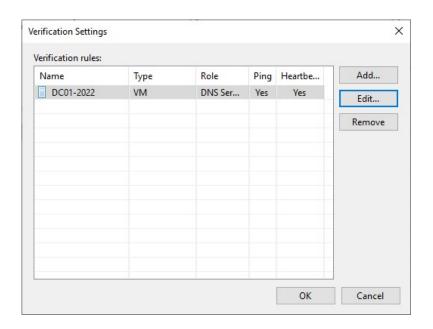
- 26. On the Verification
 Options page, select Test
 Scripts.
- 27. You can add or edit the script if needed.



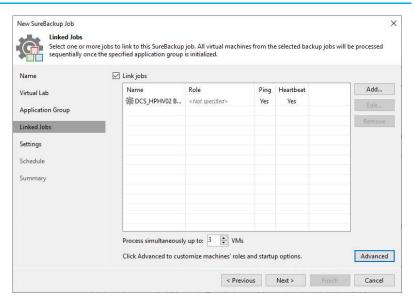
- 28. On the Verification Options page, select Credentials.
- 29. Select a user account from the Credentials drop-down list.
- 30. Click OK.



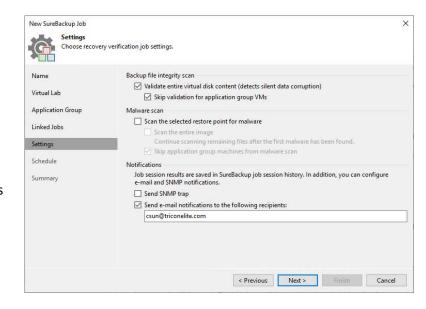
31. On the Verification Settings page, click OK.



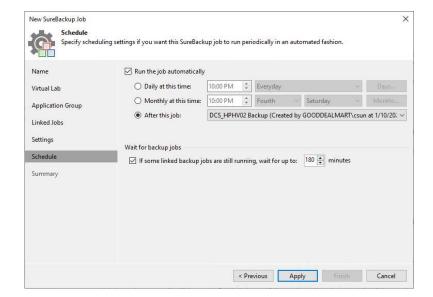
32. On the Linked Jobs page, click Next.



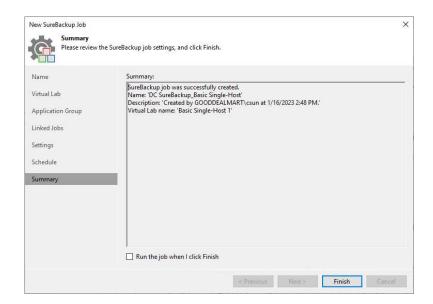
- 33. On the Settings page,
 Select the Validate entire
 virtual disc contents
 checkbox to validate the
 backup file with a CRC
 check to ensure it is not
 corrupted.
- 34. Select the Skip validation for application group VMs checkbox to exclude VMs being a part of the application group from this test.
- 35. Select the Scan the selected restore point for malware check box if you want Veeam Backup & Replication to scan VM data with antivirus software.
- 36. Select the Scan the entire image checkbox if you want the antivirus software to continue scanning VM data after the first malware is detected.
- 37. Select the Skip application group machines from the malware scan check box if you do not want to scan VMs in the application group.



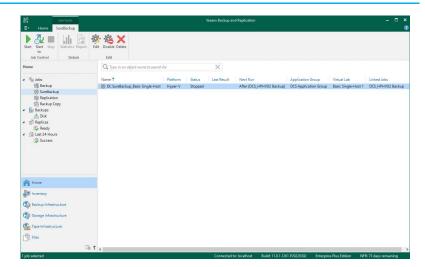
- 38. Select the Send SNMP trap check box to receive SNMP traps.
- 39. Select Send email notifications to the following recipients checkbox if you want to receive notifications via email.
- 40. Enter the recipient's email address and click Next.
- 41. On the Schedule page, select Run the job automatically.
- 42. Select After this job and choose the preceding job from the list.
- 43. Select If some linked backup jobs are still running, wait until the checkbox and enter the necessary period in the field on the right. For example, suppose the linked job is still running.
- 44. Click Apply.



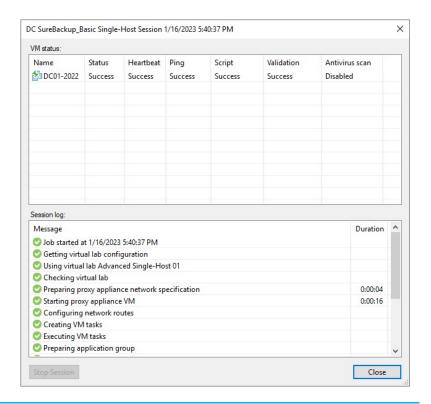
45. On the Summary page, click Finish.



- 46. On the Home page, expand Jobs and select SureBackup.
- 47. Ensure the new SureBackup job is created.



48. Verify the SureBackup job result.



Chapter 6

Creating On-Demand Sandbox

You can use an On-Demand Sandbox to perform tests on production VMs. The On-Demand Sandbox is a virtual environment where you can launch multiple VMs from backups. The On-Demand Sandbox can be used for the following tasks:

- Troubleshoot problems with VMs
- Test software patches and upgrades
- Install new software

The On-Demand Sandbox uses a virtual lab, which is utterly different from the production environment. Veeam Backup & Replication uses Instant Recovery to Microsoft Hyper-V to start a VM in the virtual lab. When you finish the On-Demand Sandbox, Veeam Backup & Replication unpublishes the VM and shuts down the virtual lab.

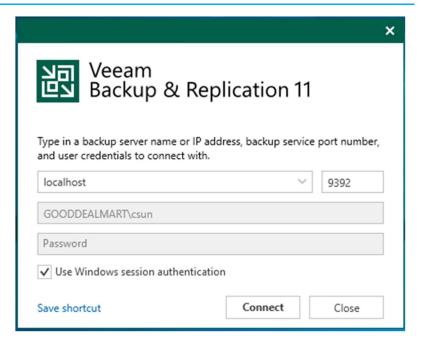
Configuring On-Demand Sandbox

You must configure the following objects to create the On-Demand Sandbox.

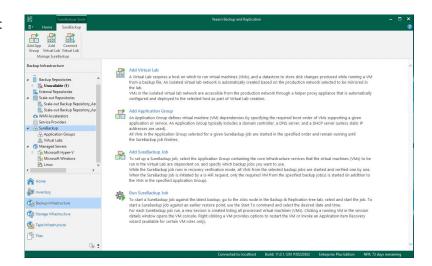
- Virtual lab
- Application Group
- SureBackup job

Instructions Screenshot (if applicable)

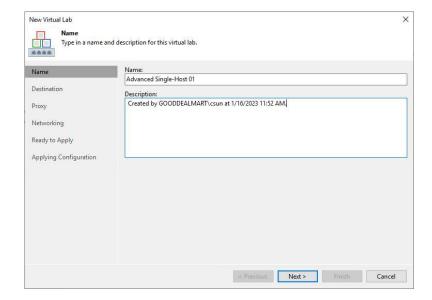
- Log in to the Veeam
 Backup and replication
 manager server.
- Open the Veeam Backup & Replication Console, and click Connect.



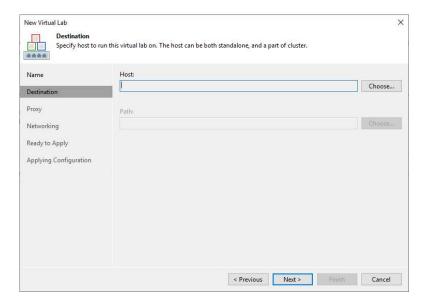
- 3. On the Home page, select Backup Infrastructure.
- 4. Select SureBackup and click Add Virtual Lab on the Backup Infrastructure page.



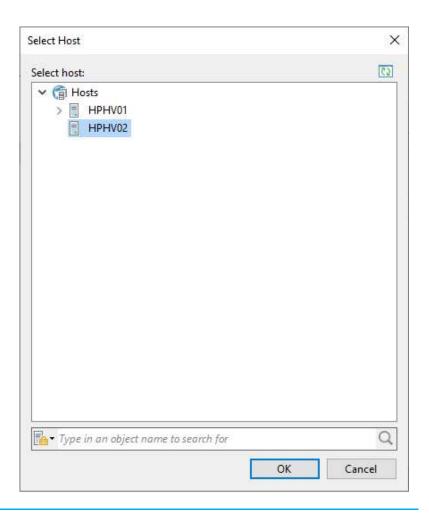
- 5. On the Name page, specify in the Name field.
- 6. In the Description field, describe future references.
- 7. Click Next.



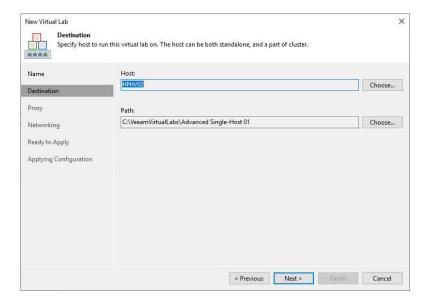
8. On the Destination page, click Choose in the Hose field.



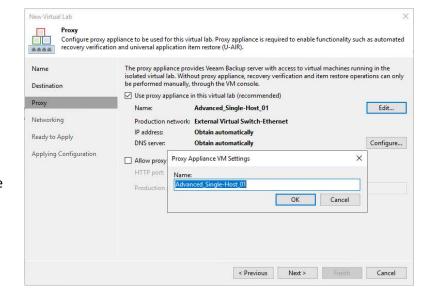
9. Select the host on the Select Host page, and click OK.



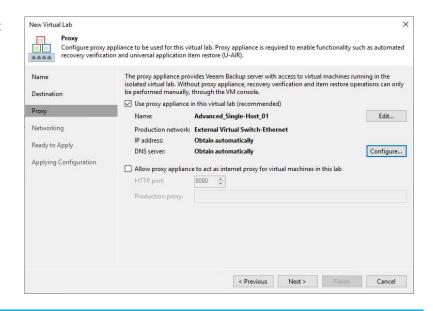
Select the path from Choose in the Path field, and click Next.



- 11. On the Proxy page, select the Use proxy appliance in this virtual lab (recommend) check box to enable automatic recovery verification of VMs.
- 12. The proxy appliance uses the same virtual lab name by default. However, you can click Edit to change the Name.



13. Click Configure and select a production network.

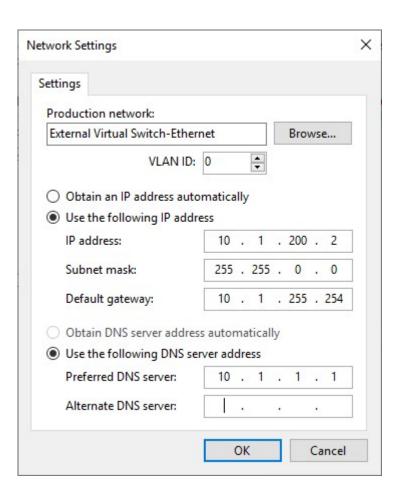


14. On the Network Settings page, enter the IP address of the proxy appliance in the production network and the DNS server settings.

15. Click OK.

Note:

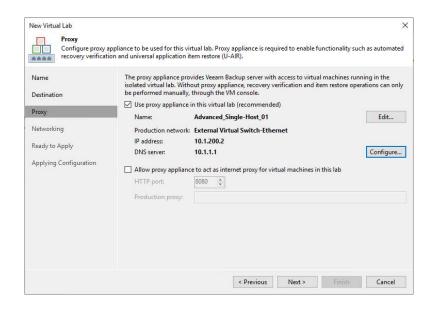
This ensures that the backup server can communicate with the proxy appliance and the VMs running in the virtual lab. Still, it's essential to keep in mind that the routing table of the backup server must be updated accordingly, depending on the IP address assigned to the proxy appliance.



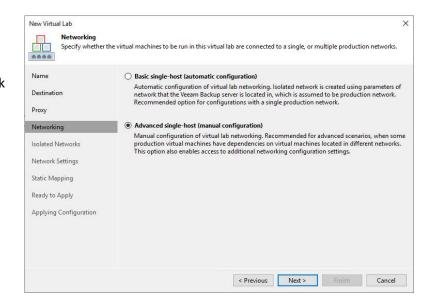
- 16. Select Allow proxy appliance to act as internet proxy for virtual machines in this lab checkbox if VMs in the virtual lab can connect to the internet.
- 17. Enter a port number for HTTP traffic in the Port field.
- 18. If the VMs need to access the Internet, enter the proxy server name in the Production proxy field.
- 19. Click Next.

Note:

Allowing the proxy appliance to act as an internet proxy enables HTTP(S) internet access for virtual lab VMs. Other protocols (such as the ICMP protocol used for ping tests) are not proxied by the proxy appliance for VMs in the virtual lab.

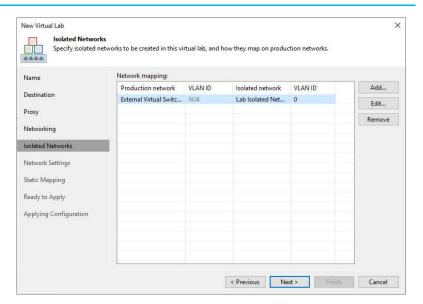


20. Select Advanced singlehost (manual configuration) on the Networking page and click Next.

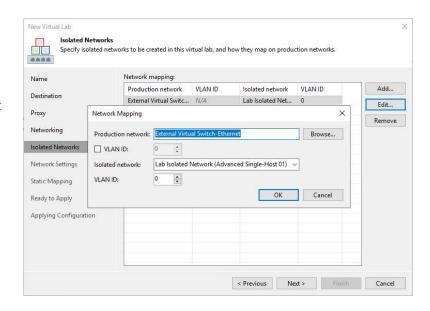


21. Select the existing

Network mapping on the
Isolated Networks page,
and click Edit.



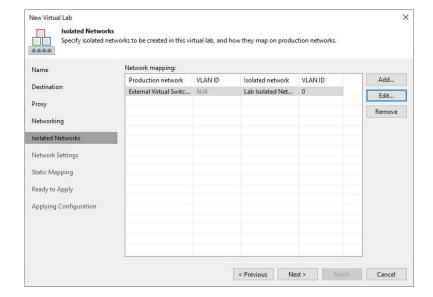
- 22. On the Networking
 Mapping page, enter a
 name for isolated and
 select a production
 network from the list that
 contains VMs from the
 application group and
 verified VMs.
- 23. Select the VLAN ID check box and enter an ID number if necessary.
- 24. Select an isolated network from the drop-down list.
- 25. Select the VLAN ID check box and enter an ID number if necessary.



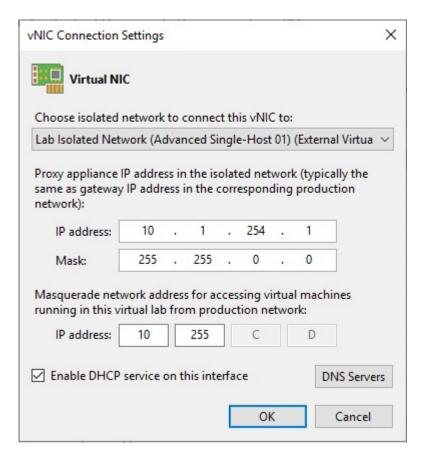
26. On the Isolate Networks page, click Next.

Note:

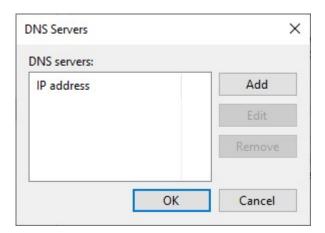
Several production networks can be mapped to the same isolated network. However, the production networks you intend to map must use the same network masks and IP address pools.



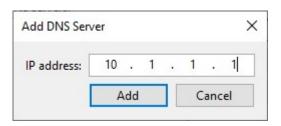
- 27. Choose an isolated network to connect this vNIC from the drop-down list on the Virtual NIC page.
- 28. Assign the IP address and Mask to the proxy appliance in the isolated network.
- 29. This allows for seamless connectivity between the VMs in the virtual lab and the production network, as the masquerade IP address is used to mask the actual IP address of the VMs running in the virtual lab, making them appear as if they are running on the production network. In cases where a specific IP address range needs to be used, the masquerade network IP address can be easily changed in the virtual lab settings.
- Select Enable DHCP service on this interface check box, and click DNS Servers.



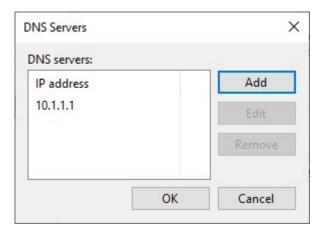
31. On the DNS Servers page, click Add.



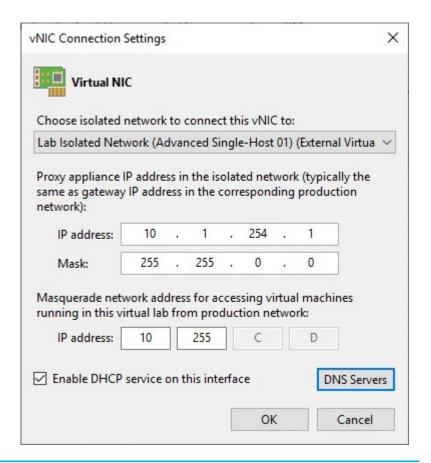
32. Enter the DNS Server IP address on the Add DNS Server page and click Add.



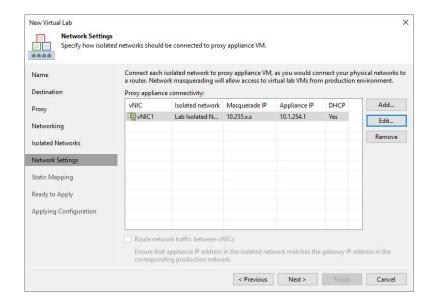
33. On the DNS Servers page, click OK.



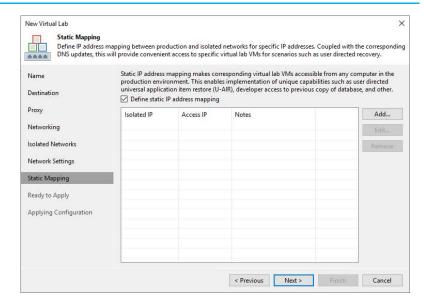
34. On the Virtual NIC page, click OK.



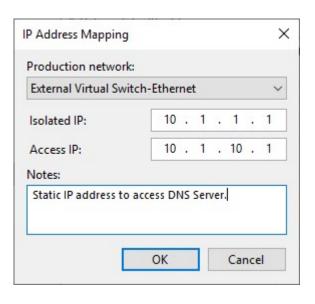
- 35. On the Network Settings page, click Next.
- 36. Select the Route network traffic between vNICs checkbox if you have multiple vNICs.

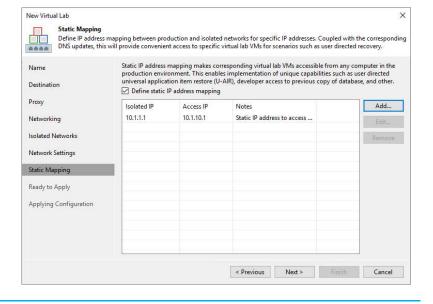


37. Select Define static IP address mapping on the Static page and click Add if necessary.

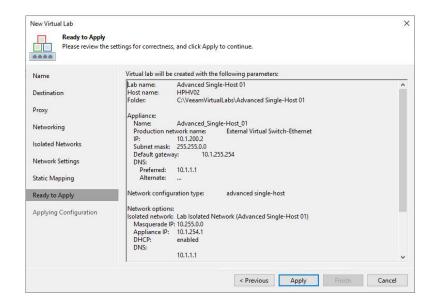


- 38. Select the production network from the drop-down list on the IP address Mapping page.
- 39. In the Isolated IP field, enter an IP address of the VM in the production.
- 40. In the Access IP field, enter the production network's IP address to access the virtual lab VM.
- 41. Describe the future references in the Notes.
- 42. Click OK.
- 43. On the Static Mapping page, click Next.

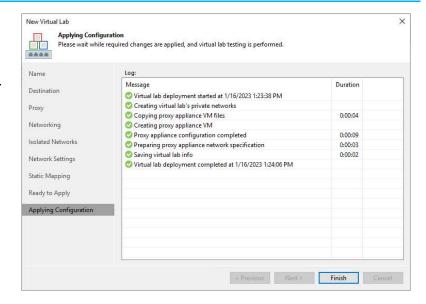




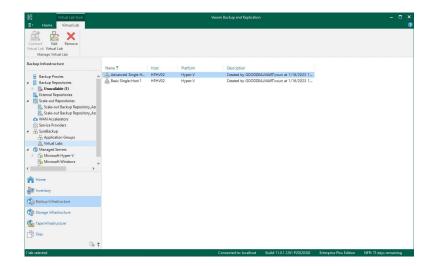
44. Review the parameters on the Ready to Apply page, and click Apply.



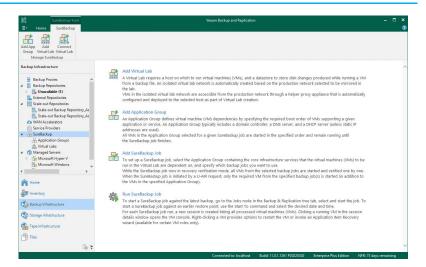
- 45. On the Applying
 Configuration page,
 ensure the processes are
 completed and successful.
- 46. Click Finish.



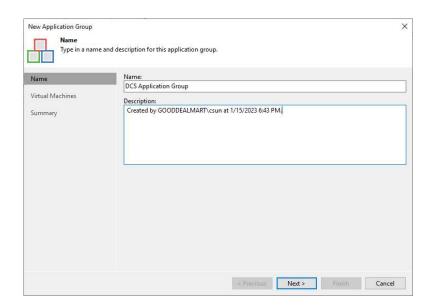
- 47. On the Backup
 Infrastructure page,
 expand SureBackup and
 select Virtual Labs.
- 48. Ensure the new Virtual Lab is created.



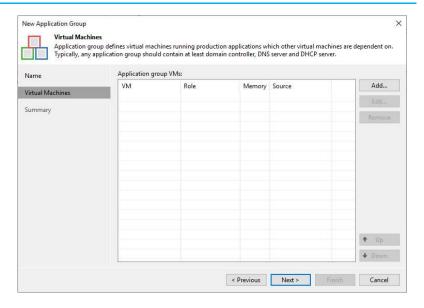
- 49. On the management console, Select Backup Infrastructure.
- 50. On the Backup
 Infrastructure page, select
 SureBack and click Add
 Application Group.



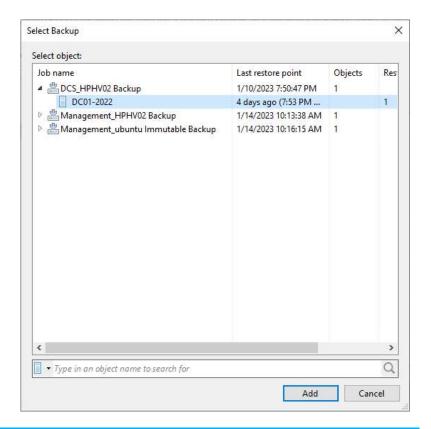
- 51. Enter the specified application group name on the Name page in the name field.
- 52. Describe the group information in the Description, and click Next.



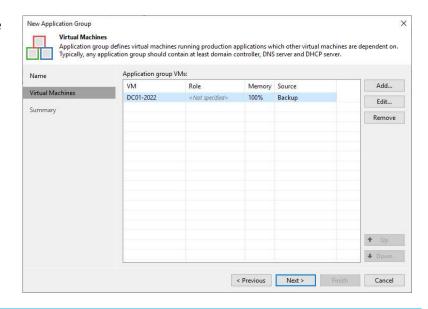
53. On the Virtual Machines page, click Add.



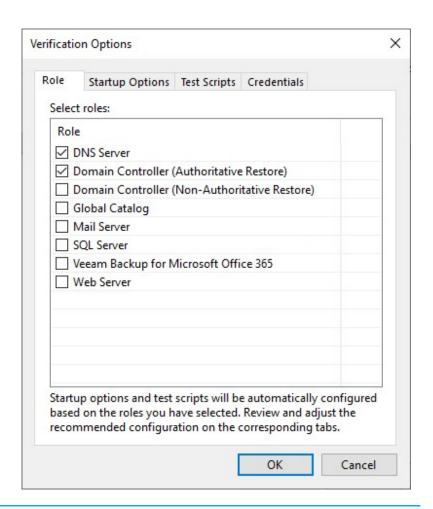
- 54. Expand the jobs name on the Select Backup page and select the machine.
- 55. Click Add.



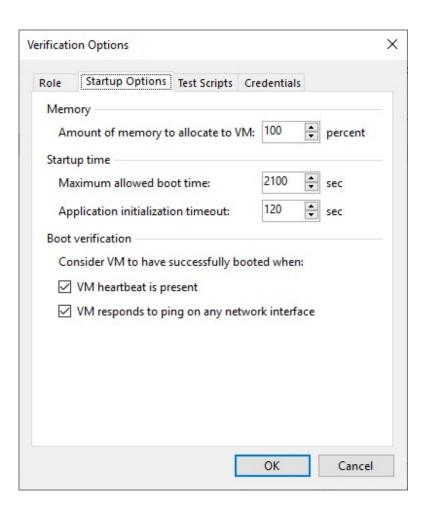
56. Select the machine on the Virtual Machines page and click Edit.



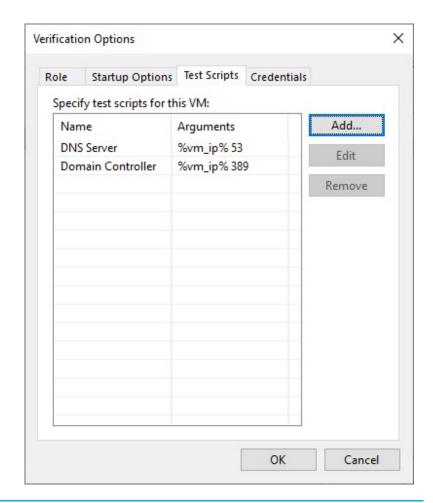
- 57. On the Verification Options page, select Role.
- 58. Select the roles check box in the Role.



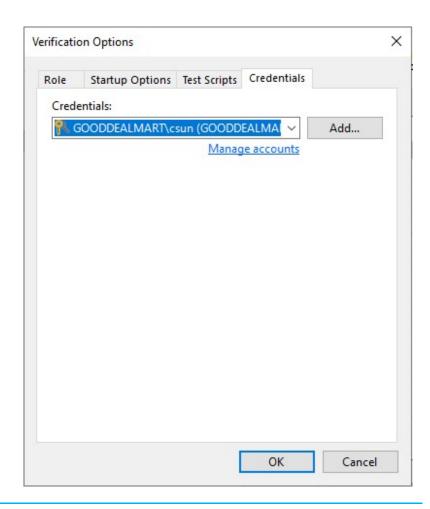
- 59. On the Verification Options page, select Startup Options.
- 60. On the Start Options page, In the Memory section, specify how much memory you want to pre-allocate to the VM when the system boots.
- 61. In the Startup time section, enter the maximum boot time for the VM and the timeout for the VM to initialize applications.
- 62. In the Boot verification section, select the VM heartbeat is present checkbox.
- 63. Select VM responds to ping on any network interface checkbox.



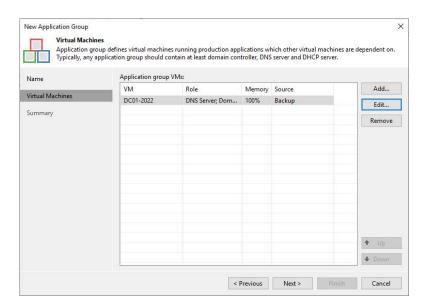
- 64. On the Verification
 Options page, select Test
 Scripts.
- 65. You can add or edit the script if needed.



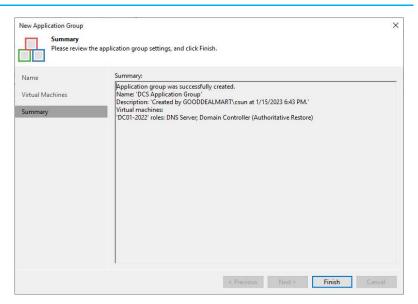
- 66. On the Verification Options page, select Credentials.
- 67. Select a user account from the Credentials drop-down list.
- 68. Click OK.



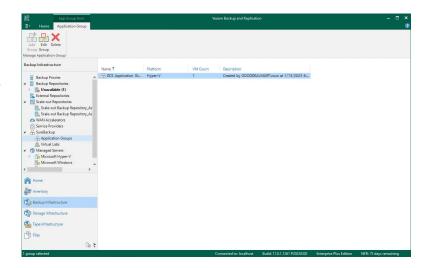
69. On the Virtual Machines page, click Next.



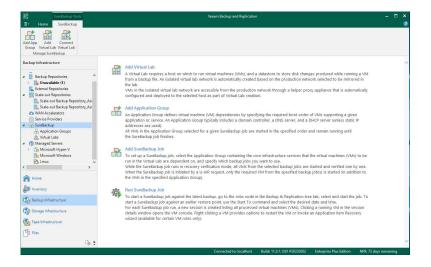
70. On the Summary page, click Finish.



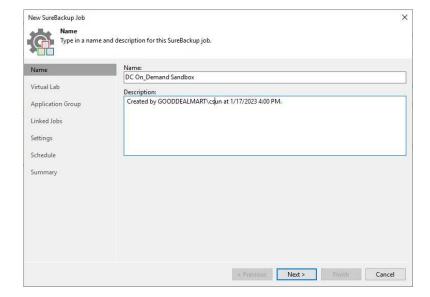
- 71. On the Backup
 Infrastructure page,
 expand SureBackup and
 select Application Groups.
- 72. Ensure the new application group is created.



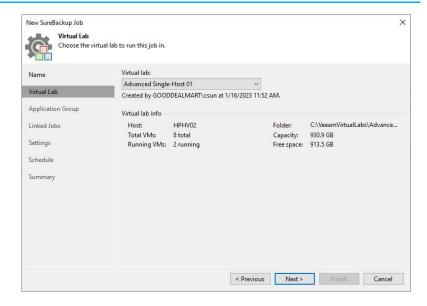
- 73. On the Home page, select Backup Infrastructure.
- 74. Select SureBackup and click Add SureBackup Job on the Backup Infrastructure page.



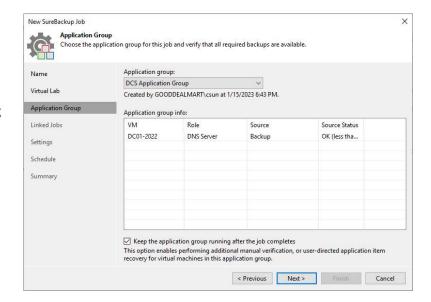
- 75. On the Name page, specify in the Name field.
- 76. In the Description field, describe future references.
- 77. Click Next.



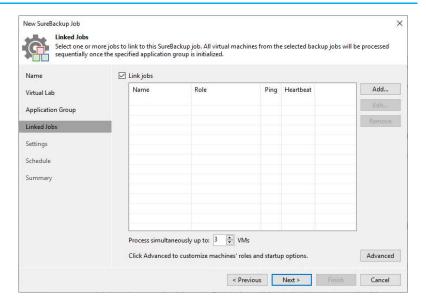
78. Select the Advanced
Single-Host 01 virtual lab
from the drop-down list
on the Virtual Lab page
and click Next.



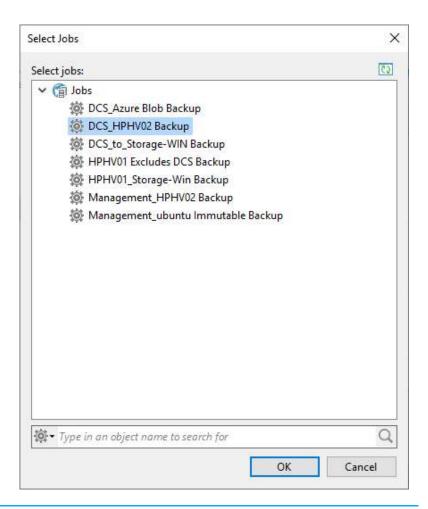
- 79. Select the application group from the drop-down list.
- 80. Select Keep the application group running after the job completes.
- 81. Click Next.



82. Select the Link jobs check box on the Linked jobs page and click Add.



83. On the Select page, select the job and click OK.

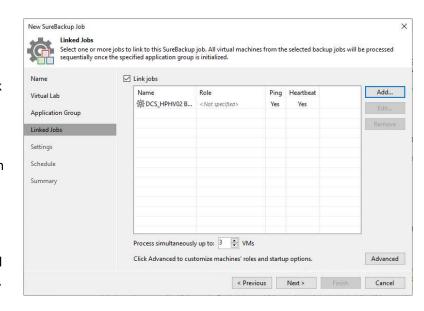


84. On the Link jobs page, change the number of processes simultaneously VMs if necessary and click Advanced.

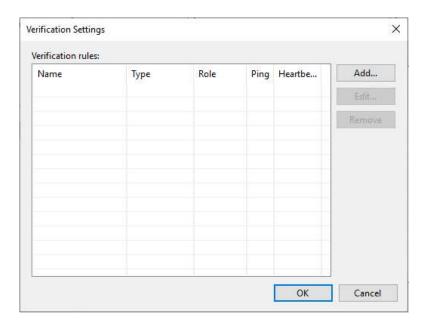
Note:

By default, you can launch and test up to three virtual machines simultaneously. You can also increase the number of VMs started and tested simultaneously. However, if these VMs are resource intensive, the performance of the SureBackup job and the Hyper-V host on which the virtual lab is hosted may suffer.

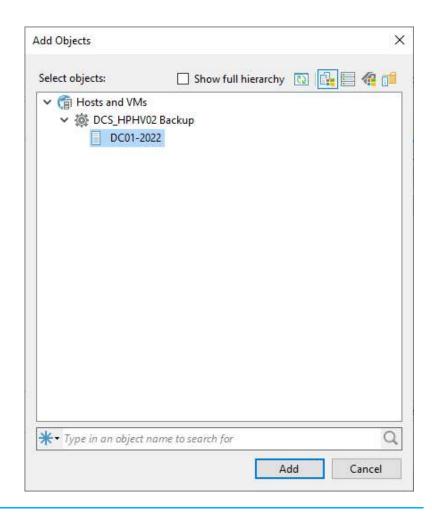
85. Click Advanced.



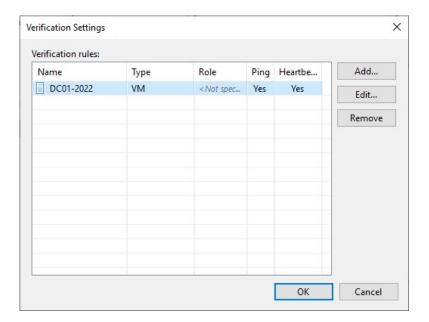
86. On the Verification Settings page, click Add.



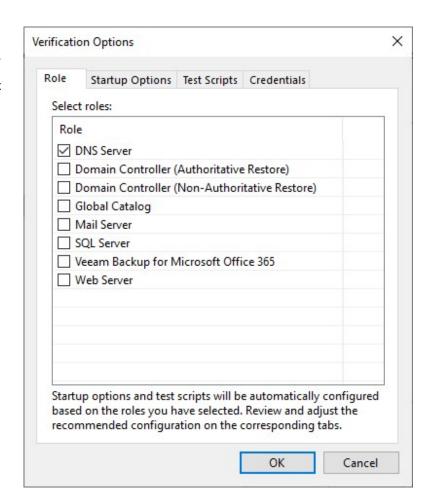
87. On the Add Objects, select the Machine and click Add.



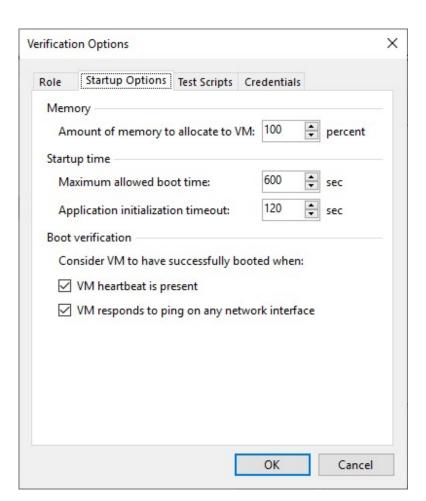
88. Select the machine and click Edit on the Verification Settings page.



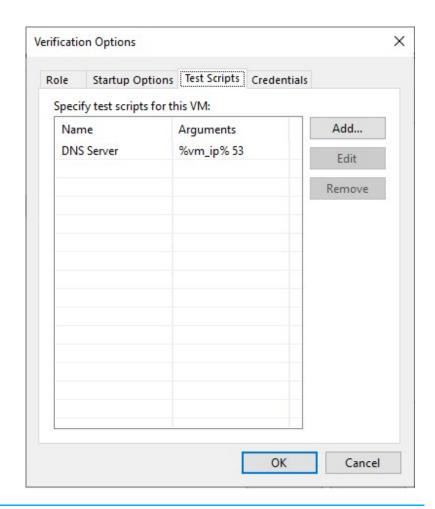
- 89. On the Verification Options page, select Role.
- 90. Select the roles check box in the Role.



- 91. On the Verification Options page, select Startup Options.
- 92. On the Start Options page, In the Memory section, specify how much memory you want to pre-allocate to the VM when the system boots.
- 93. In the Startup time section, enter the maximum boot time for the VM and the timeout for the VM to initialize applications.
- 94. In the Boot verification section, select the VM heartbeat is present checkbox.
- 95. Select VM responds to ping on any network interface.

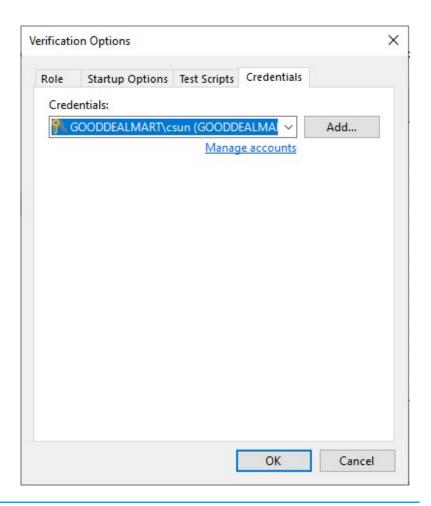


- 96. On the Verification
 Options page, select Test
 Scripts.
- 97. You can add or edit the script if needed.

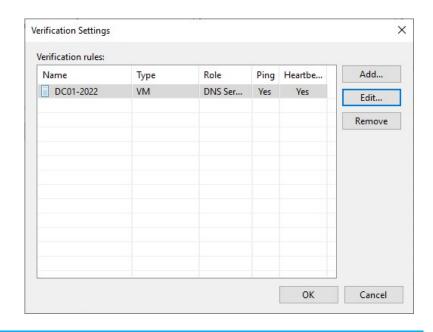


- 98. On the Verification Options page, select Credentials.
- 99. Select a user account from the Credentials drop-down list.

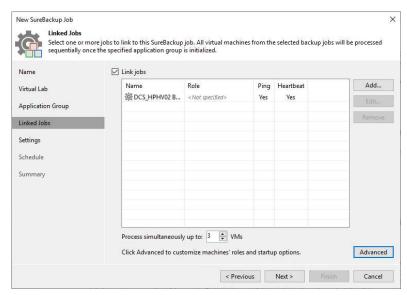
100. Click OK.



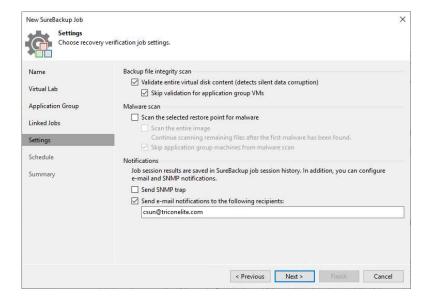
101. On the Verification Settings page, click OK.



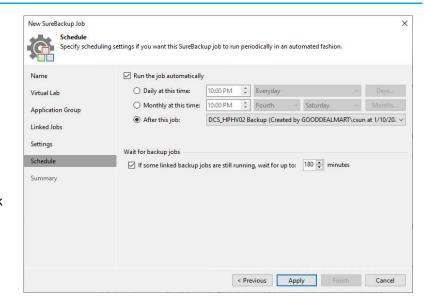
102. On the Linked Jobs page, click Next.



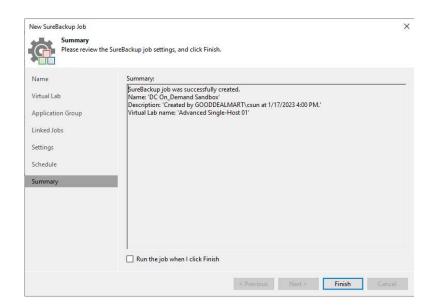
- 103. On the Settings page,
 Select the Validate entire
 virtual disc contents
 check box if you want to
 validate the backup file
 with a CRC check to
 ensure it is not corrupted.
- 104. Select the Skip validation for application group VMs checkbox if you want to exclude VMs being a part of the application group from this test.
- 105. Select the Scan the selected restore point for malware check box if you want Veeam Backup & Replication to scan VM data with antivirus software.
- 106. Select the Scan the entire image checkbox if you want the antivirus software to continue scanning VM data after the first malware is detected.
- 107. Select the Skip application group machines from the malware scan check box if you do not want to scan



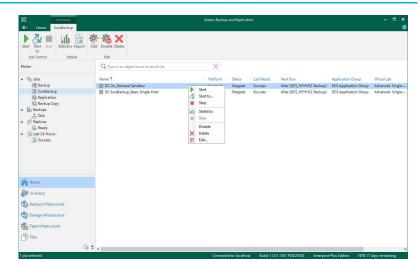
- VMs in the application group.
- 108. Select the Send SNMP trap check box to receive SNMP traps.
- 109. Select the Send email notifications to the following recipients checkbox if you want to receive notifications via email.
- 110. Enter the recipient's email address and click Next.
- 111. On the Schedule page, select Run the job automatically.
- 112. Select After this job and choose the preceding job from the list.
- 113. Select If some linked backup jobs are still running, wait to the check box and specify the necessary period in the field on the right. For example, suppose the linked job is still running.
- 114. Click Apply.



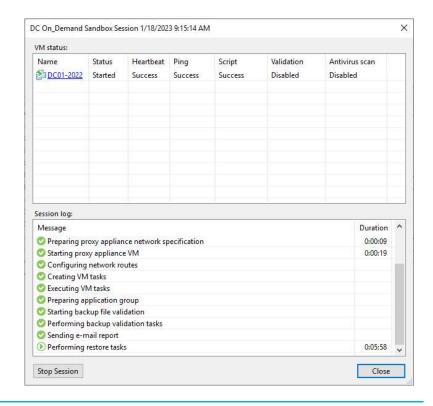
115. On the Summary page, click Finish.



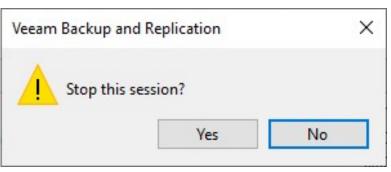
- 116. On the Home page, expand Jobs and select SureBackup.
- 117. Ensure the new SureBackup job is created.
- 118. Right-click the SureBackup job and select Start
- 119. Veeam Backup &
 Replication will launch the virtual lab and power on the VMs in the application group. Then, you can connect to virtual machines and run tests on them.



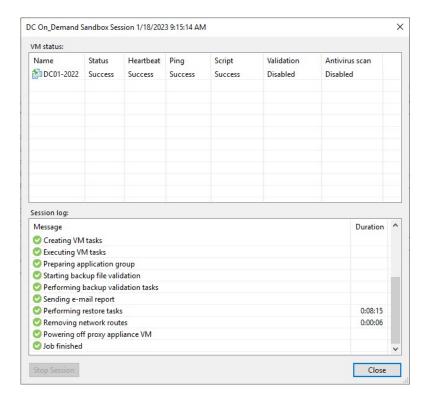
120. Click Stop Session after testing them completed.



121. Click Yes on the Stop this session warning display window.



122. Click Close on the session page after the job is finished.



Chapter 7

Reporting

Veeam Backup & Replication saves job statistics and operation data to the configuration database when you run a job. As a result, you can view real-time statistics for any job completed and generate reports with statistics data for any job or separate job session.

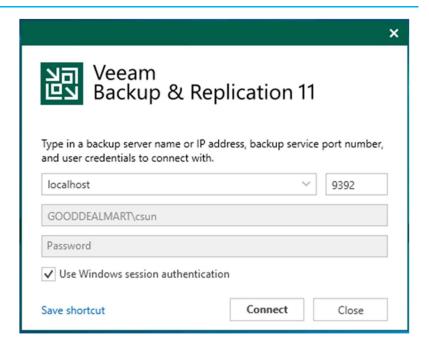
Data collected from Microsoft Hyper-V and Veeam Backup & Replication servers are analyzed in Veeam ONE reports. Reports provide structured data to assist you in examining historical data for the managed backup infrastructure and virtual environment. Reports can be viewed in a web browser, exported to various formats, or scheduled for automatic delivery via email, disc, or network share.

Real-Time Job Statistics Report from Veeam Backup & Replication Console

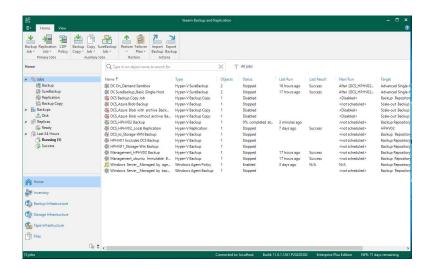
The real-time statistics provide detailed data on job sessions, such as job progress, duration, processing rate, performance bottlenecks, amount of processed data, read and transferred data, and session performance details, such as warnings and errors during the operation.

Instructions Screenshot (if applicable)

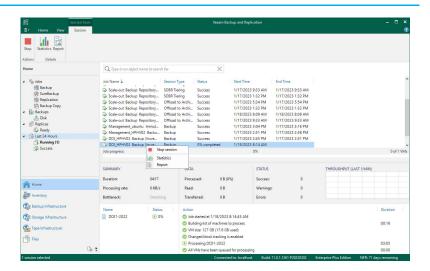
- Log in to the Veeam
 Backup and replication
 manager server.
- Open the Veeam Backup & Replication Console, and click Connect.



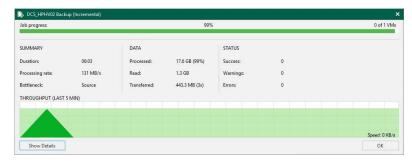
3. On the Home page, select the Last 24 Hours (or expand the Last 24 Hours and select) Running.

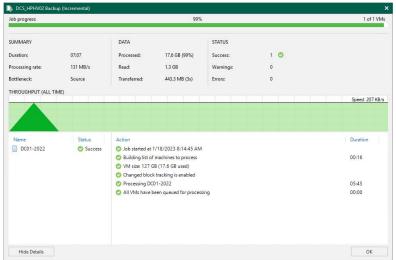


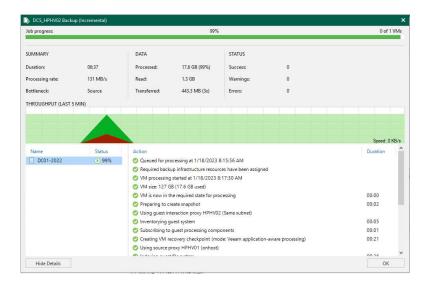
4. Right-click the job and select Statistics.



5. Real-time statistics provide detailed data on job sessions, such as job progress, duration, processing rate, performance bottlenecks, amount of processed data, read and transferred data, and session performance details, such as warnings and errors during the operation.







Job Session History Report from Backup and Replication Console

You can view detailed historical statistics on every job session.

Instructions Screenshot (if applicable) 1. Log in to the Veeam × Backup and replication manager server. Veeam Backup & Replication 11 2. Open the Veeam Backup & Replication Console, and click Connect. Type in a backup server name or IP address, backup service port number, and user credentials to connect with. localhost 9392 GOODDEALMART\csun Password

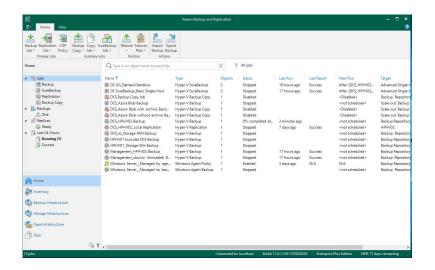
✓ Use Windows session authentication

Save shortcut

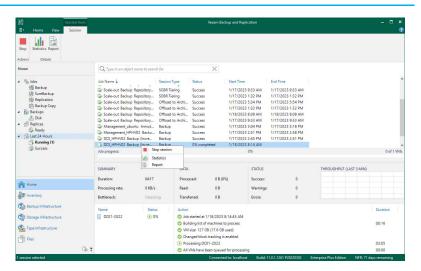
Connect

Close

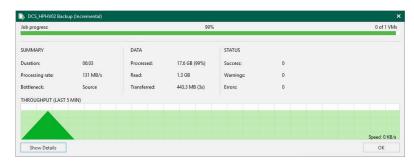
3. On the Home page, select the Last 24 Hours (or expand the Last 24 Hours and select) Running.

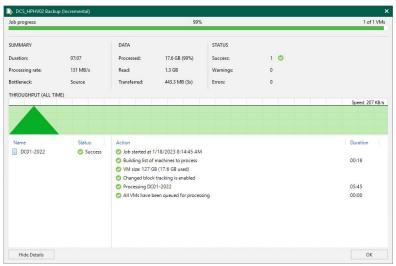


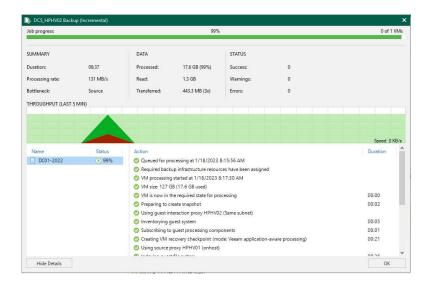
4. Right-click the job and select Statistics.



- 5. Real-time statistics provide detailed data on job sessions, such as job progress, duration, processing rate, performance bottlenecks, amount of processed data, read and transferred data, and session performance details, such as warnings and errors during the operation.
- 6. I use the keyboard's left and right arrow keys to switch between previous job sessions.







Job and Job Session Report from Veeam Backup and Replication Console

You can generate reports with information from all job sessions or just one.

Screenshot (if applicable) Instructions 1. Log in to the Veeam × Backup and replication manager server. Veeam Backup & Replication 11 2. Open the Veeam Backup & Replication Console, and click Connect. Type in a backup server name or IP address, backup service port number, and user credentials to connect with. localhost 9392 GOODDEALMART\csun Password

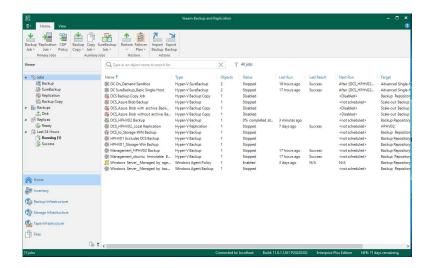
✓ Use Windows session authentication

Save shortcut

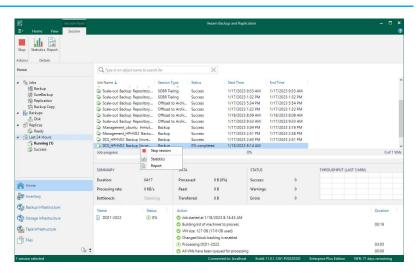
Connect

Close

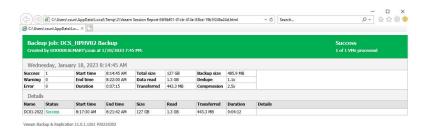
3. On the Home page, select the Last 24 Hours (or expand the Last 24 Hours and select) Running.



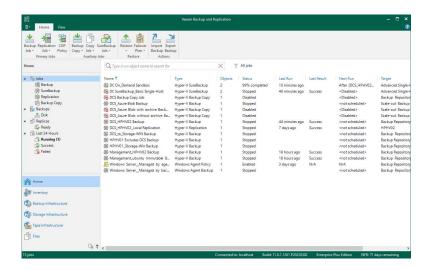
4. Right-click the job and select Report.



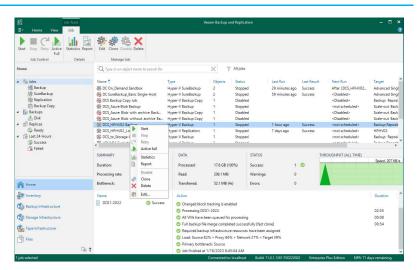
5. Review the single session report.



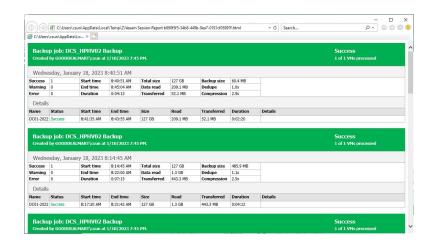
- 6. To create all sessions reported.
- 7. On the Home page, select Jobs.



8. Right-click the job and select Report.



9. Review all job sessions report.



Generate Backup Reports from Veeam ONE

You can generate reports with detailed backup information from Veeam ONE.

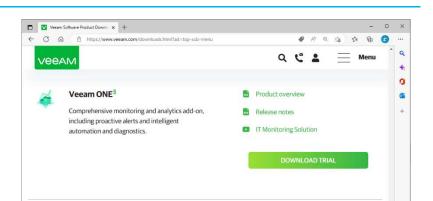
Instructions

Screenshot (if applicable)

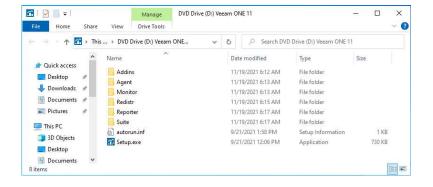
- Log in to the Veeam ONE Server.
- Sign in to your Veeam account and download the Veeam ONE software.

Note:

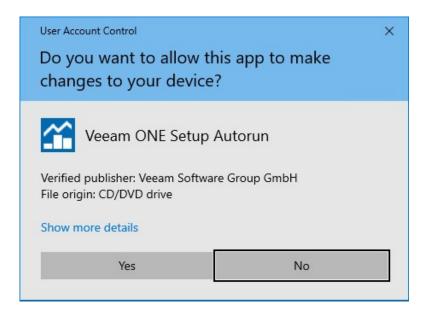
Included with the purchase of Veeam Availability Suite and Veeam Backup Essentials and is included in the download package.



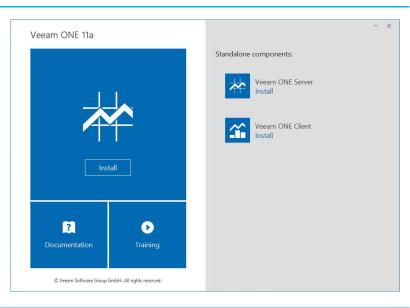
- 3. Mount the Veeam ONE ISO image file.
- 4. Run Setup.exe.



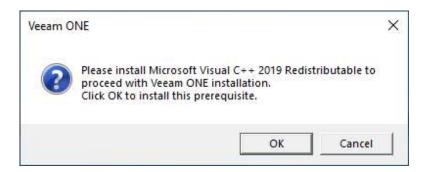
5. On the User Account Control page, click Yes.



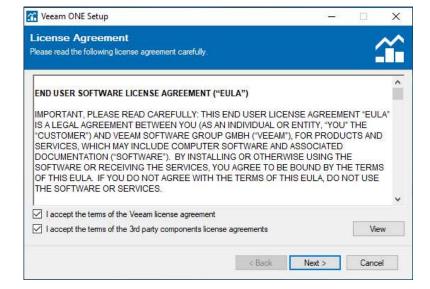
On the Veeam ONE 11a page, click Veeam ONE Server Install.



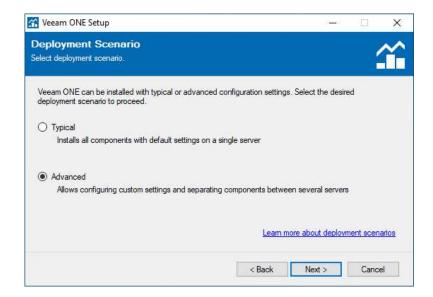
 Click OK on the Please install Microsoft Virtual C++ 2019 Redistributable to proceed with the Veeam ONE installation page



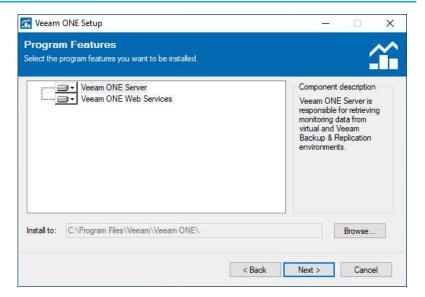
- 8. Select I accept the Veeam license agreement checkbox terms on the License Agreement page.
- Select I accept the terms of the 3rd party components license agreements checkboxes
- 10. Click Next.



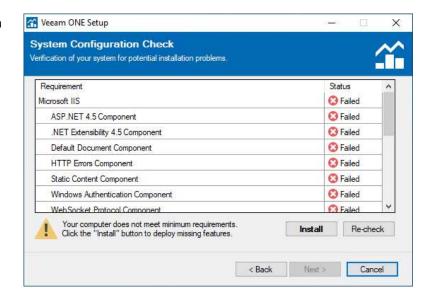
On the Deployment
 Scenario page, select
 Advanced and click Next.



- 12. On the Program Features page, Choose the features you want to install.
- 13. Click Browse and select the path.
- 14. Click Next.

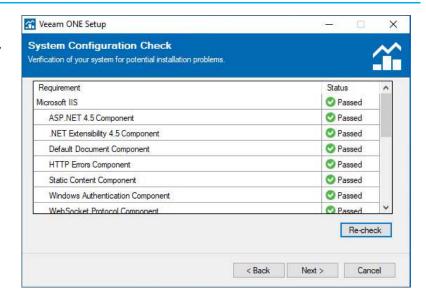


 Click Install on the System Configuration Check page and install the missing software components.

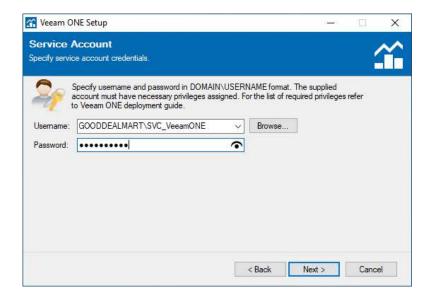


16. On the System

Configuration Check page,
ensure the requirements
installed are completed
successfully and click
Next.



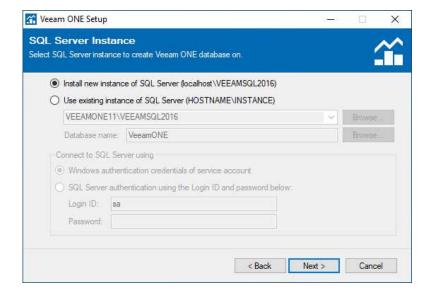
- 17. On the Service Account page, click Browse and select the user account as a service account.
- 18. Enter the password and click Next.



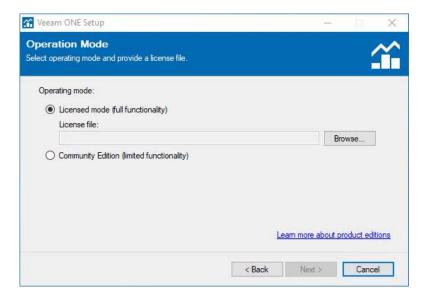
On the SQL Server
 Instance page, select
 Install a new instance of
 SQL Server
 (localhost\VEEAMSQL201
 6) and click Next.

Note:

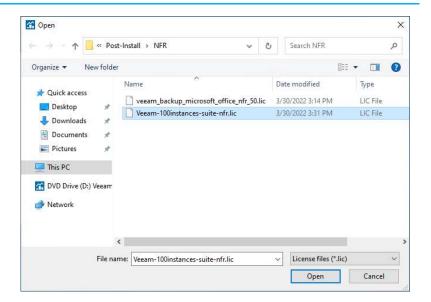
If you install Veeam ONE Server and Veeam ONE Web Services on separate machines, ensure that both components share a database.



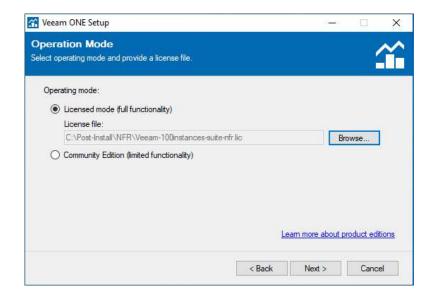
20. Select Licensed mode (full function) on the Operation Mode page and click Browse.



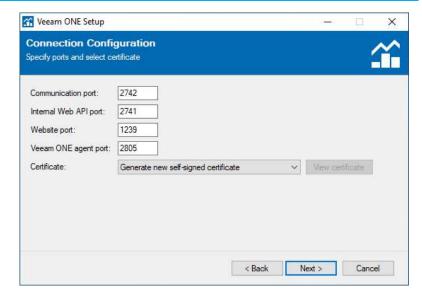
21. Select the license file and click Open.



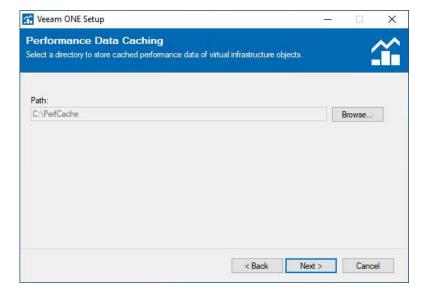
22. Click Next on the Operation Mode page.



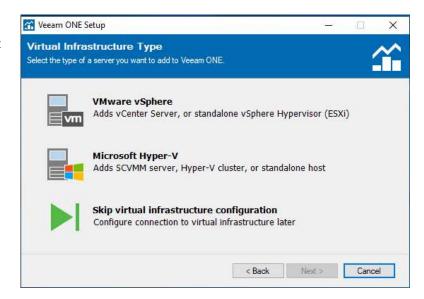
- 23. On the Connection
 Configuration page, enter
 the port number used to
 communicate with the
 Veeam ONE Reporting
 service in the
 Communication port field.
- 24. Enter the number of ports the Veeam ONE Server component will use to communicate with the Veeam ONE Web API component in the Internal Web API port field.
- 25. Enter the Veeam ONE
 Agent's port number to
 collect data from
 connected Veeam Backup
 & Replication servers in



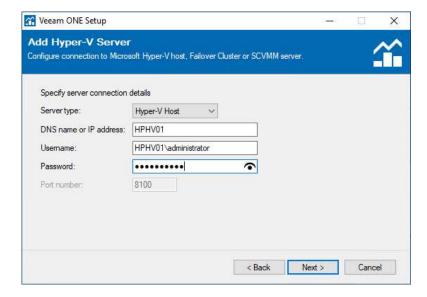
- the Veeam ONE agent port field.
- 26. Select a certificate from the Certificate list to secure the Veeam ONE internal Web API connection. The certificate can be changed later in the Veeam ONE Settings utility.
- 27. Click Next.
- 28. The performance cache is stored in the C:\PerfCache folder by default on the Performance Data Caching page. To save the cache to a different location, click Browse next to the Path field and enter the path to the new location.
- 29. Click Next.



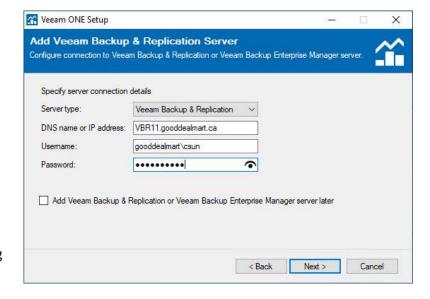
30. On the Virtual Infrastructure Type, select Microsoft Hyper-V.



- 31. Select the Server type from the drop-down list on the Add Hyper-V Server page.
- 32. Enter the FQDN or IP address of the virtualization server to connect in the DNS name or IP address field.
- 33. Enter the account credentials for connecting to the server in the Username and Password fields.
- 34. Click Next.

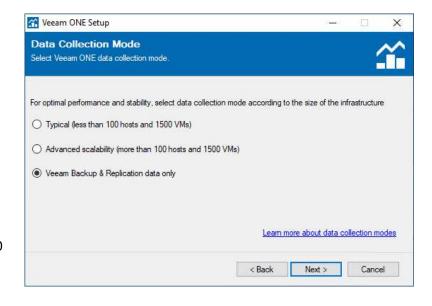


- 35. Select the Server type from the drop-down list on the Add Veeam Backup & Replication Server page.
- 36. Enter the FQDN or IP address of the backup manager server to connect in the DNS name or IP address field.
- 37. Enter the account credentials for connecting to the server in the Username and Password fields.
- 38. Select Add Veeam Backup & Replication or Veeam Backup Enterprise Manager server later if you do not want to configure backup server connection settings during installation.
- 39. Click Next.

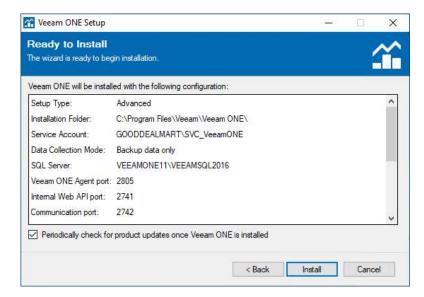


- 40. On the Data Collection Mode, select Typical (less than 100 hosts and 1500 VMs). Veeam ONE collects all inventory, configuration, and performance metrics and displays them in dashboards, reports, and alarms.
- 41. Select Advanced scalability (more than 100 hosts and 1500 VMs).

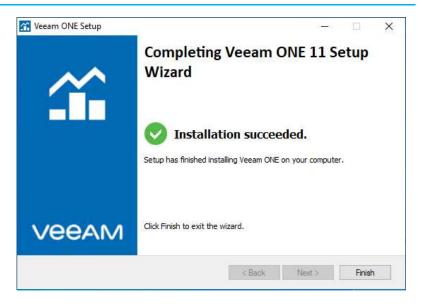
 Veeam ONE collects all metrics required for alarms and reports.
- 42. Select Veeam Backup & Replication Data Only. The Veeam ONE collects all inventory, configuration and performance metrics from Veeam Backup & Replication servers. It also collects inventory and configuration metrics from virtualization servers but skips virtual infrastructure performance metrics.
- 43. Click Next.



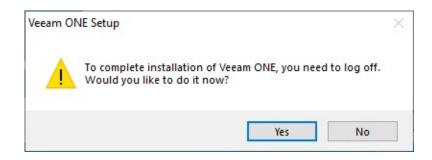
44. Select the Periodically check for product updates once Veeam ONE is installed check box on the Ready to Install page and click Next.



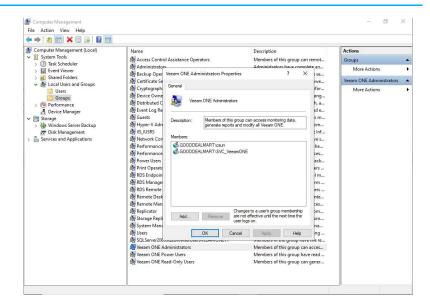
45. Click Finish on theCompleting Veeam ONE11 Setup Wizard page.



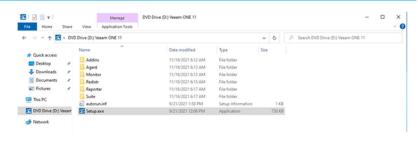
46. Click Yes on the logoff warning display window.



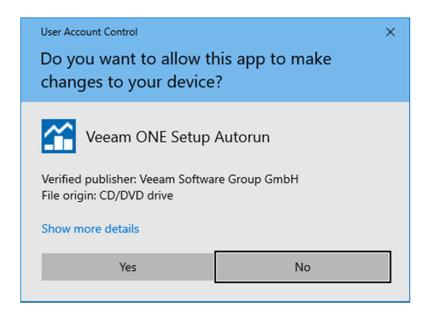
- 47. Log in to the Veeam ONE server.
- 48. Double-click the Veeam
 ONE Administrators
 group, and ensure the
 Veeam ONE service
 account is included in the
 Veeam ONE
 Administrators user
 group.



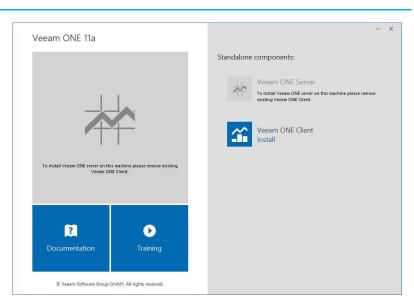
- 49. Log in to the Veeam ONE Client machine.
- 50. Mount the Veeam ONE ISO image file.
- 51. Run Setup.exe.



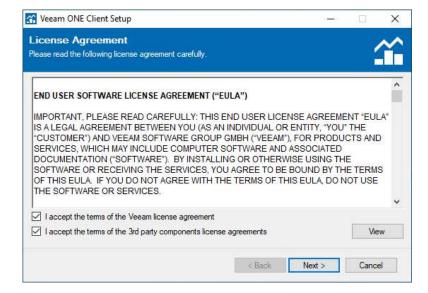
52. On the User Account Control page, click Yes.



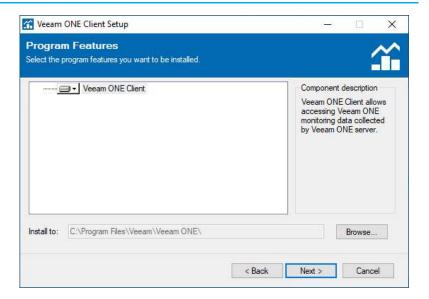
53. On the Veeam ONE 11a page, click Veeam ONE Client Install.



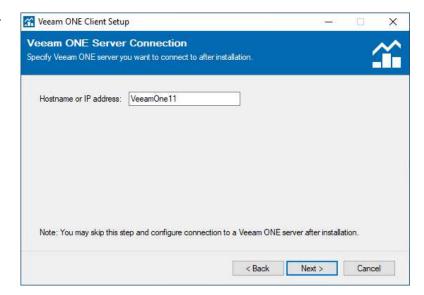
- 54. Select I accept the Veeam license agreement checkbox terms on the License Agreement page.
- 55. Select I accept the terms of the 3rd party components license agreements checkbox.
- 56. Click Next.



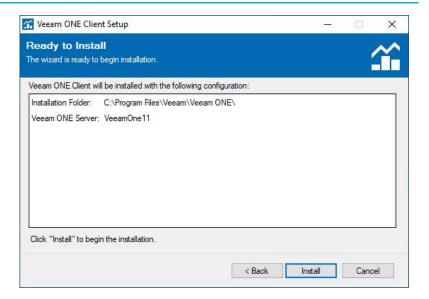
57. On the Program Features page, click Next.



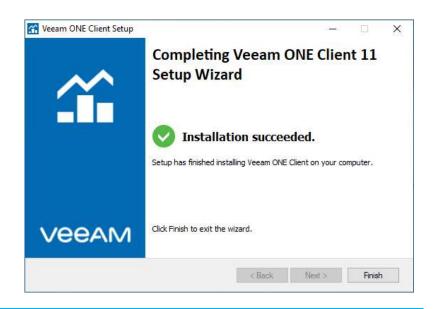
- 58. On the Veeam ONE Server Connection page, enter the Veeam ONE server's Hostname or IP address.
- 59. Click Next.



60. Click Install on the Ready to Install page.



61. Click Finish on the Completing Veeam ONE Client 11 Setup Wizard page.



- 62. Open Veeam ONE Client.
- 63. On the Sign in to Veeam ONE page, click Connect.



64. On the SMTP Server page, enter the SMTP server name in the SMTP server field.

Notification Settings

Email Notifications

Notification Policy

SNMP Settings

Summary

SMTP server:

smtp.sendgrid.net

VeeamONE11@gooddealmart.ca

✓ Use secure connection

apikey

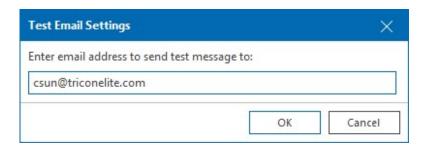
✓ Use authentication

Send Test Email

Skip email notifications configuration

SMTP Server

- 65. Enter the port number in the Port field.
- 66. Enter the sender's email address in the From field.
- 67. Select Use secure connection.
- 68. Select Use authentication.
- 69. Enter a username in the login field.
- 70. Enter a password in the Password field.
- 71. Click Send Test Email.
- 72. Enter a recipient's email address in the Test Email Settings page and click OK.



Configure SMTP server settings to receive alarm notifications, scheduled reports and dashboards

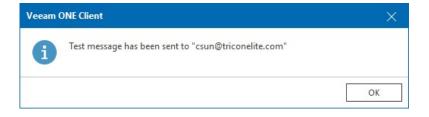
Password:

Previous Next

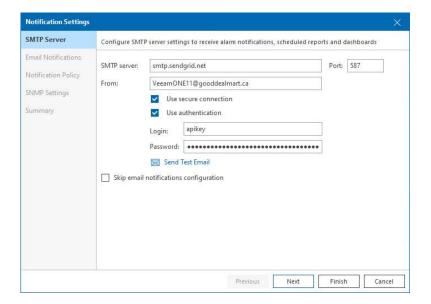
Port: 587

Finish Cancel

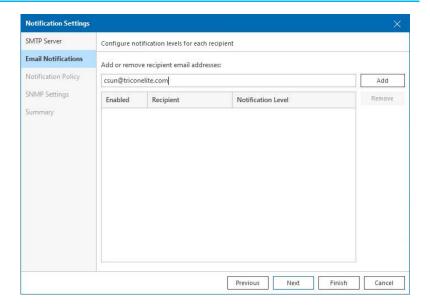
73. Ensure the test message is sent without issues and click OK.



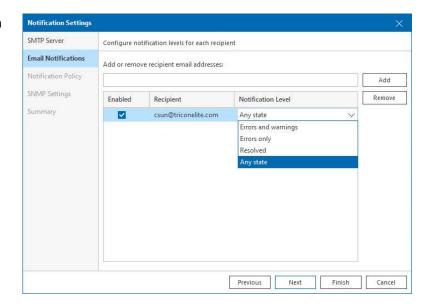
74. On the SMTP Server page, click Next.



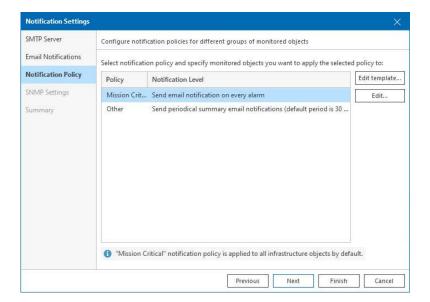
75. Enter a recipient email address on the Email Notifications page and click Add.



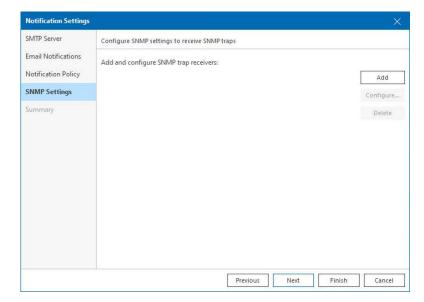
76. Select the state level from the Notification Level drop-down list and click Next.



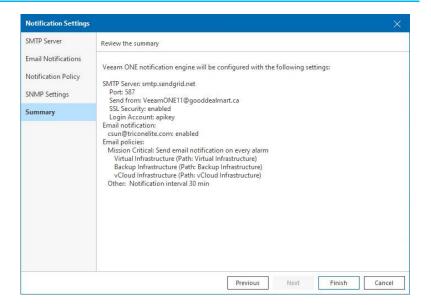
- 77. Click Edit template or edit if necessary on the Notification Policy page.
- 78. Click Next.



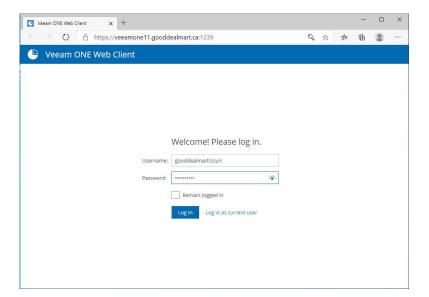
- 79. On the SNMP Settings page, click Add to add and configure SNMP trap receivers if necessary.
- 80. Click Next.



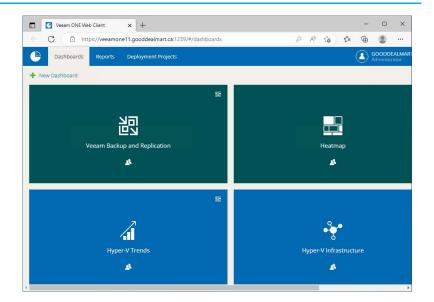
81. On the Summary page, click Finish.



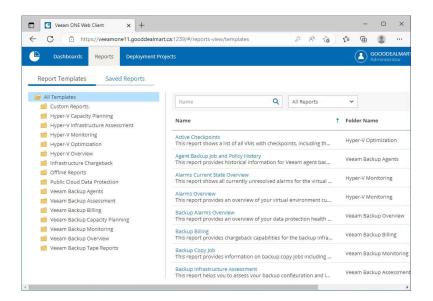
- 82. Open the Veeam ONE Web Client.
- 83. Enter your username and password on the Veeam ONE Web Client page and click Log in.



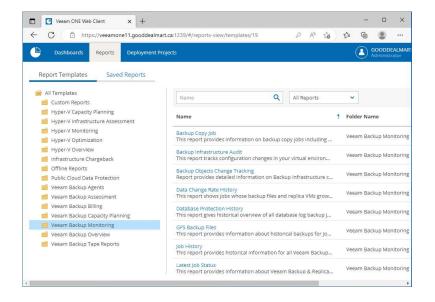
84. On the Veeam ONE Web Client page, select Reports.



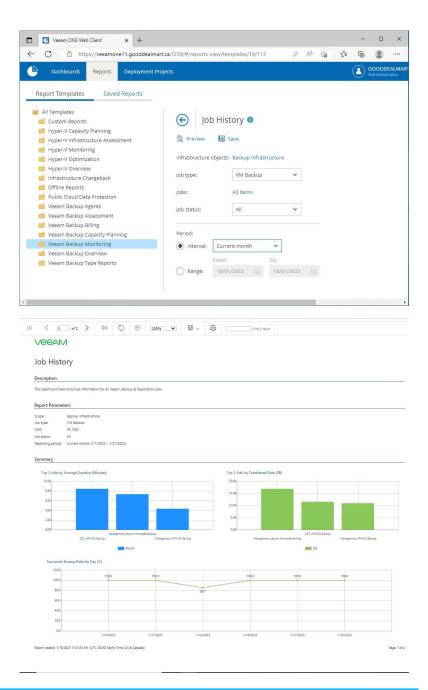
85. On the Report Templates page, select Veeam Backup Monitoring.

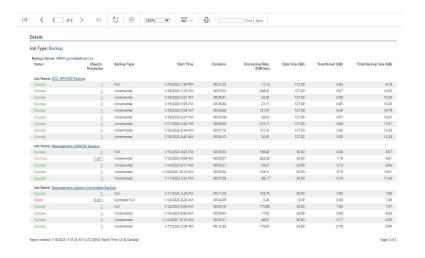


86. On the Veeam Backup Monitor page, select Job History.

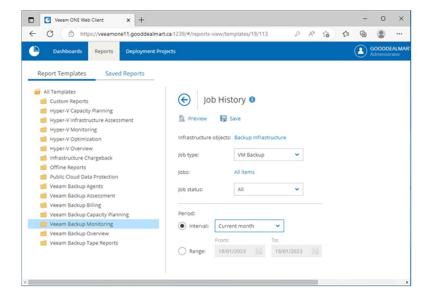


- 87. Select VM Backup from the job type drop-down list on the Job History page.
- 88. Select All Items in the Jobs field.
- 89. Select All from the Job Status drop-down list.
- 90. Select the period from the Interval drop-down list.
- 91. You also can select the period rage from the Rang.
- 92. Click Preview to view the report.

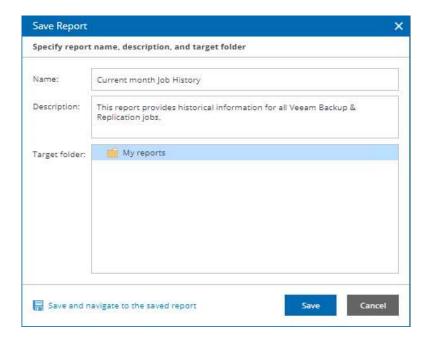




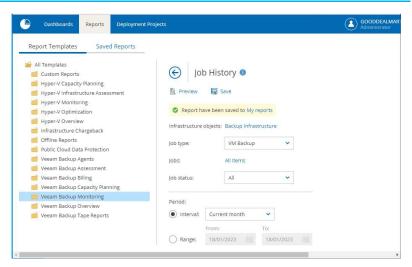
93. On the Job History page, click Save.



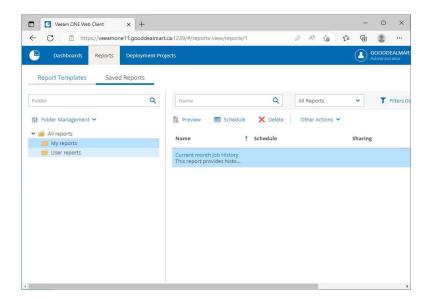
- 94. Enter the report name in the Name field on the Save Report page.
- 95. Describe the comments in the Description field.
- 96. Click Save.



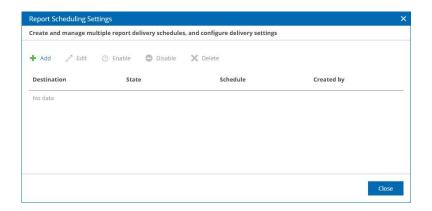
97. On the Job History page, click My reports.



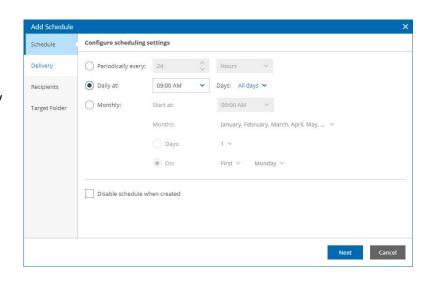
98. Select the saved report and click Schedule on the Saved Reports page.



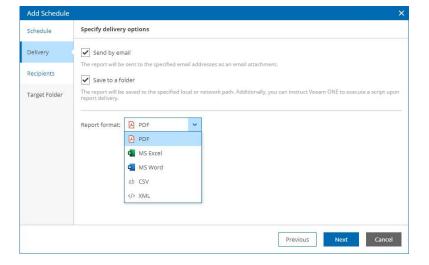
99. On the Report Schedule Settings page, click Add.



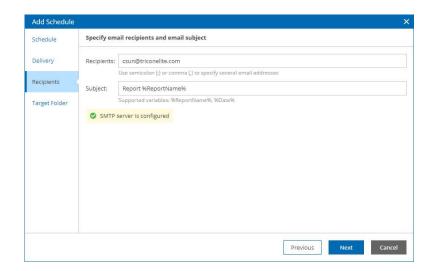
- 100. On the Configure scheduling settings page, there are three options.
- 101. Select Periodically every option and define the necessary interval.
- 102. Select the Daily option and enter the time and days of the week the object must be delivered.
- 103. Select the Monthly option and choose the necessary months, dates or weekdays.
- 104. Select the Disable schedule when created check box to configure a schedule without enabling it.



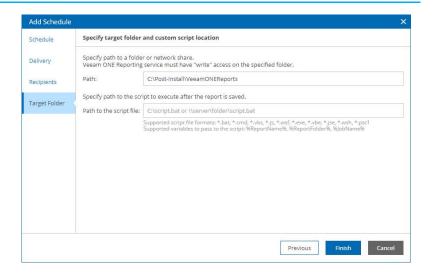
- 105. On the Specify delivery options page, select the Send by an email check box to send objects by email.
- 106. Select the Save to a folder checkbox.
- 107. Select the report file format from the Report format drop-down list.
- 108. Click Next.



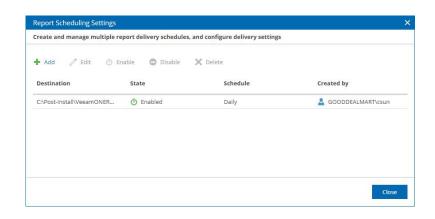
109. On the Recipients page, enter the recipient's address and click Next.



110. Enter the folder path in the Path field on the Target Folder page and click Finish.



- 111. Enable the schedule state on the Report Scheduling Settings page and click Close.
- 112. The report file will save to the path and email recipients.



Chapter 8

Join us at MVPDays and meet great MVPs like this in person

If you liked their book, you would love to hear them in person.

Live Presentations

Dave frequently speaks at Microsoft conferences around North America, such as TechEd, VeeamOn, TechDays, and MVPDays Community Roadshow.

Cristal runs the MVPDays Community Roadshow.

You can find additional information on the following blog:

www.checkyourlogs.net

www.mvpdays.com

Video Training

For video-based training, see the following site:

www.mvpdays.com

Live Instructor-led Classes

Dave has been a Microsoft Certified Trainer (MCT) for over 15 years and presents scheduled instructor-led classes in the US and Canada. For current dates and locations, see the following sites:

418

- www.truesec.com
- www.checkyourlogs.net

Consulting Services

Dave and Cristal have worked with some of the largest companies in the world and have a wealth of experience and expertise. Customer engagements are typically between two weeks and six months.