

# Master Veeam Tricks

Volume 1

By Veeam Vanguard:

Dave Kawula   Cristal Kawula   Ian Sanderson

Karl Widmer   Rhys Hammond   Markus Kraus

Mike Conjoice   Craig Dalrymple   Didier Van Hoya

Eugene Kashperovetskyi

Foreword by: Rick Vanover aka RICKATRON



PUBLISHED BY

MVPDays Publishing

<http://www.mvpdays.com>

Copyright © 2018 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 prior written permission of the publisher.

ISBN: TBD

### **Warning and Disclaimer**

Every effort has been made to make this manual as complete and as 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](http://www.checkyourlogs.net) or sending an email to [feedback@mvpdays.com](mailto:feedback@mvpdays.com).



# Foreword by: Rick Vanover

# Acknowledgments

## From Dave

Cristal, you are my rock and my source of inspiration. For the past 20 + years you have been there with me every step of the way. Not only are you the “BEST Wife” in the world you are my partner in crime. Christian, Trinity, Keira, Serena, Mickaila, Mackenzie, and Rycker you kids are so patient with your dear old dad when he locks himself away in the office for yet another book. Taking the time to watch you grow in life, sports, and become little leaders of this new world is incredible to watch.

Thank you, Mom and Dad,, (Frank and Audry) and my brother Joe. You got me started in this crazy IT world when I was so young. Brother, you mentored me along the way both coaching me in hockey and helping me learn what you knew about PCs and Servers. I’ll never forget us as teenage kids working the IT Support contract for the local municipal government. Remember dad had to drive us to site because you weren’t old enough to drive ourselves yet. A great career starts with the support of your family, and I’m so lucky because I have all the support one could ever want.

Last but not least, the MVPDays volunteers, you have donated your time and expertise and helped us run the event in over 20 cities across North America. Our latest journey has us expanding the conference worldwide as a virtual conference. For those of you that will read this book, your potential is limitless just expand your horizons, and you never know where life will take you.

# About the Authors

## Dave Kawula – MVP / Veeam Vanguard

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 Center, 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 Center and operating system topics.

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

Recently Dave has been honored to take on the role of Conference Co-Chair of TechMentor with fellow MVP Sami Laiho. The lineup of speakers and attendees that have been to this conference over the past 20 years is fantastic. Come down to Redmond or Orlando in 2018, and you can meet him in person. Checkout his speaking site at [www.davekawula.com](http://www.davekawula.com)

He recently tied for 1<sup>st</sup> 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 both local customers and large international enterprises, providing optimal guidance and methodologies to achieve and maintain an efficient infrastructure.

BLOG: [www.checkyourlogs.net](http://www.checkyourlogs.net)

Twitter: @DaveKawula



## Cristal Kawula – MVP / Veeam Vanguard

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

Cristal can be found speaking at Microsoft Ignite, MVPDays, and other local user groups. She is extremely active in the community and has recently helped publish a book for other Women MVP's called Voices from the Data Platform.

This year at Microsoft Ignite she lead community meetups for various topics such as: Women in IT, Parenting in IT, Diversity in Tech, and becoming a Community Rockstar.

BLOG: <http://www.checkyourlogs.net>

Twitter: @supercristal1



## Ian Sanderson – VMWare vEXPERT / Veeam Vanguard

I am an IT professional with over 11 years IT experience with a focus on virtualization software and surrounding ecosystem products. This blog is a place for me to jot down any interesting things I find along the way. If you would like to get in touch, please use the social links at the top of this blog post.

BLOG: <https://www.snurf.co.uk>

Twitter: @ian0x0r



## Karl Widmer – VMWare vEXPERT / Veeam Vanguard

ICT system specialist in server and virtualization, responsible for all vSphere, HCI and cloud infrastructure at Oerlikon Group. Veeam and VMware addicted, hot on certifications. Veeam Vanguard 2017-2018, vExpert 2017-2018, VCA, MCSA, VMCE. Blogger.

BLOG: <https://www.driftar.ch>

Twitter: @widmerkarl



## Rhys Hammond – VMWare vEXPERT / Veeam Vanguard

My name is Rhys Hammond, and this blog is a collection of tips and tricks for solving issues that I've encountered designing, installing and supporting Veeam solutions for various clients.

I learned early on that what motivates me is working with technologies and products that I truly believe in. I am truly passionate about helping organizations make the most of the best backup & disaster recovery product available today.

Currently, I am working at Data#3 as a Senior Systems Engineer specializing in Veeam solutions. One of my goals is to help the IT community understand more of what Veeam can offer, as part of that effort, I maintain <https://rhyshammond.com>, a Veeam and Virtualization-centric weblog providing information, insight, and technical knowledge on all things Veeam related.

In 2017, I was awarded VMware vExpert for having demonstrated significant contributions to the community and a willingness to share expertise with others. VMware grants the vExpert as an honorary title for advocates of the company's products.

BLOG: <https://rhyshammond.com>

Twitter: @HammondRhys



## Markus Kraus – VMWare vEXPERT / Veeam Vanguard

Markus is living the IT Pro life since more than 10 years as a consultant, Architect and Systems Engineer. He is focusing on virtualization and automation technologies. In addition to his second Vanguard nomination, he was 2016-2018 awarded as a vExpert.

BLOG: <https://mycloudrevolution.com/en/>

Twitter: @vMarkus\_K



## Mike Conjoice – VMWare vEXPERT / Veeam Vanguard

Mike works as a Solution Architect at Bupa Dental UK, one of the UK's largest dental chains. Having worked with Veeam since 2012, initially as a partner, now as a customer, he is a strong advocate for Veeam as both a product and a company.

BLOG: <http://mikeconjoice.com/>

Twitter: @mikeconjoice



## Craig Dalrymple – VMWare vEXPERT / Veeam Vanguard

Senior Systems Engineer at brightsolid, Dundee, Scotland and co-leader/founder of the UK Veeam User Group. 20+ years experience within the IT industry. Enjoy rugby/gaming in spare time.

BLOG: <https://www.cragdoo.co.uk/>

Twitter: @cragdoo



## Didier Van Hoya – MVP / Veeam Vanguard

Didier Van Hoya is an IT veteran with over 20 years of experience and accumulated hands on expertise in ICT. He specializes mainly in Microsoft technologies (storage, virtualization, networking, cloud) to design and build highly available, high performance solutions that scale both efficiently and effectively without breaking the budget. He works mainly as a subject matter expert advisor and infrastructure architect in Wintel & Hybrid Azure environments, often leveraging DELL EMC hardware to deliver exceptional value for the money. As a Microsoft MVP in Cloud and Datacenter, a member of the Microsoft Extended Experts Team in Belgium, a DELL TechCenter Rockstar and a Veeam Vanguard he contributes his experience and knowledge to the global community. He's a trusted adviser, blogger, writer and public speaker on his areas of expertise.

Blog: <https://blog.workinghardinit.work>

Twitter: @workinghardinit

LinkedIn: <https://be.linkedin.com/in/didiervanhoye>



## Eugene Kashperovetskyi – MVP / Veeam Vanguard

Eugene Kashperovetskyi, affectionately known as Eugene K, is a highly experienced web hosting professional with 15 years of hosting solutions support, design, implementation and automation. Whether performing minor system optimizations or handling big projects with different unknowns involved, Eugene believes there is always a way to make it run smoothly, seamlessly and at a high rate of automation applied for routinely performed operations. He likes to be challenged with the tasks and research to approach the goal desired.

BLOG: <https://www.singlehop.com/blog/author/eugenek/>



## Technical Editors

### Emile Cabot - MVP

Emile started in the industry during the mid-90s working at an ISP and designing celebrity web sites. He has a strong operational background specializing in Systems Management and collaboration solutions and 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



## **Cary Sun – Microsoft MVP / Cisco Champion**

Cary Sun is 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 multi-platform, all LAN/WAN topologies, network administration, E-mail and Internet systems, security products, PCs and Servers environment. Expertise is analyzing user's 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's is a very active blogger at [checkyourlogs.net](http://www.checkyourlogs.net) and always available online for questions from the community. He passion for technology is contagious, and he makes everyone around him better at what they do.

Blog:<http://www.checkyourlogs.net>

Twitter: @SifuSun



# Contents

<b>Foreword by: Rick Vanover .....</b>	<b>iii</b>
<b>Acknowledgments.....</b>	<b>iv</b>
From Dave .....	iv
<b>About the Authors .....</b>	<b>v</b>
Dave Kawula – MVP / Veeam Vanguard.....	v
Cristal Kawula – MVP / Veeam Vanguard.....	vi
.....	vi
Ian Sanderson – VMWare vEXPERT / Veeam Vanguard .....	vii
Karl Widmer – VMWare vEXPERT / Veeam Vanguard .....	viii
Rhys Hammond – VMWare vEXPERT / Veeam Vanguard .....	ix
Markus Kraus – VMWare vEXPERT / Veeam Vanguard.....	x
.....	x
Mike Conjoice – VMWare vEXPERT / Veeam Vanguard .....	xi
.....	xi
Craig Dalrymple – VMWare vEXPERT / Veeam Vanguard .....	xii
Didier Van Hoya – MVP / Veeam Vanguard.....	xiii
Eugene Kashperovetskyi – MVP / Veeam Vanguard .....	xiv
Technical Editors .....	xv
Emile Cabot - MVP.....	xv
Cary Sun – Microsoft MVP / Cisco Champion .....	xvi
	xvii

<b>Contents.....</b>	<b>xvii</b>
<b>Introduction .....</b>	<b>23</b>
<b>North American MVPDays Community Roadshow .....</b>	<b>23</b>
Sample Files .....	24
Additional Resources .....	24
<b>Chapter 1.....</b>	<b>26</b>
<b>Upgrading a Veeam Backup Repository from Server 2016 to Server 2019 + ReFS DeDuplication .....</b>	<b>26</b>
Upgrading to Windows Server 2019 .....	27
Install and Configure Windows Deduplication .....	37
Tweak Windows Deduplication.....	43
<b>Chapter 2.....</b>	<b>49</b>
<b>Veeam RESTful API and Postman.....</b>	<b>49</b>
Postman .....	49
Veeam Enterprise Manager.....	51
RESTful API Authentication.....	52
Using Postman GET to obtain login sessions .....	52
Using Postman POST to obtain login session authentication token.....	55
Adding a header.....	59
Example of making API call and information .....	61
Summary.....	61
<b>Chapter 3.....</b>	<b>62</b>
<b>Sizing Veeam for Office 365 Backup Storage.....</b>	<b>62</b>
Background .....	62

Getting Started .....	62
Calculation .....	67
<b>Chapter 4.....</b>	<b>68</b>
<b>HTML Report of all Veeam Backup Copy Jobs in PowerShell .....</b>	<b>68</b>
<b>Chapter 5.....</b>	<b>73</b>
<b>Troubleshooting Veeam B&amp;R Error Code: 32768 Failed to create VM Recovery Snapshot.....</b>	<b>73</b>
<b>Chapter 6.....</b>	<b>80</b>
<b>Install Veeam Backup and Replication using PowerShell .....</b>	<b>80</b>
<b>Chapter 7.....</b>	<b>97</b>
<b>Veeam for Nutanix AHV.....</b>	<b>97</b>
Background.....	97
Before you begin .....	97
Appliance Setup .....	97
Veeam and Nutanix Integration .....	103
Creating a Backup Job .....	109
<b>Chapter 8.....</b>	<b>118</b>
<b>Veeam Backup and Replication and SMB MultiChannel .....</b>	<b>118</b>
Introduction .....	118
Proving Veeam Backup and Replication leverages SMB Multichannel .....	118
Conclusion .....	121
SMB Direct Support.....	122
More SMB3 and Windows Server 2016 Goodness.....	122

Items to investigate further .....	122
<b>Chapter 9.....</b>	<b>124</b>
<b>Veeam Backup and Replication Preferred Subnet and SMB Multichannel .....</b>	<b>124</b>
Introduction .....	124
Veeam Backup and Replication Preferred Subnet & SMB Multichannel .....	125
So how do we control the NICs used with SMB Multichannel? .....	133
Some additional things to note .....	133
Conclusion .....	136
<b>Chapter 10.....</b>	<b>137</b>
<b>Using StarWind Cloud VTL for AWS and Veeam.....</b>	<b>137</b>
How does the Solution Work? .....	138
Download the Software .....	139
Installation of StarWind VTL .....	140
The configuration of StarWind VTL.....	147
Creating an Amazon S3 Bucket .....	157
Setup Cloud Replication .....	161
Configure Microsoft iSCSI Initiator .....	164
Installing HPE StoreEver Tape Drivers.....	171
Add the Virtual Tape Library to Veeam.....	176
Configure a Tape backup Job .....	179
Conclusion .....	195
<b>Chapter 11.....</b>	<b>196</b>
<b>Configure Azure Direct Restore with Veeam 9.5 .....</b>	<b>196</b>
Setup your Azure Proxy VM .....	197

Restore a VM to Azure .....	212
<b>Chapter 12.....</b>	<b>226</b>
<b>Configure 2 Factor Authentication (2FA) for Veeam with DUO .....</b>	<b>226</b>
Important Notes.....	228
Creating a DUO Account.....	228
Configuring DUO .....	235
Installing the DUO Client .....	237
Adding DUO Users.....	241
<b>Chapter 13.....</b>	<b>247</b>
<b>The Secret Life of VBM Files.....</b>	<b>247</b>
What do VMB files do? .....	247
Why do we create VMBs? .....	247
When are VMB files Created? .....	247
Where are VBM files stored?.....	248
What is inside of a VBM file? .....	248
When do we use VBM files?.....	249
What happens if I don't have a VBM file? .....	249
<b>Chapter 14.....</b>	<b>250</b>
<b>Efficient Veeam NetApp Backups from Storage Snapshots.....</b>	<b>250</b>
Storage Design for NetApp Backup from Storage Snapshots.....	251
Storage Snapshot procedure.....	252
Storage Snapshot Configuration.....	253
Veeam Setup .....	254
VSphere Infrastructure .....	263

<b>Chapter 15.....</b>	<b>264</b>
<b>Sizing your Backup Storage and Saving Money .....</b>	<b>264</b>
Step One – Assess the Job Setup and Key Requirements .....	265
Step Two – Make sure the backup Windows Align with RPO/RTO.....	266
Step Three – Make Sure the network can handle the Backup Jobs.....	267
Step Four – Calculate Total Backup Storage Requirements.....	267
Step Five – Map your Provisioning Plan .....	271
Step Six – Plan How to Spend your Savings .....	271
<b>Chapter 16.....</b>	<b>273</b>
<b>Join us at MVPDays and meet great MVP’s like this in person.....</b>	<b>273</b>
Live Presentations .....	273
Video Training.....	273
Live Instructor-led Classes .....	274
Consulting Services .....	274

## Introduction

# North American MVPDays Community Roadshow

The purpose of this book is to showcase the fantastic expertise of our guest speakers at the North American MVPDays Community Roadshow. They have so much passion, expertise, and expert knowledge that it only seemed fitting to write it down in a book.

MVPDays was founded by Cristal and Dave Kawula back in 2013. It started as a simple idea; “There’s got to be a good way for Microsoft MVPs to reach the IT community and share their vast knowledge and experience in a fun and engaging way” I mean, what is the point in recognizing these bright and inspiring individuals, and not leveraging them to inspire the community that they are a part of.

We often get asked the question “Who should attend MVPDays”?

Anyone that has an interest in technology is eager to learn and wants to meet other like-minded individuals. This Roadshow is not just for Microsoft MVP’s it is for anyone in the IT Community.

Make sure you check out the MVPDays website [at](http://www.mvpdays.com) [www.mvpdays.com](http://www.mvpdays.com). You never know maybe the roadshow will be coming to a city near you.

The goal of this particular book is to give you some fantastic Veeam tips from the amazing Veeam Vanguard Community. Each chapter is broken down into a unique tip, and we hope you find some immense value in what we have written.

## Sample Files

All sample files for this book can be downloaded from [www.checkyourlogs.net](http://www.checkyourlogs.net) and [www.github.com/mvpdays](http://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

# Upgrading a Veeam Backup Repository from Server 2016 to Server 2019 + ReFS DeDuplication

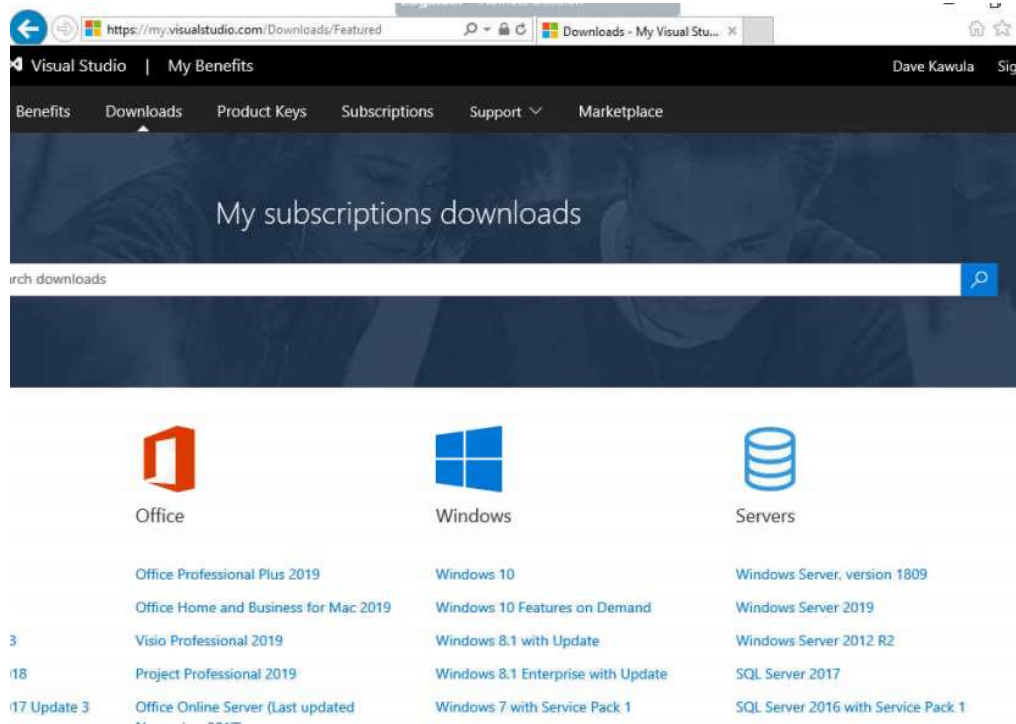
**By: Dave Kawula (Microsoft MVP / Veeam Vanguard)**

I was working with a customer today, and we are finally ready to upgrade our primary Veeam Backup Repository that is running Microsoft Storage Spaces on Windows Server 2016. The new Operating System will be Windows Server 2019 LTSC, and the whole point of this upgrade is to allow us to use ReFS + Deduplication.

The Deduplication feature is now supported in the LTSC builds of Windows Server and the purpose of this post today is to show you how to upgrade the OS Drives, Install the Deduplication Feature, Enable it and test.

## Upgrading to Windows Server 2019

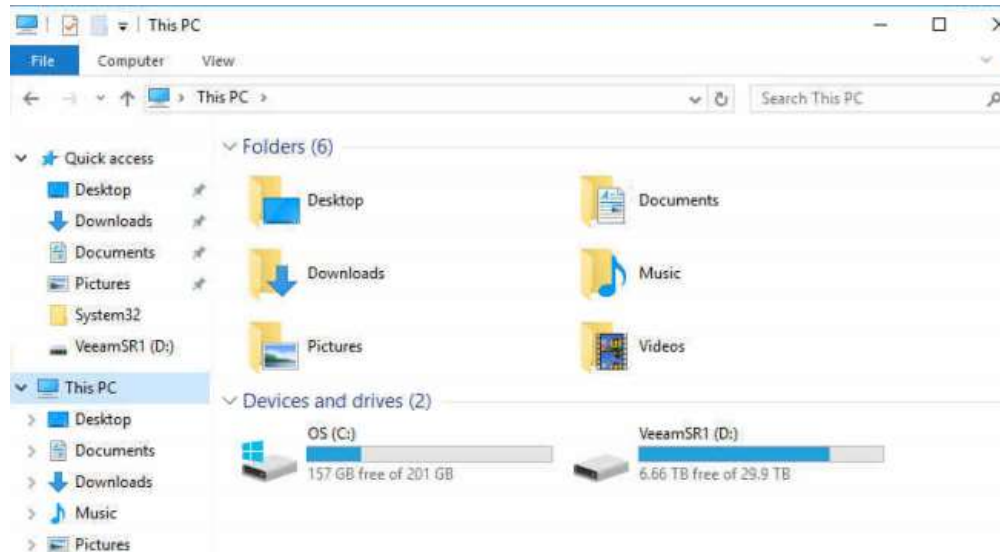
1. Download the media from your Volume License Site (VLK)



2. Mount the ISO in the Target system that you want to be upgraded.

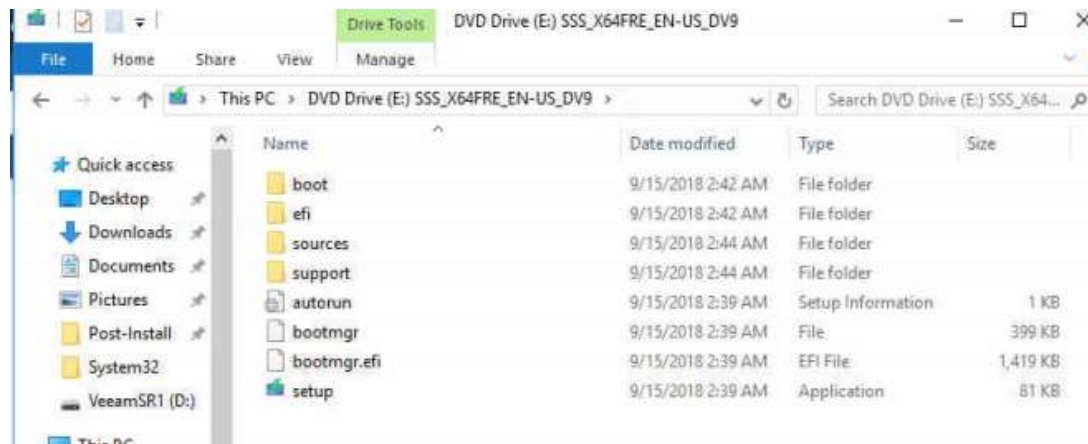


3. Before we pull the trigger on the upgrade, have a look at your Storage Pool's Virtual Disk for a before picture.

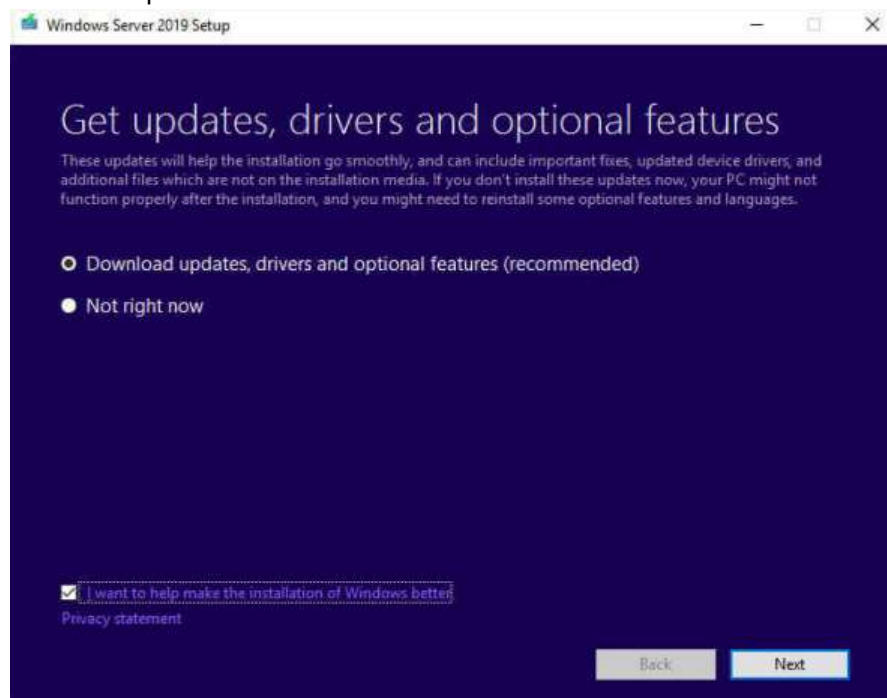


As you can see we have 6.66 TB of 29.9 TB free. We have this Storage Pool setup in a 3-Way Mirror for maximum performance for the Veeam Backup Repository.

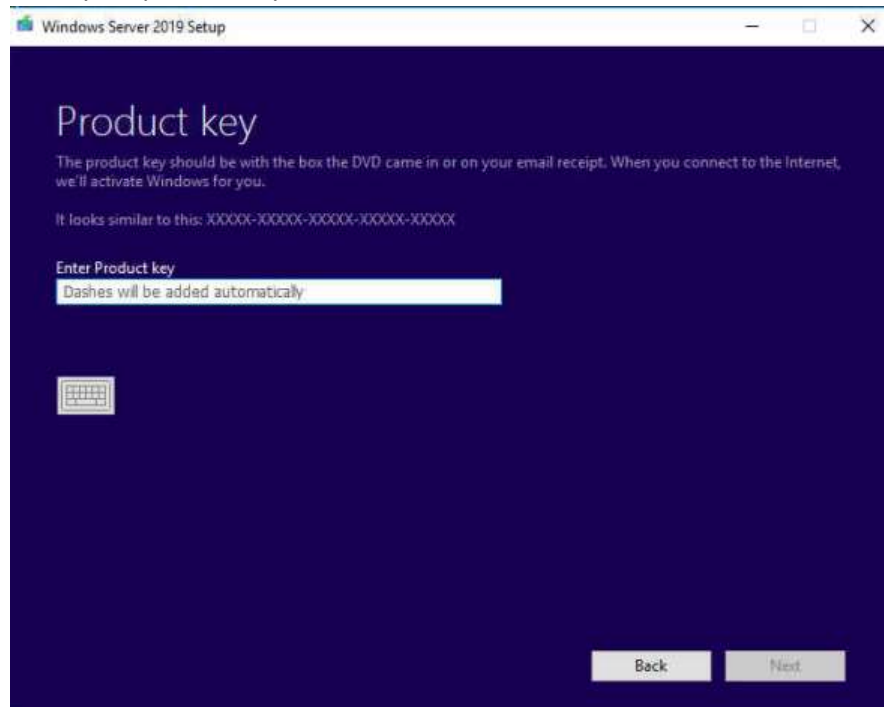
4. Run Setup.exe locally to start the upgrade



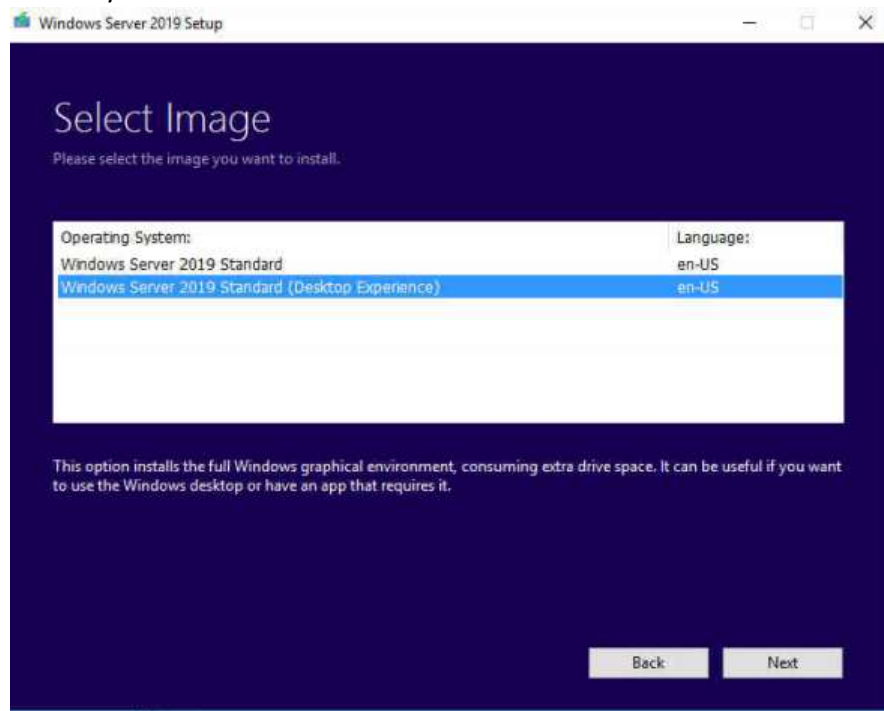
5. Choose Download updates, Drivers, and optional features (Recommended) and select I want to help make the installation of Windows Better checkbox



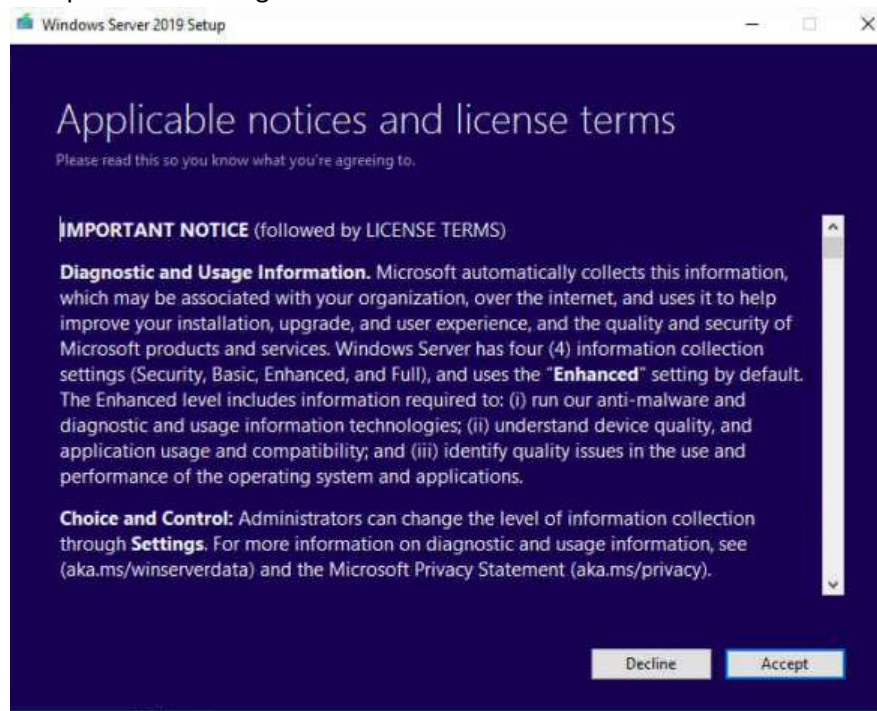
6. Enter your product key



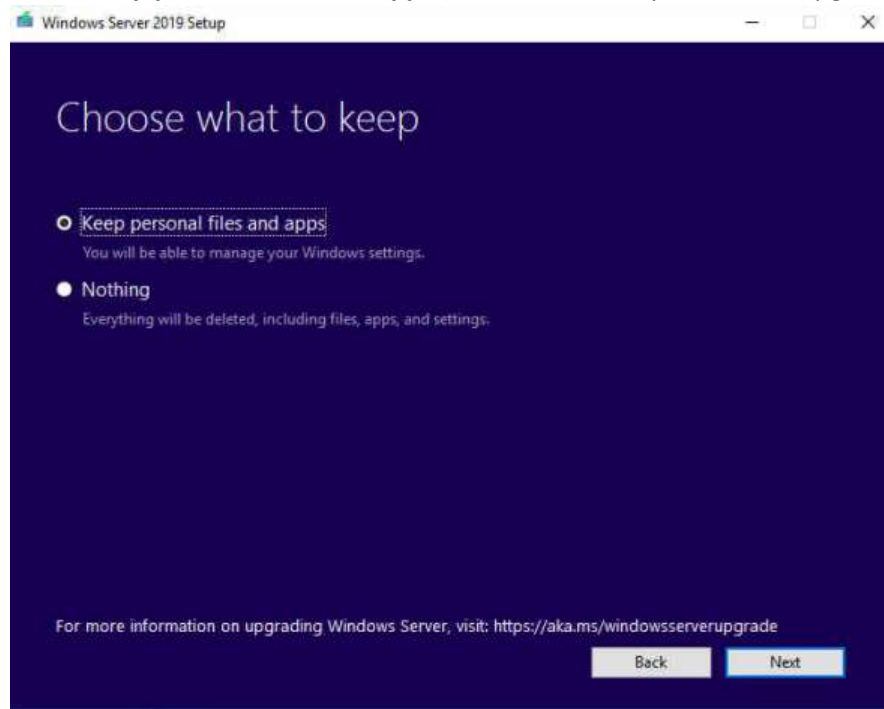
7. Choose your edition



8. Accept the License Agreement

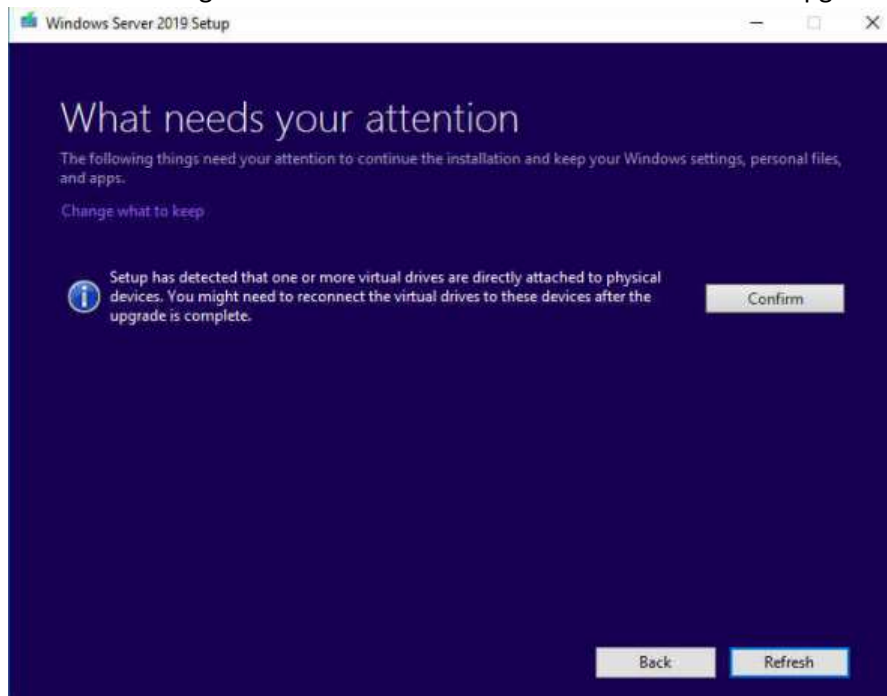


9. Select **keep personal files and apps** (These means do you want to upgrade)

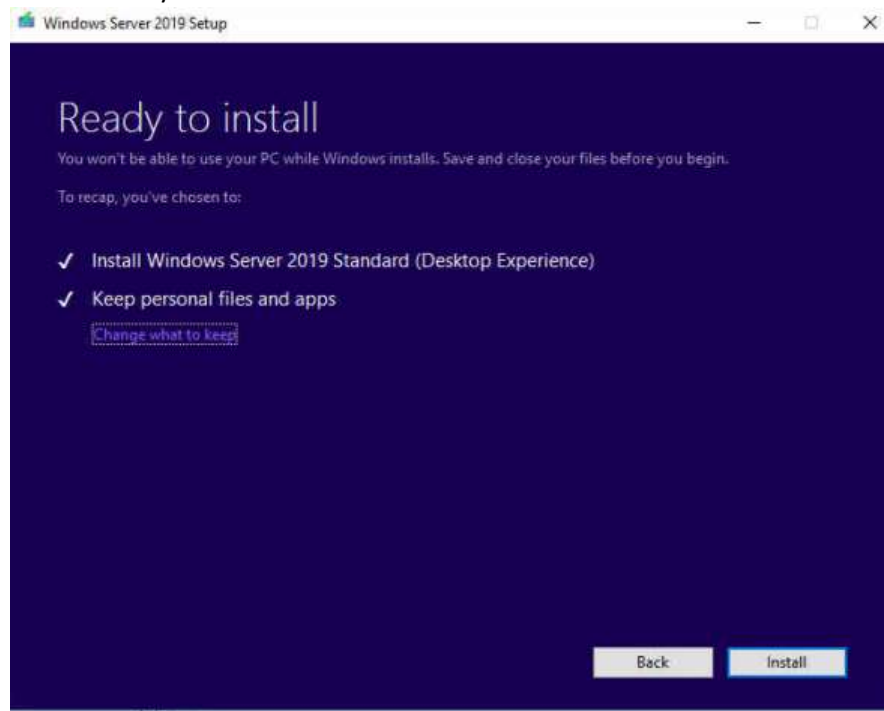


10. Because we are running a Storage Spaces Pool and have a Virtual Disk → Choose Confirm on Setup has detected that one or more virtual drives are directly attached to physical

devices. You might need to reconnect the virtual drives after the upgrade is complete.



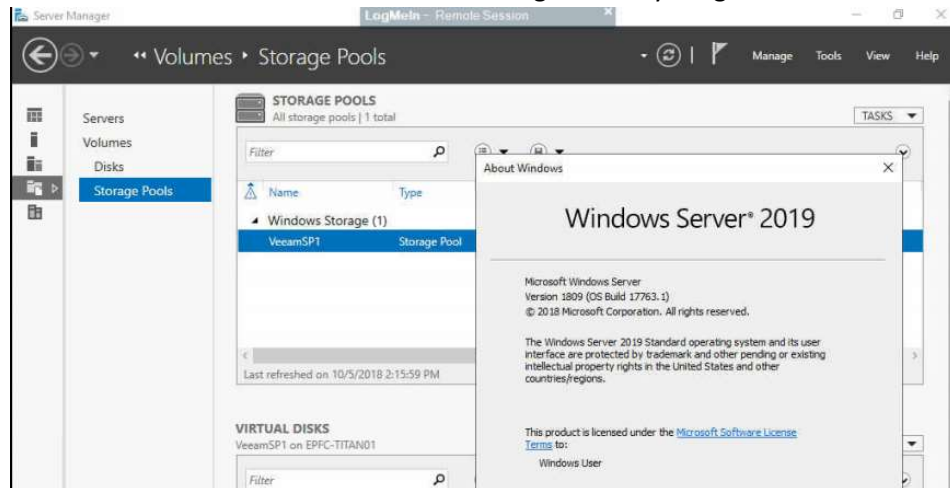
11. On the Ready to Install screen click **Install**



12. Grab a cup of coffee and wait for the upgrade to complete.

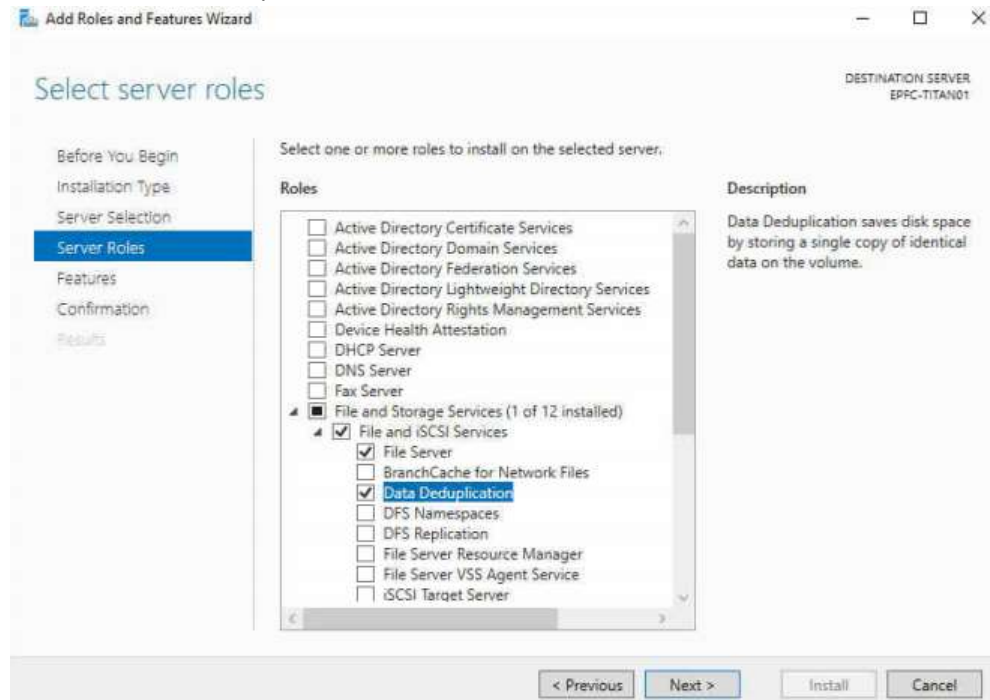


13. 30 minutes later and we are back in business. The good news is our Storage Pool and Volumes all came back online without having to do anything.

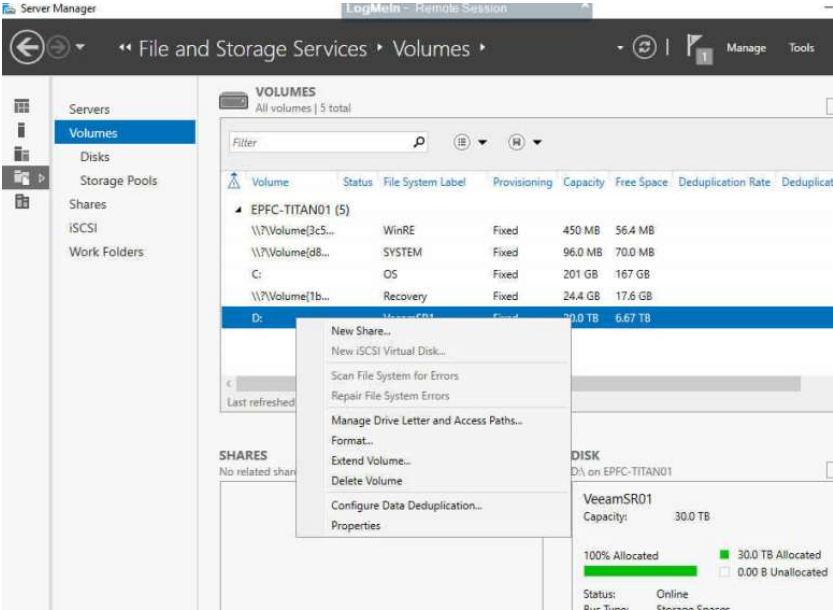


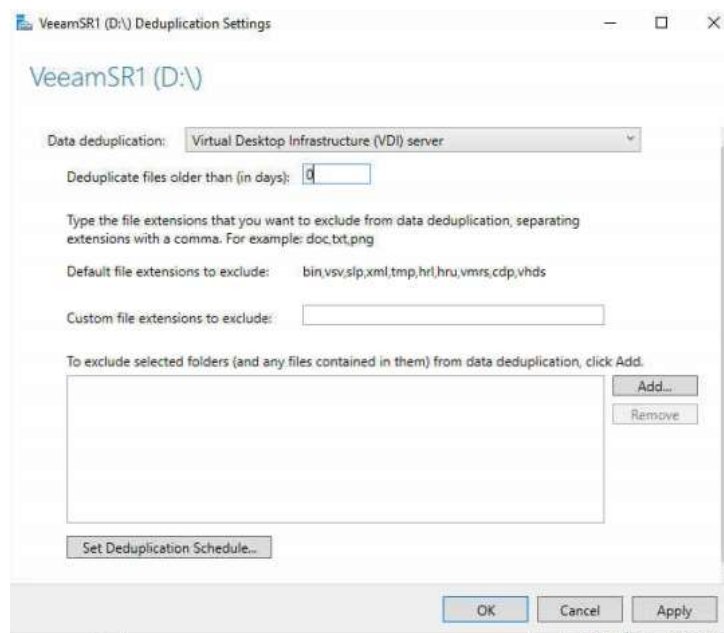
## Install and Configure Windows Deduplication

1. Let's Install the Deduplication Feature now



2. Configure Deduplication for your Veeam Repository





EPFC-TITAN01 Deduplication Schedule

EPFC-TITAN01

☒ Enable background optimization

Regularly run data deduplication at low priority and pause data deduplication when the system is busy to minimize the impact on system performance.

☒ Enable throughput optimization

During the specified hours, run data deduplication at normal priority and consume the resources required to maximize performance.

Days of the week: ☒ Sunday ☒ Monday ☒ Tuesday ☒ Wednesday  
☒ Thursday ☒ Friday ☒ Saturday

Start time: 1:45 AM

Duration (in hours): 6

☒ Create a second schedule for throughput optimization

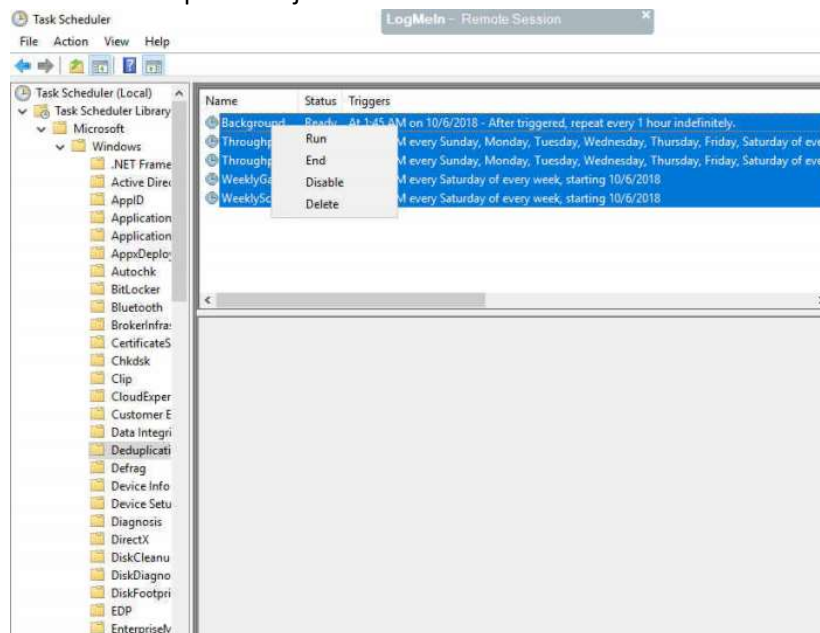
During the specified hours, run data deduplication at normal priority and consume the resources required to maximize performance.

Days of the week: ☒ Sunday ☒ Monday ☒ Tuesday ☒ Wednesday  
☒ Thursday ☒ Friday ☒ Saturday

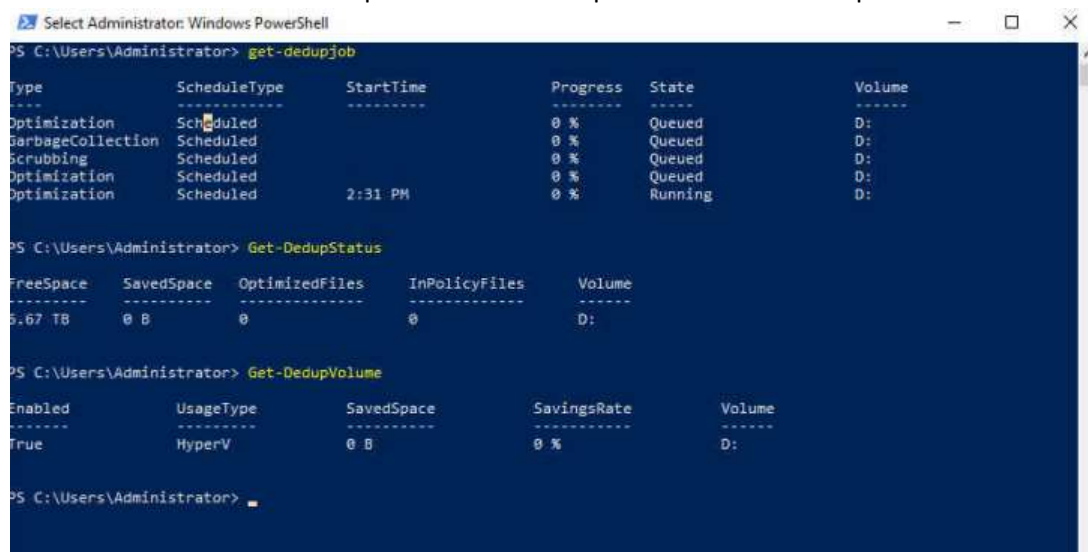
Start time: 9:00 AM

Duration (in hours): 8

### 3. Start the Deduplication jobs via Task Scheduler



### 4. View the results via Get-DedupJob and Get-DedupStatus and Get-DedupVolume

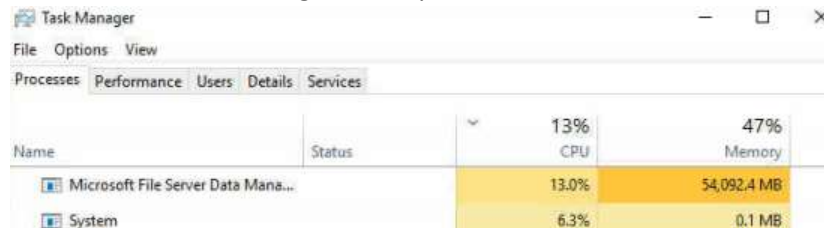


You can continue to view the progress until all of the Deduplication jobs are finished.

So, I would suggest it's coffee time again while you wait.

## Tweak Windows Deduplication

1. Monitor the progress. You can also do this by watching it from Task Manager and selecting the Microsoft File Server Data Management Host Process. This is the Deduplication engine running. Interestingly enough this was consuming 20 % of the CPU and 54 GB of RAM during the first pass.



The screenshot shows the Windows Task Manager Performance tab. The 'Processes' tab is selected, and the 'Performance' sub-tab is active. The table displays the following data:

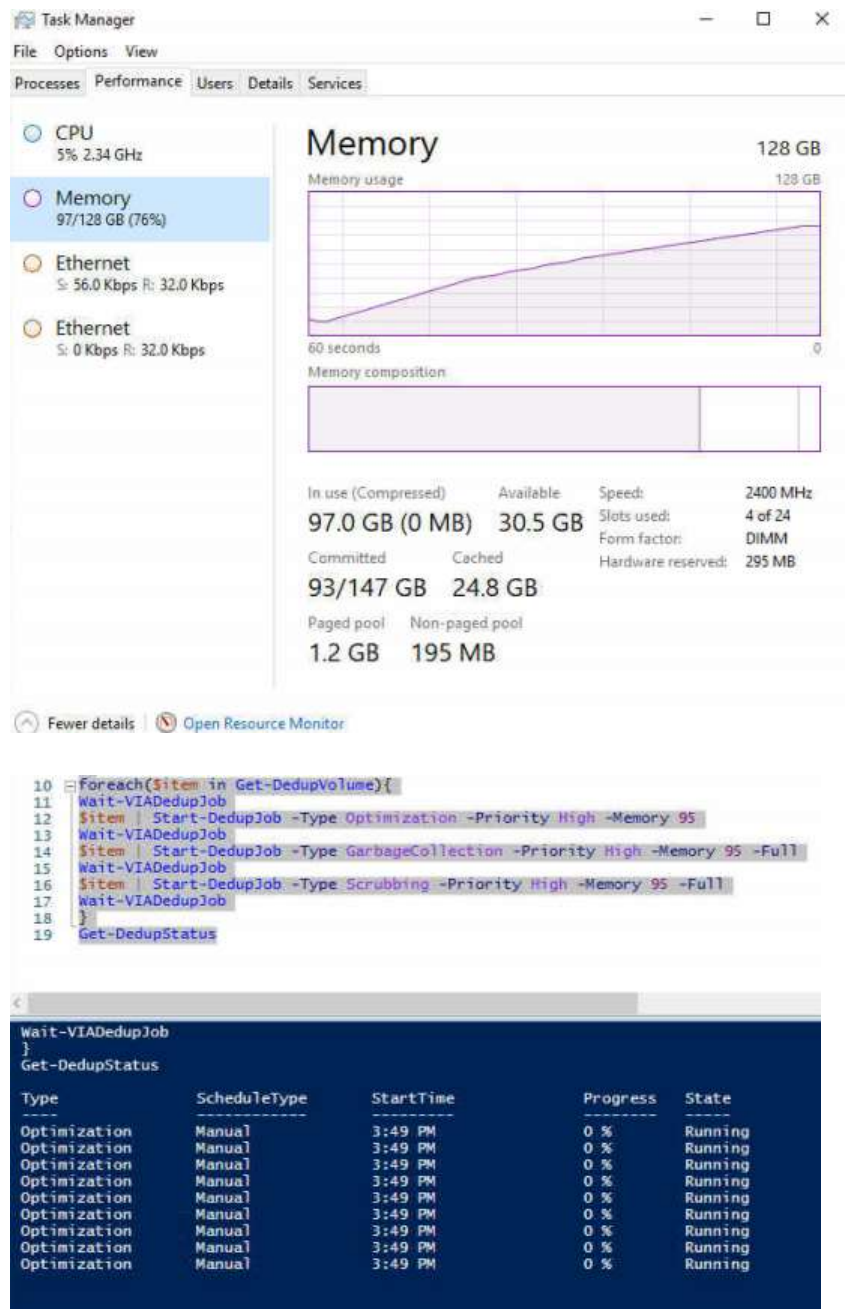
Name	Status	CPU	Memory
Microsoft File Server Data Mana...		13.0%	54,092.4 MB
System		6.3%	0.1 MB

2. View the final results. As a side note, I'm impatient when trying something out, so I wanted to push the system on the first pass to speed it up. The initial 54 GB of RAM wasn't enough, so I did this. I grabbed a piece of code from my friend Mikael Nystrom <https://deploymentbunny.com/2017/04/28/powershell-is-king-a-data-deduplication-script-that-run-the-optimizationgarbage-collection-and-scrubbing-in-a-single-sweep-including-progress/>

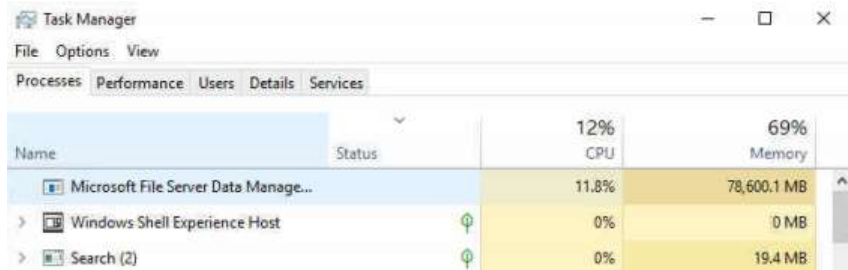
```
Function Wait-VIADedupJob
{
while ((Get-DedupJob).count -ne 0 )
{
Get-DedupJob
Start-Sleep -Seconds 30
}
}
```

```
foreach($item in Get-DedupVolume) {  
    Wait-VIADedupJob  
  
    $item | Start-DedupJob -Type Optimization -Priority High -Memory  
    80  
  
    Wait-VIADedupJob  
  
    $item | Start-DedupJob -Type GarbageCollection -Priority High -  
    Memory 80 -Full  
  
    Wait-VIADedupJob  
  
    $item | Start-DedupJob -Type Scrubbing -Priority High -Memory 80  
    -Full  
  
    Wait-VIADedupJob  
}  
  
Get-DedupStatus
```

3. There this looks better now:



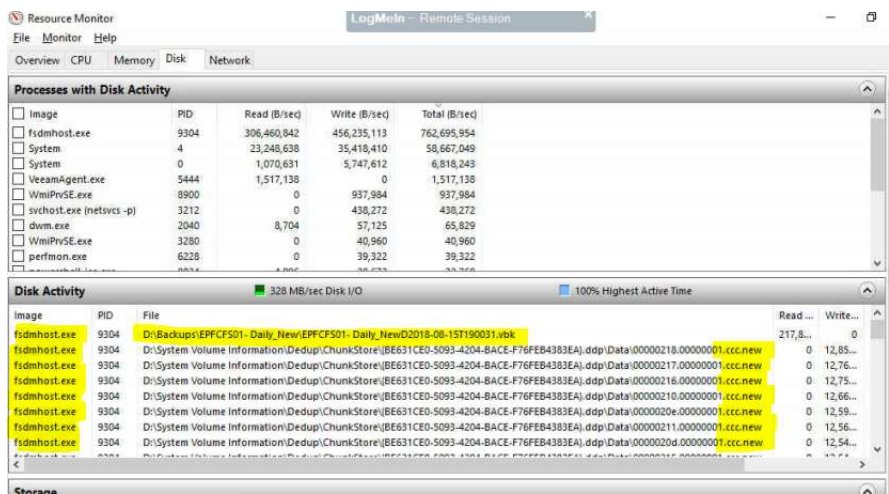
4. Here is the memory usage for the Microsoft File Server Data Management Process now.



The screenshot shows the Windows Task Manager window with the 'Processes' tab selected. The 'Microsoft File Server Data Management Process' is highlighted, showing it is using 11.8% CPU and 78,600.1 MB of memory. Other processes shown include 'Windows Shell Experience Host' and 'Search (2)'.

Name	Status	CPU	Memory
Microsoft File Server Data Manage...		11.8%	78,600.1 MB
Windows Shell Experience Host		0%	0 MB
Search (2)		0%	19.4 MB

5. If you are curious about what is happening and which file the Deduplication engine is working on your can open Resource Monitor and look for the process that is running **fsdmhost.exe**. You will see the file that is being read and then you will see it being broken up into the **Chuck Store**. In this case, we can see the files floating into the **D:\System Volume Information\Dedup\ChunkStore** → With an extension of **ccc.new**.



The screenshot shows the Windows Resource Monitor window with the 'Disk' tab selected. The 'Processes with Disk Activity' section lists several processes, including 'fsdmhost.exe'. The 'Disk Activity' section shows the file being read by 'fsdmhost.exe' is 'D:\System Volume Information\Dedup\ChunkStore\B6631CE0-5093-4204-BACE-F76FEB4383EA\ddp\Data\00000218.00000001.ccc.new'.

Image	PID	Read (B/sec)	Write (B/sec)	Total (B/sec)
fsdmhost.exe	9304	306,460,842	456,235,113	762,695,954
System	4	23,248,638	35,418,410	58,667,049
System	0	1,070,631	5,747,612	6,818,243
VeeamAgent.exe	5444	1,517,138	0	1,517,138
WmiPnVSE.exe	8900	0	937,984	937,984
svchost.exe (netsvc -p)	3212	0	438,272	438,272
dwm.exe	2040	8,704	57,125	65,829
WmiPnVSE.exe	3280	0	40,960	40,960
perfmon.exe	6228	0	39,322	39,322

Image	PID	File	Read...	Write...
fsdmhost.exe	9304	D:\Backups\EPFCF501-Daily_New\EPFCF501-Daily_New02018-08-15T190031.vbk	217,8...	0
fsdmhost.exe	9304	D:\System Volume Information\Dedup\ChunkStore\B6631CE0-5093-4204-BACE-F76FEB4383EA\ddp\Data\00000217.00000001.ccc.new	0	12,85...
fsdmhost.exe	9304	D:\System Volume Information\Dedup\ChunkStore\B6631CE0-5093-4204-BACE-F76FEB4383EA\ddp\Data\00000218.00000001.ccc.new	0	12,76...
fsdmhost.exe	9304	D:\System Volume Information\Dedup\ChunkStore\B6631CE0-5093-4204-BACE-F76FEB4383EA\ddp\Data\00000219.00000001.ccc.new	0	12,75...
fsdmhost.exe	9304	D:\System Volume Information\Dedup\ChunkStore\B6631CE0-5093-4204-BACE-F76FEB4383EA\ddp\Data\00000220.00000001.ccc.new	0	12,66...
fsdmhost.exe	9304	D:\System Volume Information\Dedup\ChunkStore\B6631CE0-5093-4204-BACE-F76FEB4383EA\ddp\Data\00000221.00000001.ccc.new	0	12,59...
fsdmhost.exe	9304	D:\System Volume Information\Dedup\ChunkStore\B6631CE0-5093-4204-BACE-F76FEB4383EA\ddp\Data\00000222.00000001.ccc.new	0	12,56...
fsdmhost.exe	9304	D:\System Volume Information\Dedup\ChunkStore\B6631CE0-5093-4204-BACE-F76FEB4383EA\ddp\Data\00000223.00000001.ccc.new	0	12,54...

Until I checked this, I thought that the jobs were stalled. I was wrong the Deduplication Engine was working through some very large Veeam Backup.VBK files.

NOTE: My estimates on how long the initial pass will take are as follows. I have seen the Disk MB/Sec running consistently at around 300 MB/sec. Your calculation will look something like this  $300\text{MB/sec} * 60 \text{ Seconds} * 60 \text{ Minutes} = \text{MB Per Hours Processing} = \text{Roughly 1 TB per Hour}$ . If there are 24 TB's of data which in our case there is. It will take approximately 24 Hours for the initial pass to complete.

6. You can see the time the job started and stopped by checking in the Deduplication Event Log. Look for Event ID 6148 you will see the message: **Optimization job has started.**

The screenshot displays the Windows Event Viewer interface. On the left, the 'Deduplication' folder is expanded under 'System' logs. The main pane shows a list of events with columns: Level, Date and Time, Source, Event ID, and Task Category. Event ID 6148 is highlighted. Below the list, the 'Details' tab for Event 6148 is open, showing the message 'Optimization job has started.' and a volume path. Metadata includes Log Name: Data Deduplication/Operational, Source: Data Deduplication, Event ID: 6148, Level: Information, User: SYSTEM, OpCode: Info, and Computer: EPFC-TITAN01.

Level	Date and Time	Source	Event ID	Task Category
Info	10/5/2018 3:49:07 PM	Data Deduplication	6148	None
Info	10/5/2018 3:49:07 PM	Data Deduplication	6158	None
Info	10/5/2018 3:49:07 PM	Data Deduplication	6153	None
Info	10/5/2018 3:49:07 PM	Data Deduplication	4185	None
Info	10/5/2018 3:49:07 PM	Data Deduplication	4103	None
Info	10/5/2018 3:49:07 PM	Data Deduplication	4185	None
Info	10/5/2018 3:49:07 PM	Data Deduplication	4185	None
Info	10/5/2018 3:49:07 PM	Data Deduplication	4103	None
Info	10/5/2018 3:49:07 PM	Data Deduplication	4185	None
Info	10/5/2018 3:49:07 PM	Data Deduplication	4185	None
Info	10/5/2018 3:49:07 PM	Data Deduplication	4103	None
Info	10/5/2018 3:49:07 PM	Data Deduplication	4185	None
Info	10/5/2018 3:49:07 PM	Data Deduplication	4103	None

Event 6148, Data Deduplication

General Details

Optimization job has started.

Volume: D: (\Volume{ef089fb2-d42f-4448-8a35-71fdaef2af9f}\) )

Log Name: Data Deduplication/Operational

Source: Data Deduplication Logged: 10/5/2018 3:49:07 PM

Event ID: 6148 Task Category: None

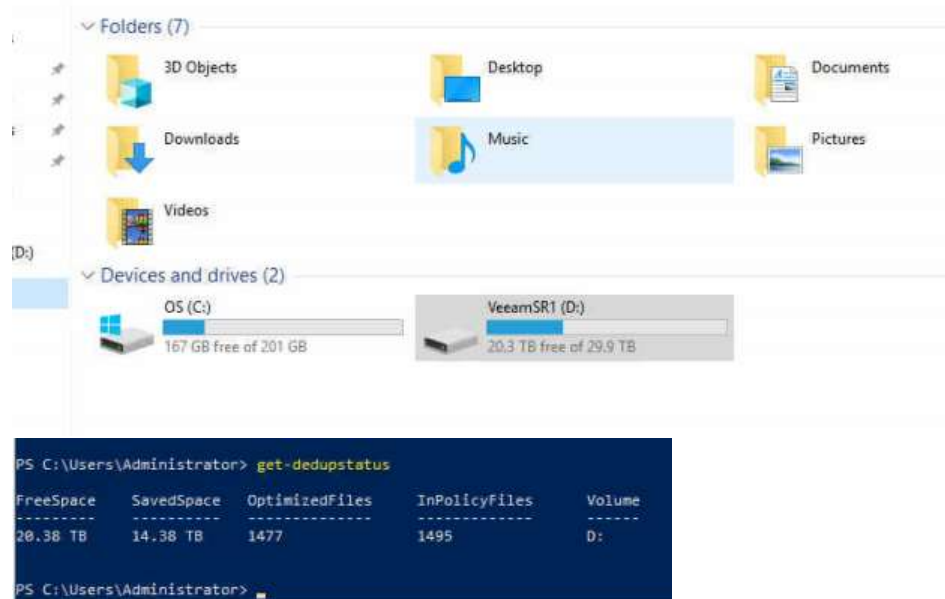
Level: Information Keywords:

User: SYSTEM Computer: EPFC-TITAN01

OpCode: Info

More Information: [Event Log Online Help](#)

7. We can check for the completion event ID when we come back to this later for the complete duration of the post-processing. Moreover, after one day which was pretty much my estimate the initial pass completed. I've now gained back about 14 TB of space post-Windows 2019 Deduplication.



I hope you enjoyed reading this and happy upgrading to Windows Server 2019 with Deduplication on ReFS.

Dave

## Chapter 2

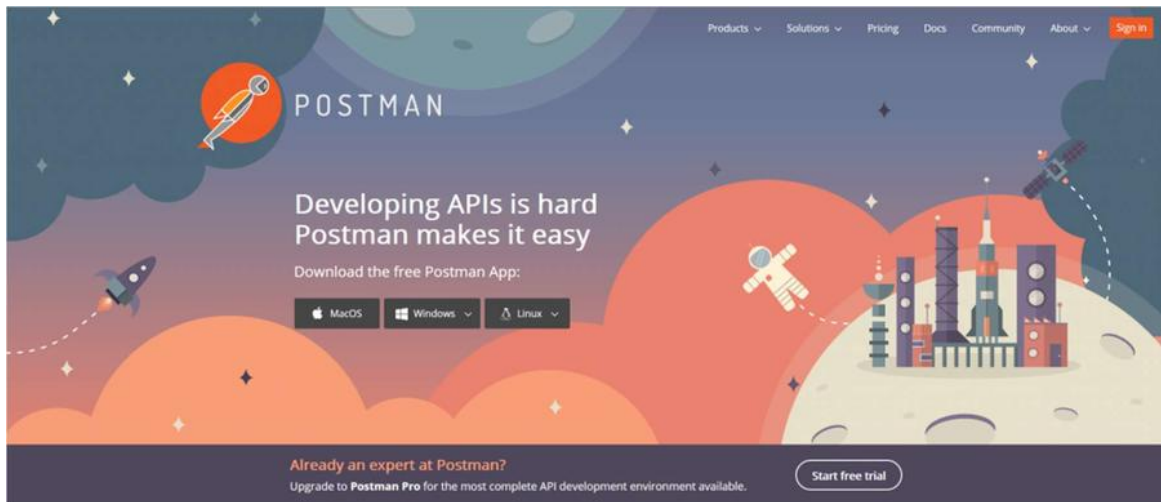
# Veeam RESTful API and Postman

**By: Craig Dalrymple (VMWare vEXPERT / Veeam Vanguard)**

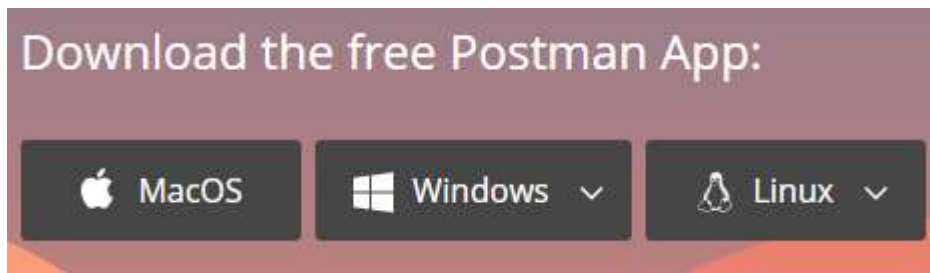
Recently I've been trying to get to grips with RESTful API on a couple of solutions I use at work, and as a relative newcomer to the whole world of RESTful APIs, I was looking for a simple/easy way to query systems and view the responses. One of my colleagues, Matt Thompson, pointed me in the direction of Postman [www.getpostman.com](http://www.getpostman.com).

## Postman

Postman is a popular REST client, with an intuitive user interface that allows users to send requests, save responses, add tests and create workflows

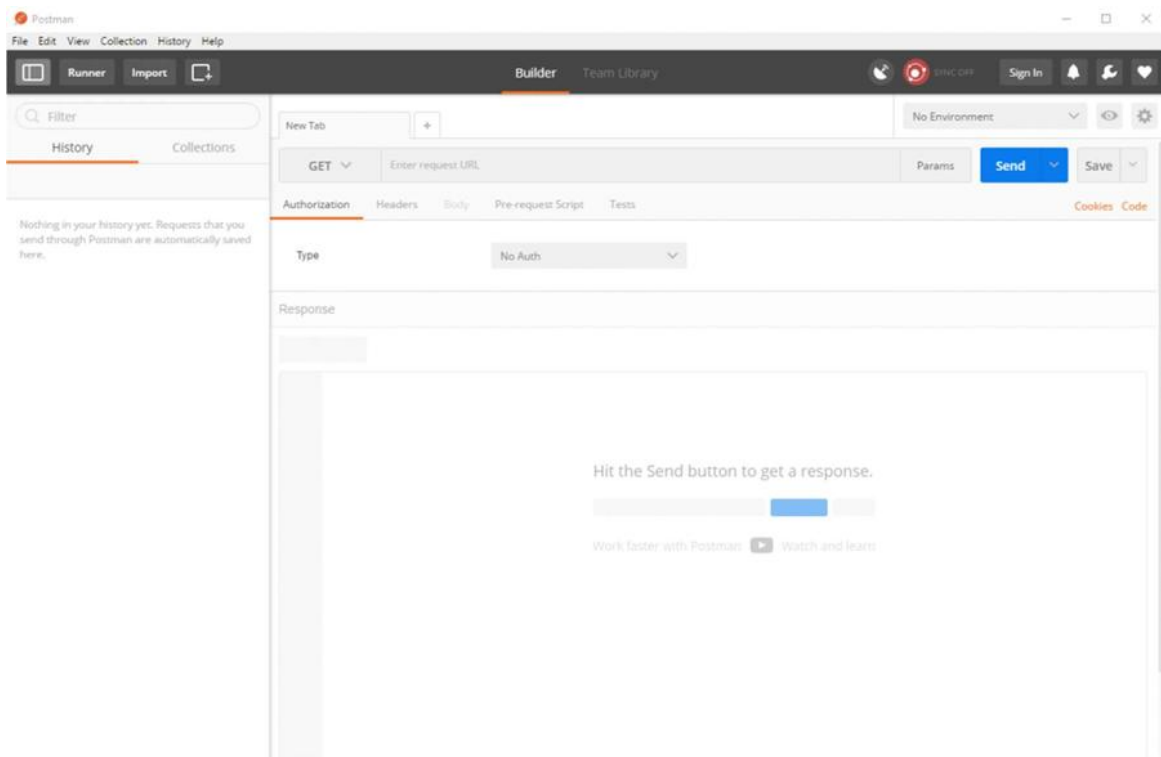


You can download your flavor of the app, in my case Windows x64, from <https://www.getpostman.com/>



Installation is incredibly straightforward, with no options to change/set, so a simple 'next, next, finish.'

Once installed the Postman dashboard looks like this



## Veeam Enterprise Manager

One of the applications I am interested in using RESTful APIs with is Veeam Backup and Replication (VBR). Veeam does have some excellent documentation around this:-

<https://helpcenter.veeam.com/docs/backup/rest/overview.html?ver=95>

The requirements for using the Veeam RESTful API is an Enterprise Manager installation, as this acts as a kind of 'proxy' for the API requests.

It should be pointed out that installing Enterprise Manager is very straightforward, is included in the installation ISO of VBR and is covered by your existing VBR license. I'm not going to cover the

installation and configuration of Enterprise Manager in this post, hoping to do so in a separate blog post.

## **RESTful API Authentication**

Before you can make API calls, you will need to authenticate your session. Veeam operates a 2 step authentication method when using the Veeam Backup Enterprise Manager RESTful API:

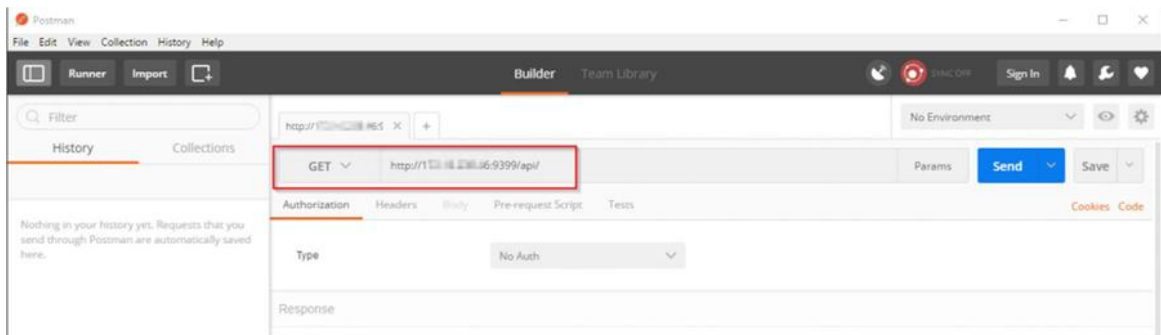
1. The user/client must log on to Veeam Backup Enterprise Manager with a valid username and password. These credentials will have been set previously on the Veeam Backup Enterprise Manager portal.
2. Using the credentials, the client then obtains an authorization token that must be used while making all API calls during the current login session.

## **Using Postman GET to obtain login sessions**

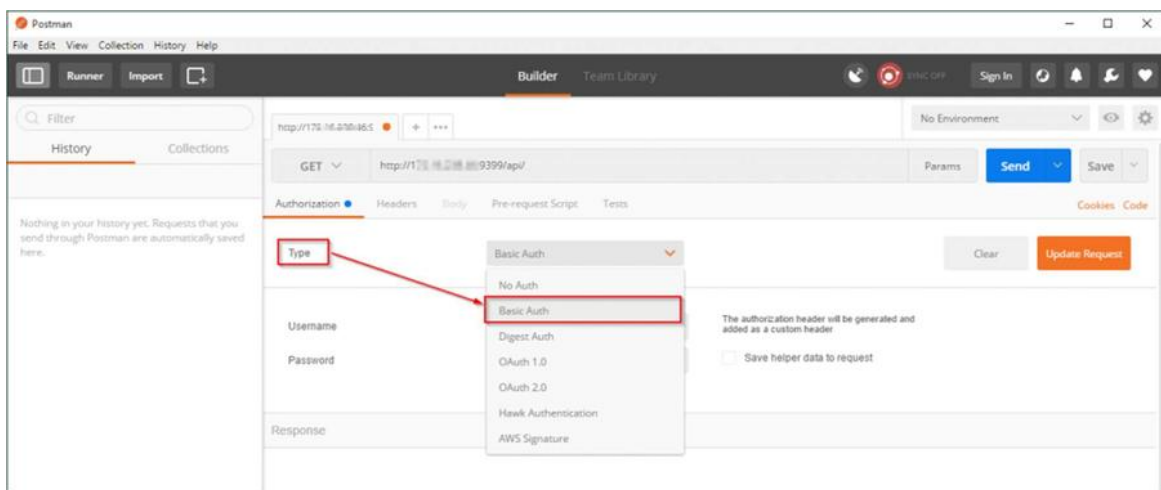
In the Postman dashboard, ensure you have set the action to 'POST' and then enter the following details in the 'Enter request URL' bar.

HTTP://<address\_of\_your\_Enterprise\_Manager>:9399/api/

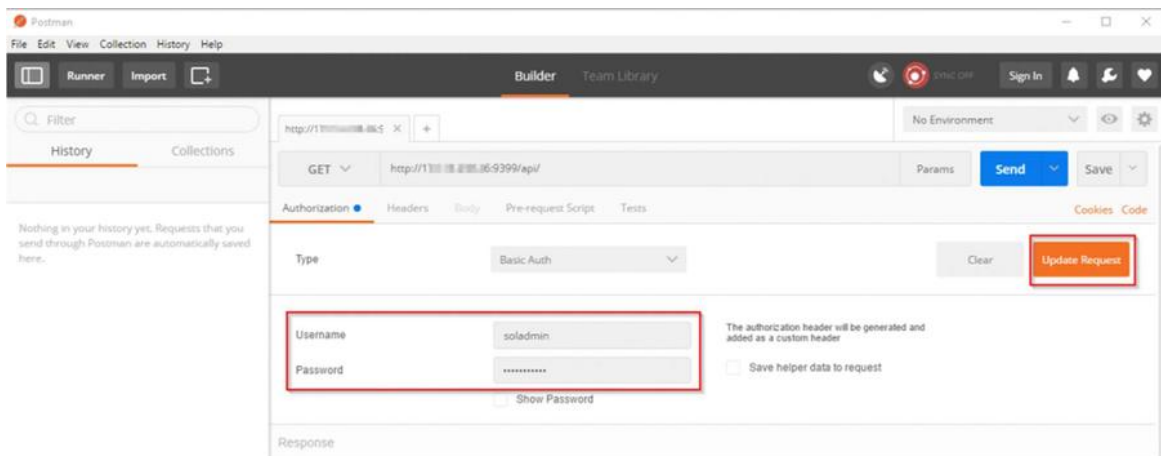
NOTE: the address can be either IP address or FQDN, in my examples I'm using the IP address



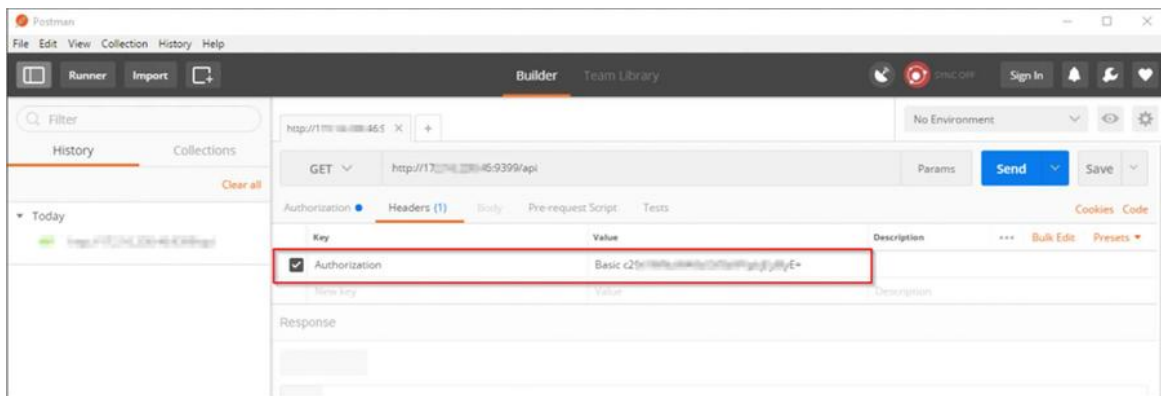
Now click on the 'Authorization' tab, and from the 'Type' drop down box select 'Basic Auth.'



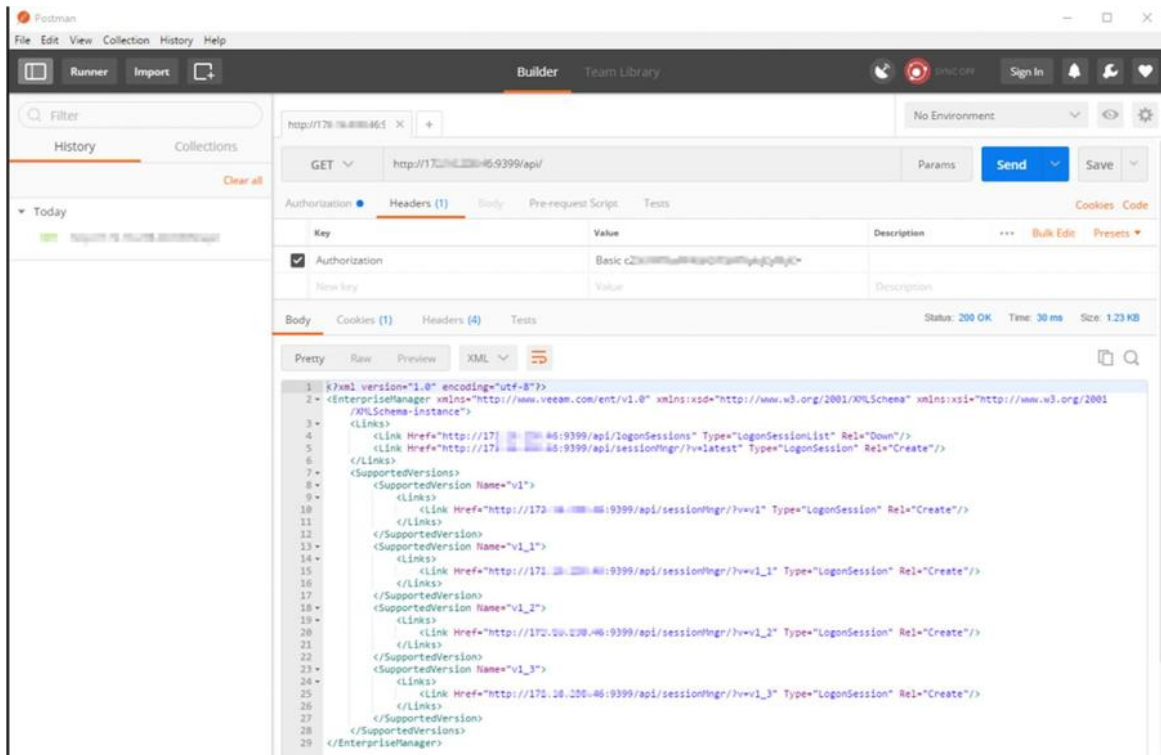
Enter the credentials you use to access the Veeam Backup Enterprise Manager console and click 'Update Request.'



The Veeam RESTful API looks for base64 encoded username and password, and thankfully Postman will add the correct encoding then add the relevant header. If you click on the **'Headers'** section, then you will see the relevant header format.



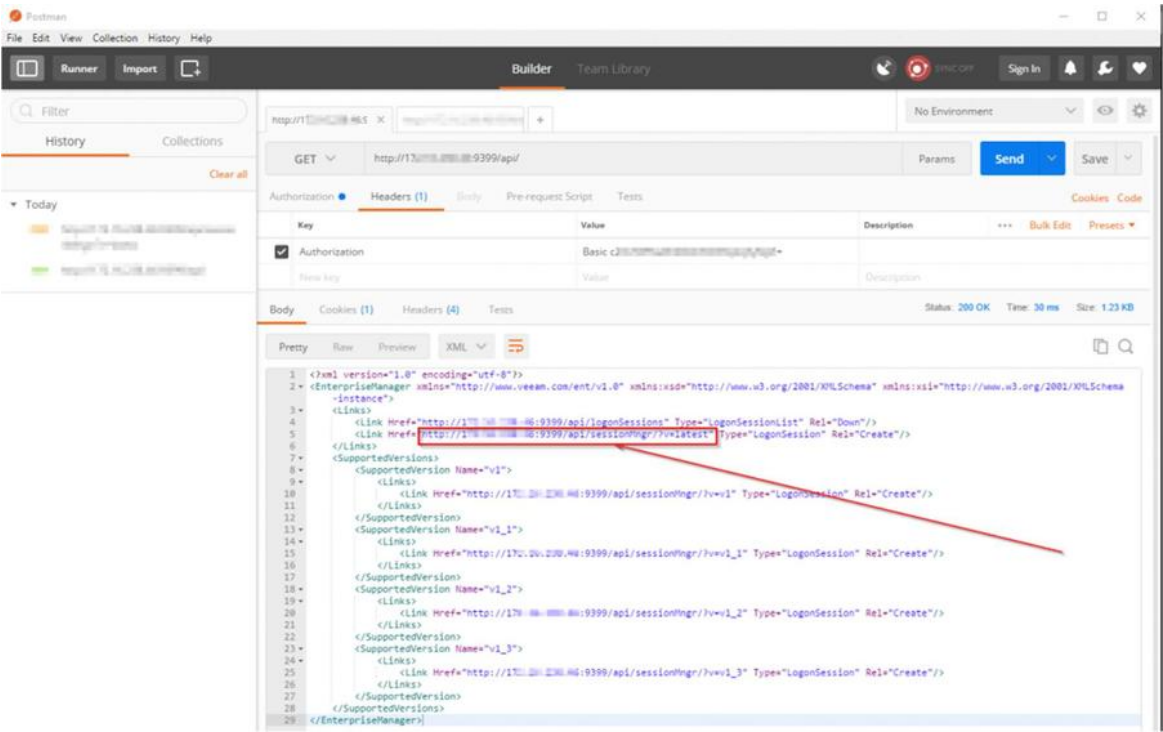
Click on 'Send' and if your session successfully authenticates you should see a similar response as follows:



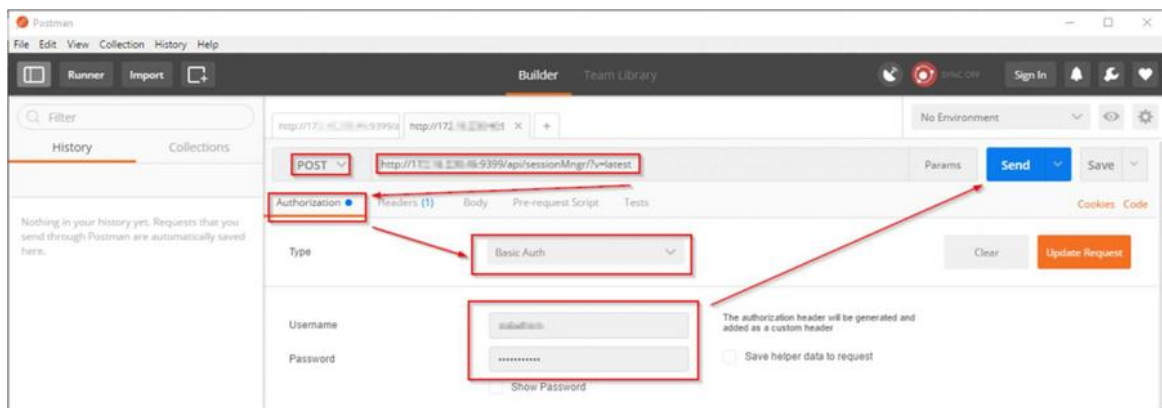
This verifies that the credentials used to access Veeam Backup Enterprise Manager are valid.

## Using Postman POST to obtain login session authentication token

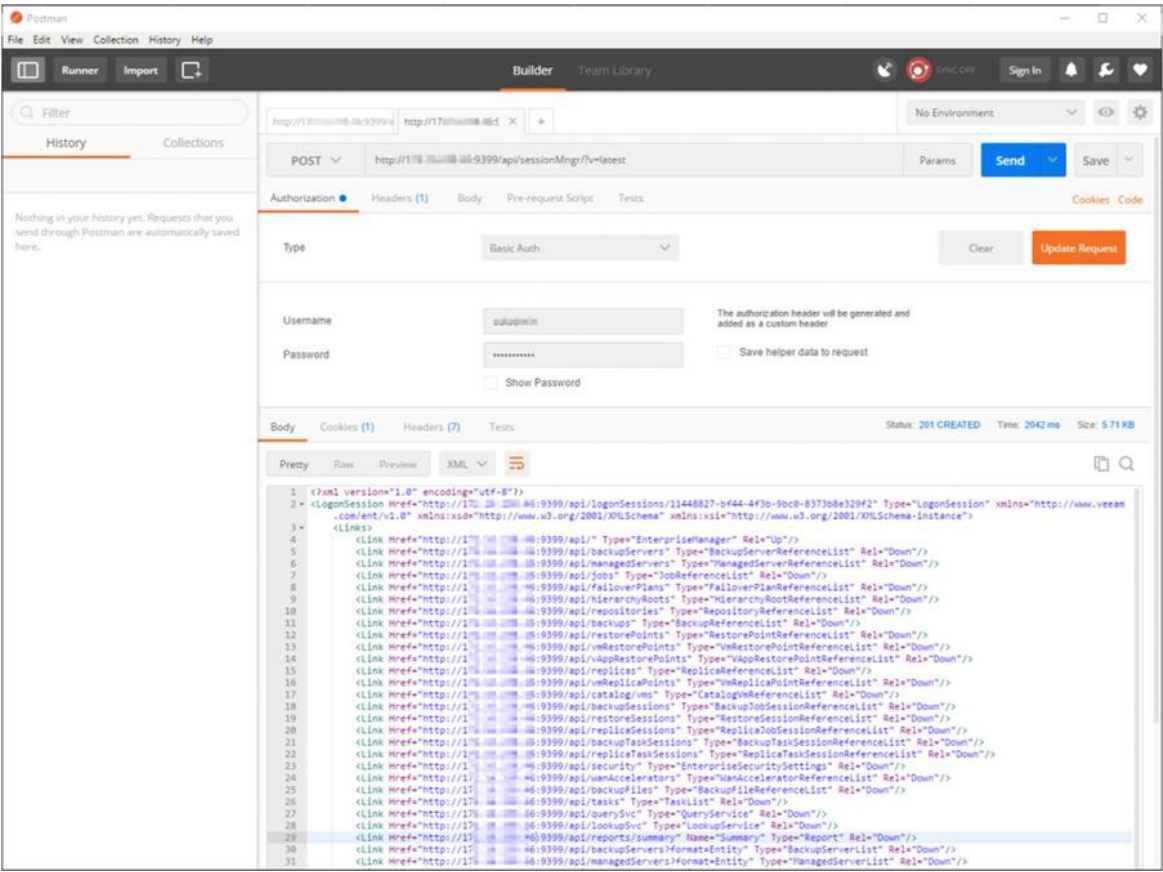
In the list of results on the previous screen, you can click on the link [http://<address\\_of\\_your\\_Enterprise\\_Manager>:9399/api/sessionMgr/?v=latest](http://<address_of_your_Enterprise_Manager>:9399/api/sessionMgr/?v=latest)



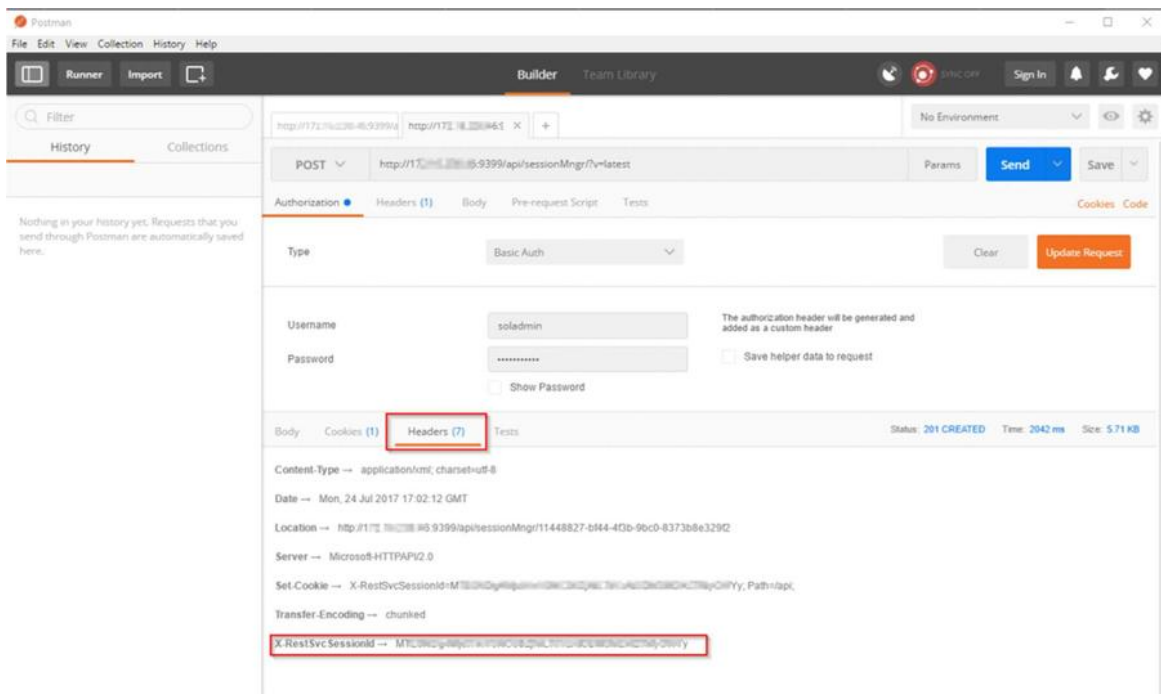
This will open a new tab within Postman. Change the action to '**POST**' and set the Authorization to '**Basic Auth.**' Re-enter the credentials for the Veeam Backup Enterprise Manager portal and click '**Send.**'



If successful you should receive back a list of the all the valid API sections



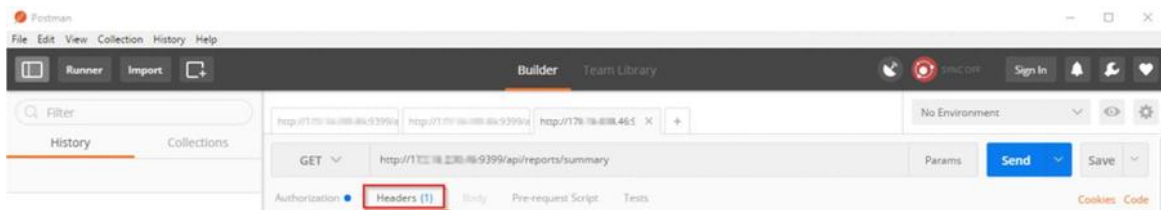
The authorization token can be found in the **‘Headers’** section in the body of the results. Click on **‘Headers’** section and look for the **X-RestSvcSessionId**



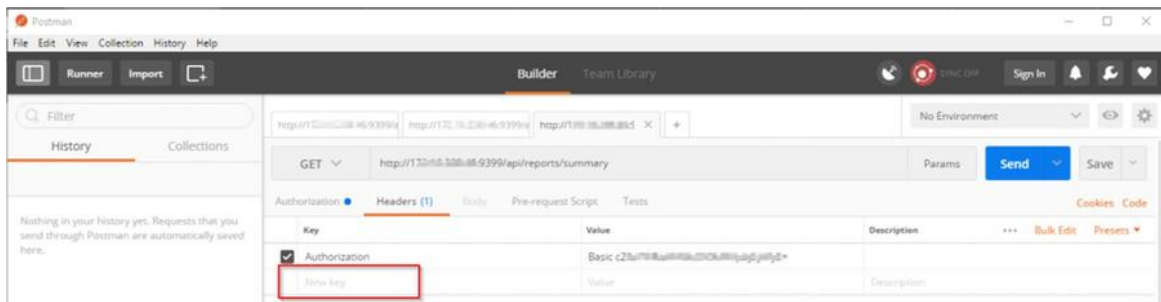
Any subsequent API calls made by the client must include this session id contained within the header. Note this value for later use.

## Adding a header

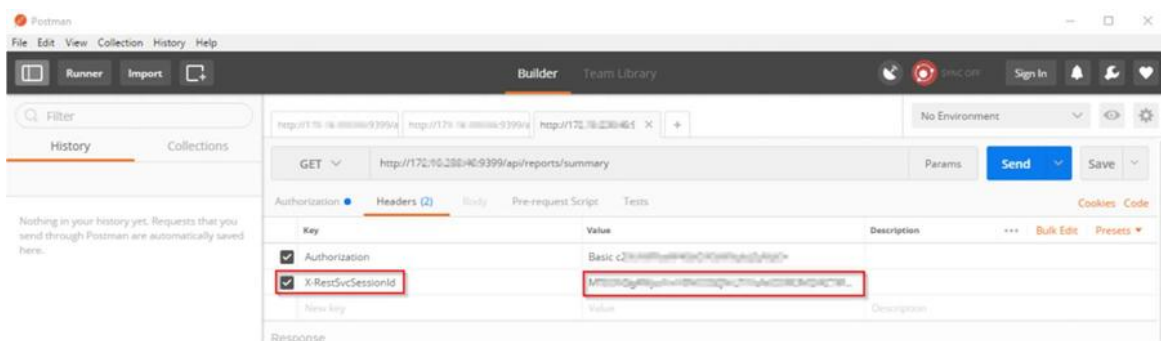
Once you have formed your request in the bar, click on 'Headers.'



Click on the 'New Key' line. NB if you set Basic Auth you will see the existing entry for the Authorization header.



Enter 'X-RestSvcSessionId' as the new key and the value recorded from the previous steps



Click send, and you're set to go.

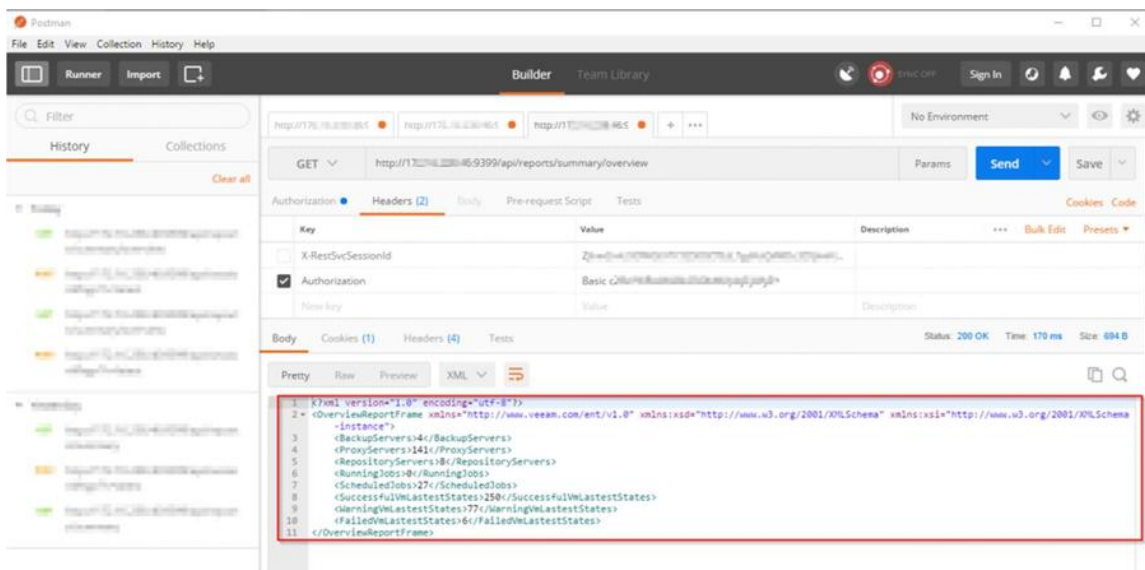
## Example of making API call and information

So now we can make API calls to our Veeam Backup Enterprise Manager, what kind of information can we retrieve?

Well, we can retrieve a summary overview of the Veeam Backup infrastructure components and jobs using the call:

**GET**

[HTTP://<ADDRESS OF YOUR ENTERPRISE MANAGER>:9399/api/reports/summary/overview](http://<ADDRESS OF YOUR ENTERPRISE MANAGER>:9399/api/reports/summary/overview)



The hope now is that I can leverage this information into a useful dashboard, possibly using Grafana or something similar.

## Summary

As a total newbie to RESTful API, I found using Postman very easy to use and extremely helpful at providing a relatively simple way to test API queries and view the results.

## Chapter 3

# Sizing Veeam for Office 365 Backup Storage

By: Ian Sanderson (VMWare vEXPERT / Veeam Vanguard)

## Background

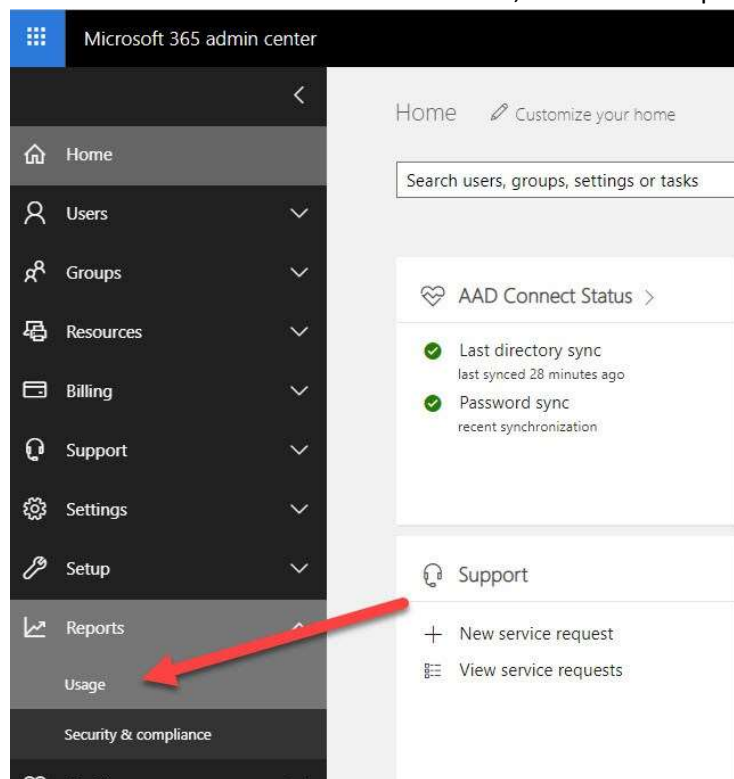
I recently expanded my Veeam for Office 365 deployment to protect new organizations but it struck me that I had no idea how much data I needed to protect from Office 365! If this was on-premises Exchange or SharePoint data, I could have looked at the database sizes to get an idea of how much data I would likely need to protect or run some PowerShell scripts to see how much space is consumed.

So how does one check the consumed storage space for Exchange Online, SharePoint Online and OneDrive for business?

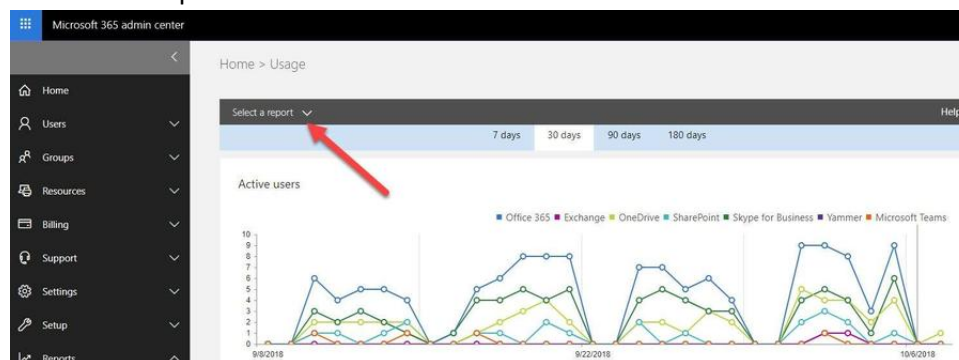
## Getting Started

1. First things first, get yourself logged into the Office 365 admin portal.

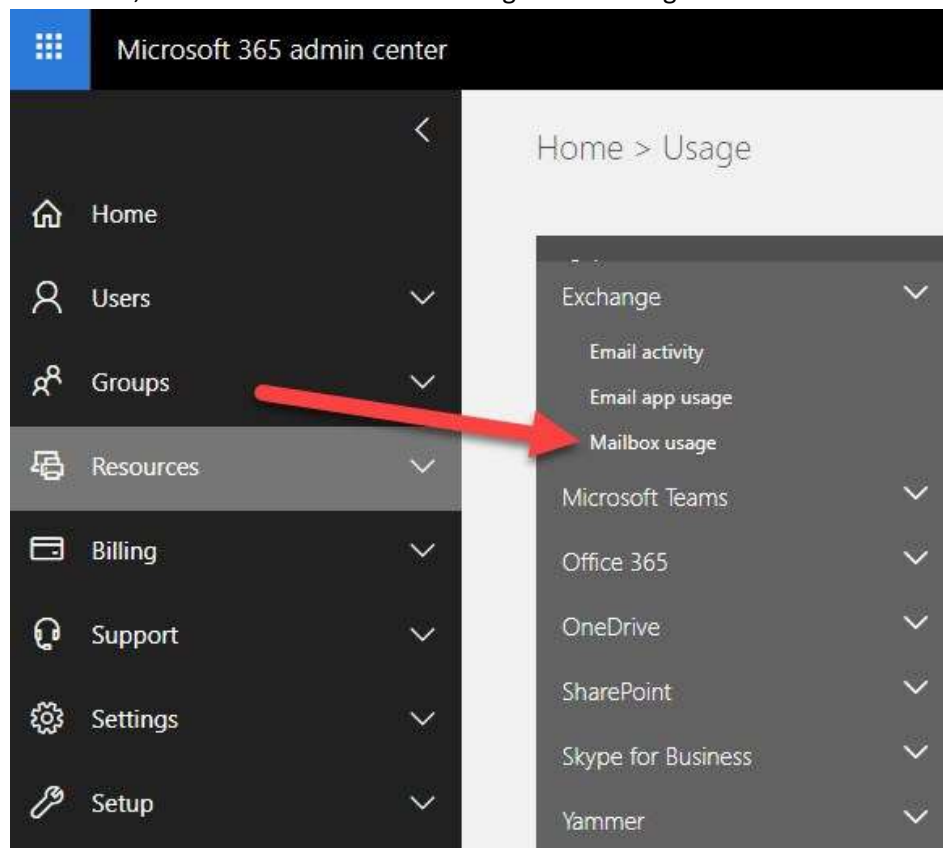
2. On the Microsoft Office 365 Admin Center, click on the Reports/Usage Page.



3. You will now see some usage stats for each application consumed from Office 365. Click on 'Select a report.'

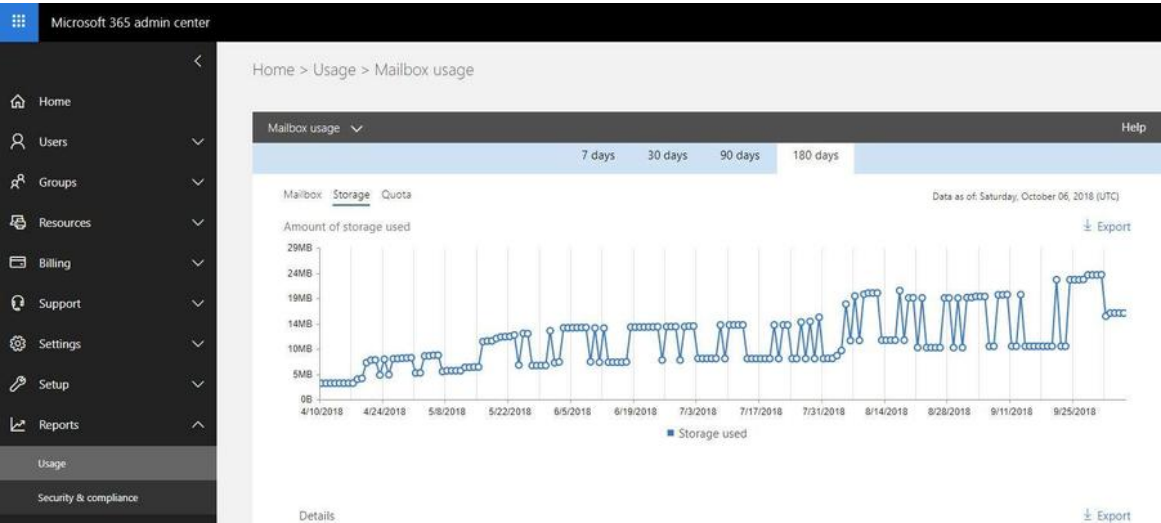


4. From here, let's take a look at the Exchange Online usage stats.

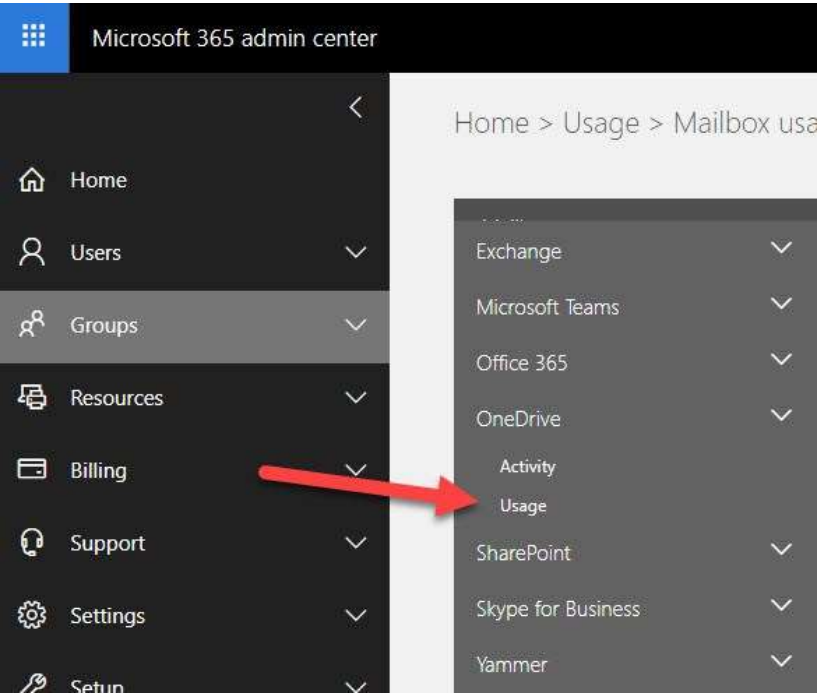


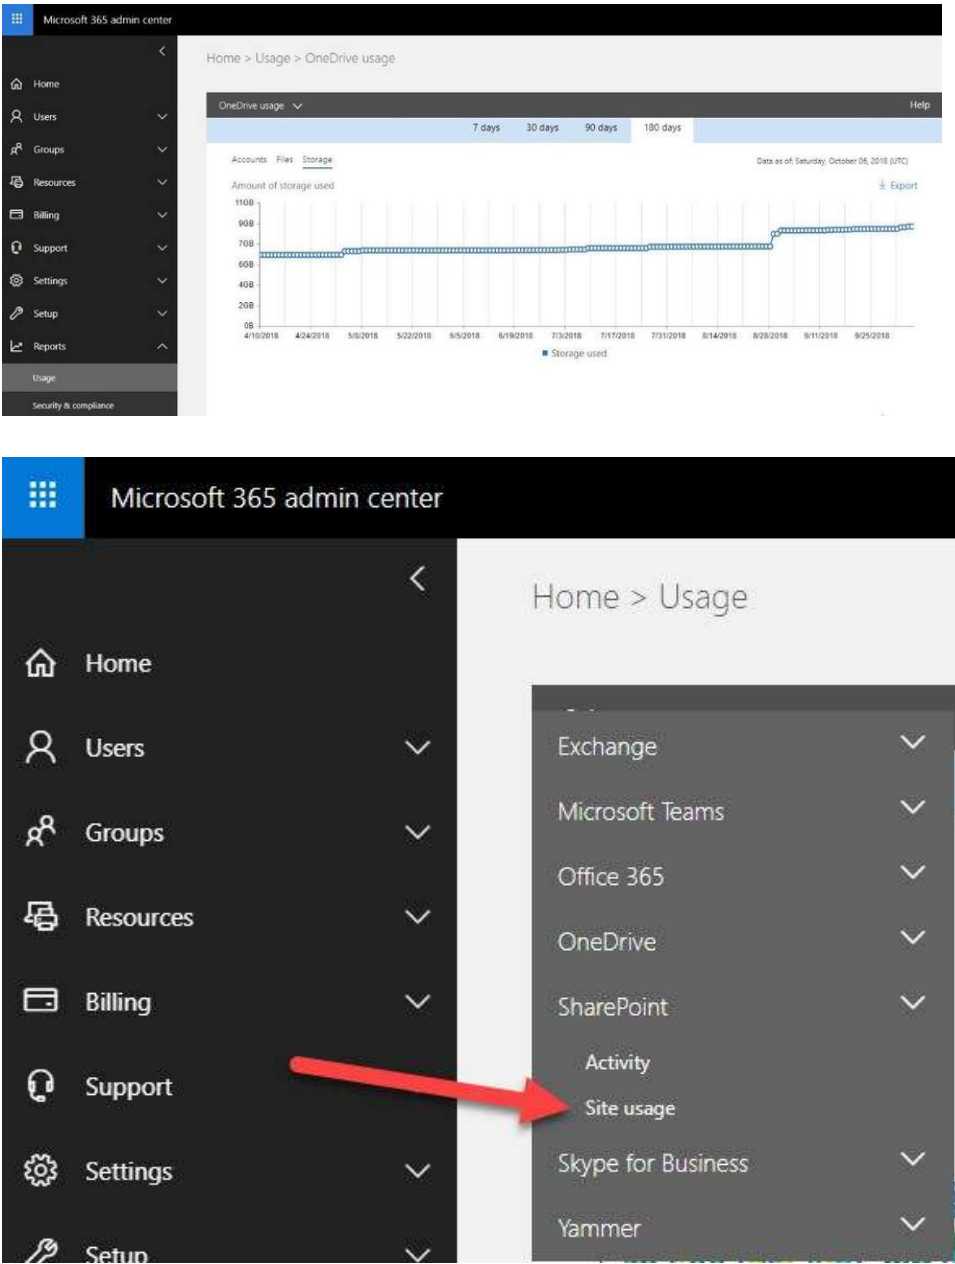
5. If you click on Storage, you will be able to check consumed space, but more importantly, you will see growth over the last 180 days. Using the growth figures, we can roughly calculate growth as a percentage. This is important for determining the space that will be

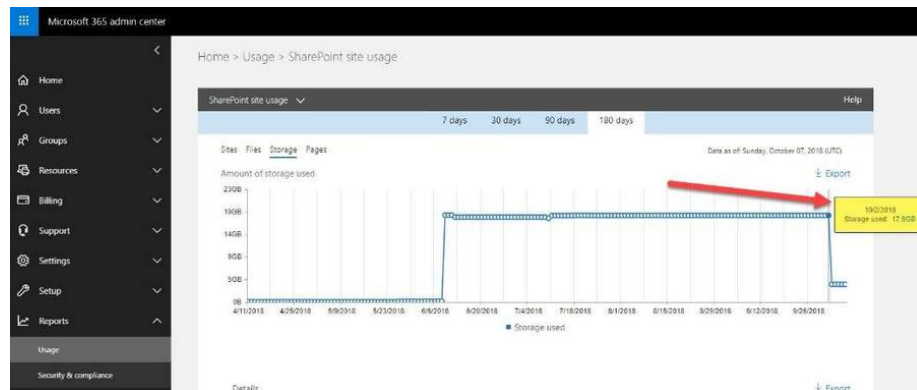
required to meet your backup retention requirements.



6. It's a similar story for OneDrive for business and SharePoint Online.







Using my examples above, I needed 36GB of space to protect all data at its peak consumption. After the first backup, all subsequent backup passes are incremental backups with Veeam backup for Office 365. Assuming 1% growth, then the next pass will be an additional 360MB of data and so on.

### Calculation

Following a recent trip to the Veeam Vanguard summit, Veeam shared their calculation for sizing backups which are:

$$(\text{current primary mailbox total size}) + ((\text{daily change rate} * 2) * (\text{days of retention})) + (10\% \text{ working area})$$

Thanks to Karl Widmer for capturing the formula in his blog post here:

<https://www.driftar.ch/index.php/2018/10/16/veeam-vanguard-summit-2018-in-prague-day-1-summary/>

## Chapter 4

# HTML Report of all Veeam Backup Copy Jobs in PowerShell

**By: Mike Conjocie (Veeam Vanguard)**

Due to the way our infrastructure operates, we have some remote sites backing up to a central Veeam repository (~300 branches at present)

We get some requests from the network team on bandwidth usage issues which is always finger pointed at the Veeam Backup Copy jobs.

I put together the following Powershell script to report on all the Backup Copy jobs currently configured and return their status.

There are some parameters which can be configured on the script:

-Server - This specifies the Veeam server to run the report against (Defaults to localhost if no parameter is specified)

-Outfile - This specifies the location in which to save the HTML report (Defaults to C:\Reports\CopyJobStatus.html if no parameter is specified)

If you're running this on a machine other than your Veeam server, you'll need to have the VeeamPSSNapin Snapin available (Installed alongside the Veeam B&R Console)

Internally, we have this running as a scheduled task every 30 minutes saving the report to a web server accessible to all teams. A screenshot of the report is available below:

## Veeam OffSite Backups Information

The following report was run on 01/12/2017 17:39:52.

JobName	Status	Progress
to Core	Stopped	100%
to Core	Stopped	100%
to Core	Stopped	100%
to Core	Stopped	100%
to Core	Stopped	100%
to Core	Running	31%
to Core	Running	8%
to Core	Running	20%
to Core	Running	24%
to Core	Running	59%
to Core	Running	15%
to Core	Running	37%
to Core	Running	20%
to Core	Running	42%
to Core	Running	20%
to Core	Running	12%
to Core	Running	32%
to Core	Running	17%
to Core	Running	1%
to Core	Running	21%
to Core	Running	74%
	Running	21%
to Core	Running	16%
to Core	Running	22%

The code in its entirety is listed below. Enjoy!

```
<#
.SYNOPSIS
    Veeam OffSite Backups Information
.DESCRIPTION
    This scripts connects to the specified Veeam server, and returns the status
    of any Backup Copy jobs (whether the are running or idle, the duration of that
    job, and the current progress), then outputs that information to a colour coded
    HTML file.
.PARAMETER Server
    Specifies the Veeam Backup Server. If this is blank, Localhost will be used.
.PARAMETER Outfile
    Specifies the path to save the HTML report. If this is blank,
    C:\Reports\CopyJobStatus.html will be used.
.NOTES
    File Name   : BackupCopyJobStatus.ps1
    Author      : Mike Conjoice - mike@mikeconjoice.com
    Requires    : Veeam Console installed if running from a remote machine
.LINK
    http://www.mikeconjoice.com
.EXAMPLE
    BackupCopyJobStatus.ps1

    This will run the script with the default parameters.
.EXAMPLE
    BackupCopyJobStatus.ps1 -Server SRV-VB01 -Outfile c:\Reports\veeam.html

    This will run the script against the Veeam server named "SRV-VB01" and
    output the HTML report to "C:\Reports\veeam.html"
#>

#####
## TODO: Add Job Duration
## TODO: Send STOPPED Jobs to the Bottom
## -----
## DONE: Change $outfile and $server to Parameters
#####

#####
## Parameters
#####

param(
    [String]$outfile="C:\Reports\CopyJobStatus.html",
    [String]$server="localhost"
)

#####
## Load required Snapins and Modules
#####

if ((Get-PSSnapin -Name VeeamPSSnapin -ErrorAction SilentlyContinue) -eq $null)
{
    Add-PSSnapin VeeamPSSnapin
}

70
```

```
#####
## Set the CSS for the output file
#####

$head = @"
<style>
th {background-color: #00aff0}
table {border-collapse: collapse}
table, th, td {
border: 1px solid black;
padding: 5px
}
th {color: white}
body {font-family: sans-serif}
</style>
"@

#####
## Create an empty array
#####

$results = @()

#####
## Begin the script by collecting a list of all the Offsite backup jobs
#####

if ($Server -eq $null) {
    Connect-VBRServer -Server $server
} else {
    Disconnect-VBRServer
    Connect-VBRServer -Server $server
}
$JobNames = Get-VBRJob | where-Object {$_.JobType -Like "*Sync"}

#####
## Loop through all the returned jobs to find the Job Name, Current Status,
Progress, and Duration to be entered in to the array
#####

foreach ($JobName in $JobNames) {
    $Job = Get-VBRJob -name $JobName.Name
    $LastSession = $Job.FindLastSession()
    $Name = $Job.Name
    $Status = $LastSession.State
    $Progress = "$($LastSession.BaseProgress)%"

    #####
    ## Create a new PSObject and populate the array with the details
    #####

    $results += New-Object PSObject -Property @{JobName = $Name; Status =
$status; Progress = $Progress;}
}

#####
## Collate, sort, colourise, and output the results to an HTML file
#####
```

```
$results |
Select JobName, Status, Progress |
sort
@{expression="Status";Descending=$false},@{expression="JobName";Ascending=$true}
|
ConvertTo-Html -body "<H2>Veeam OffSite Backups Information</H2> <p>The
following report was run on $(get-date).</p>" -PreContent $head -Title "Veeam
OffSite Backups Information" |
    foreach {
        $PSItem -replace "<td>working</td>", "<td style='background-color:green;
color: white'>Running</td>" -replace "<td>Idle</td>", "<td style='background-
color:orange; color: white'>Idle</td>" -replace "<td>Stopped</td>", "<td
style='background-color:red; color: white'>Stopped</td>"
    } | Out-File $outfile

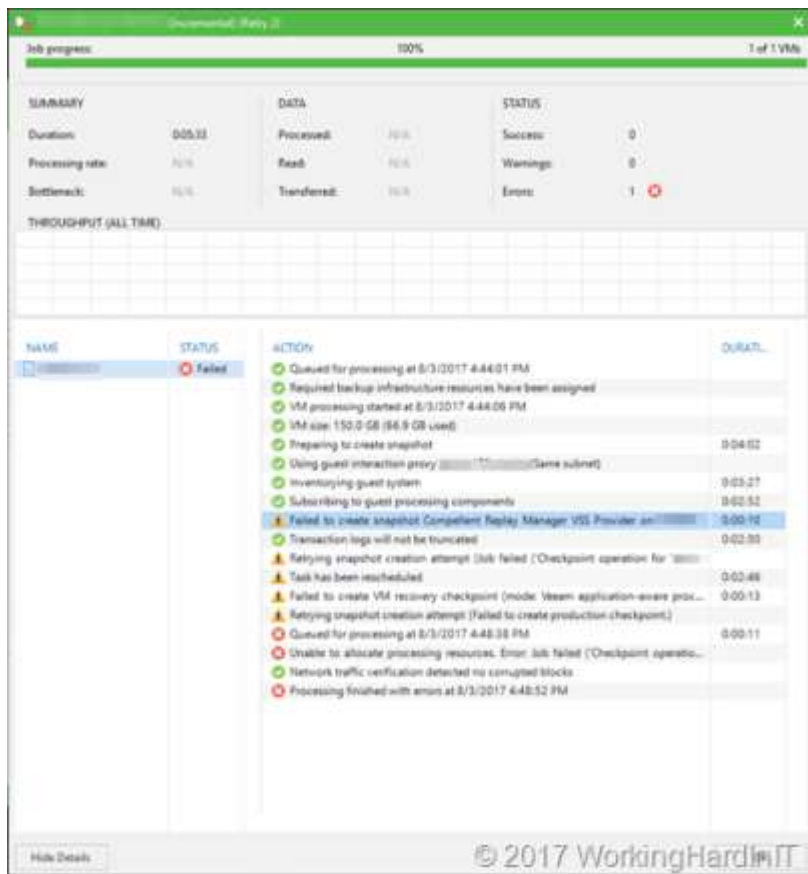
#####
## End of script
#####
```

## Chapter 5

# Troubleshooting Veeam B&R Error Code: 32768 Failed to create VM Recovery Snapshot

**By: Didier Van Hoya (Microsoft MVP / Veeam Vanguard)**

I recently had to move a Windows Server 2016 VM over to another cluster (2012R2 to 2016 cluster) and to do so I use shared-nothing live migration. After the VM was happily running on the new cluster, I kicked off a Veeam backup job to get a first restore point for that VM. Better safe than sorry right?

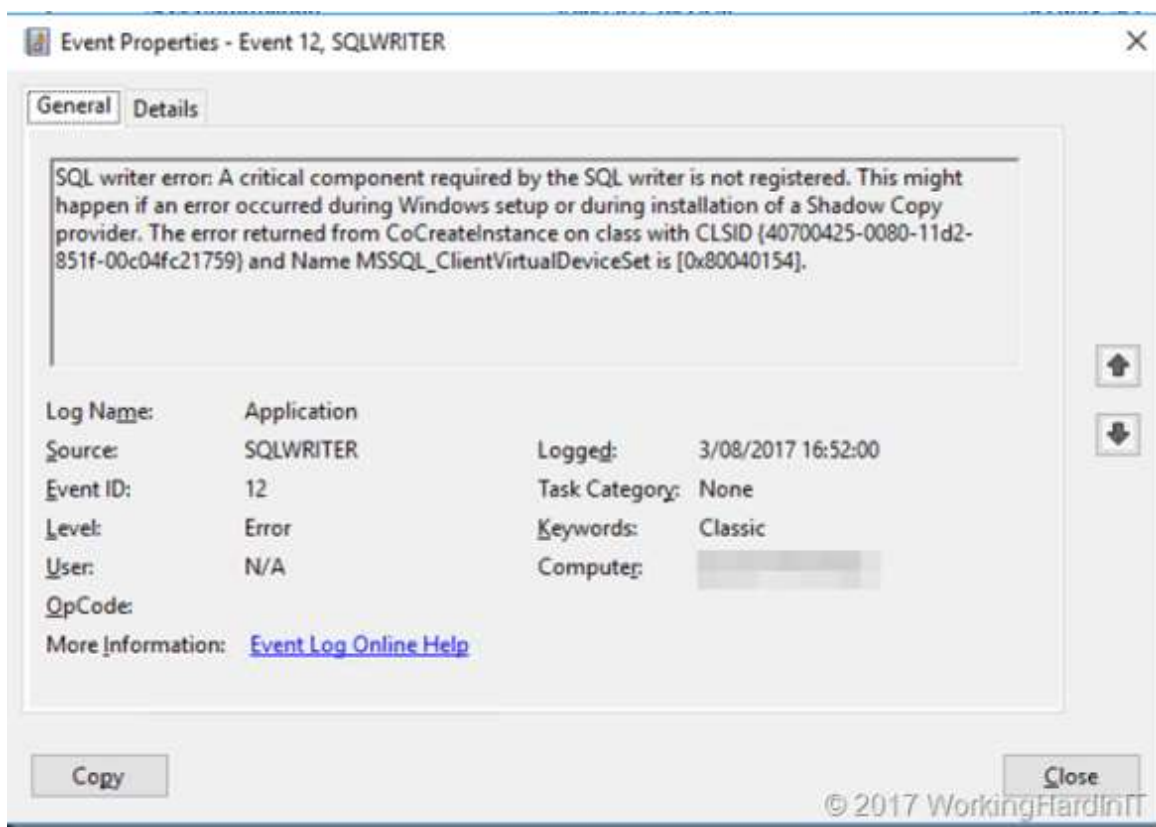


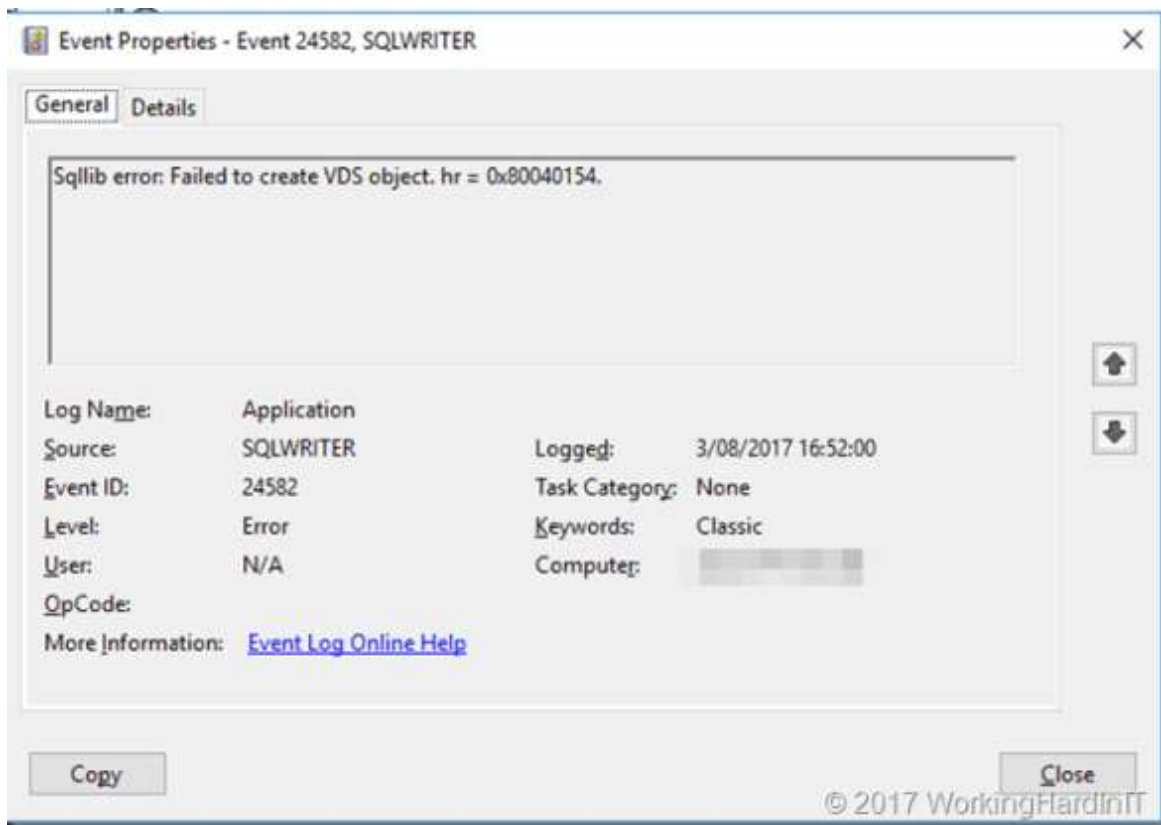
However, the job and the retries failed for that VM. The error details are:

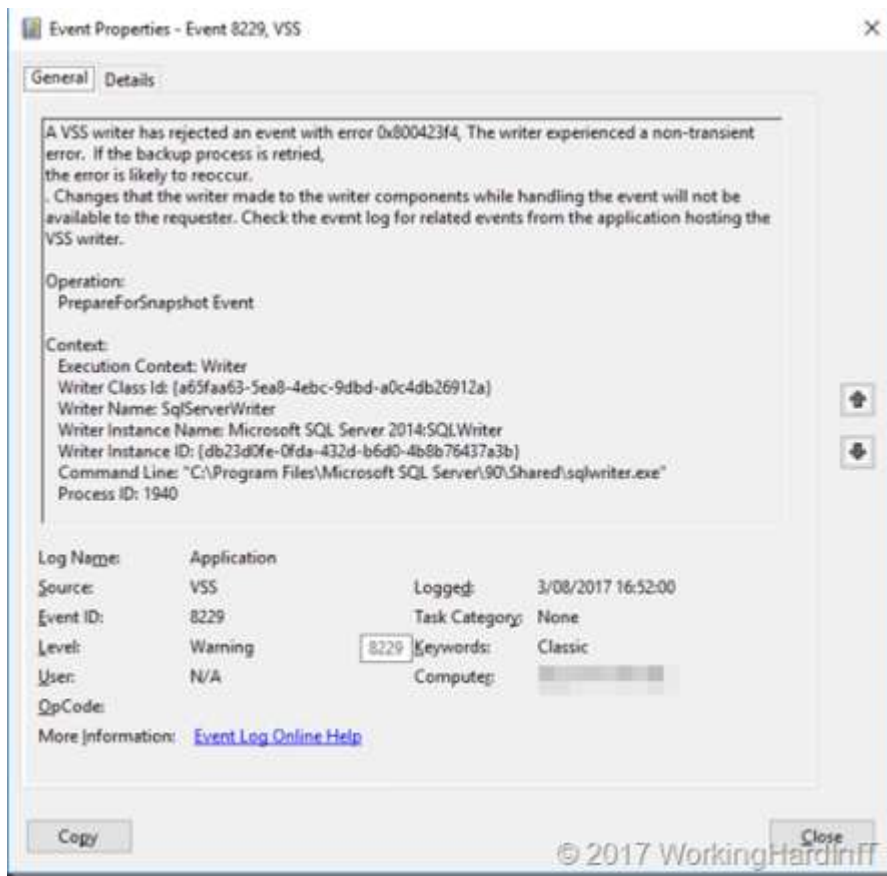
Failed to create snapshot Compellent Replay Manager VSS Provider on repository01.domain.com (mode: Veeam application-aware processing) Details: Job failed ('Checkpoint operation for 'FailedVM' failed. (Virtual machine ID 459C3068-9ED4-427B-AAEF-32A329B953AD). 'FailedVM' could not initiate a checkpoint operation: %2147754996 (0x800423F4). (Virtual machine ID 459C3068-9ED4-427B-AAEF-32A329B953AD)'). Error code: '32768'.

Failed to create VM recovery snapshot, VM ID '3459c3068-9ed4-427b-aaef-32a329b953ad'.

Also when the job fails over to the native Windows VSS approach when the HW VSS provider fails it still does not work. At first, that made me think of a bug that used to exist in Windows Server 2016 Hyper-V where a live storage migration of any kind would break RCT, and new full was needed to fix it. That bug has long since been fixed, and now a new full backup did not solve anything here. Now there are various reasons why creating a checkpoint will not succeed so we need to dive in deeper. As always the event viewer is your friend. What do we see? 3 events during a backup and they are SQL Server related.







On top of that, the SqlServerWriter is in a non-retryable error when checking with vssadmin list writers.

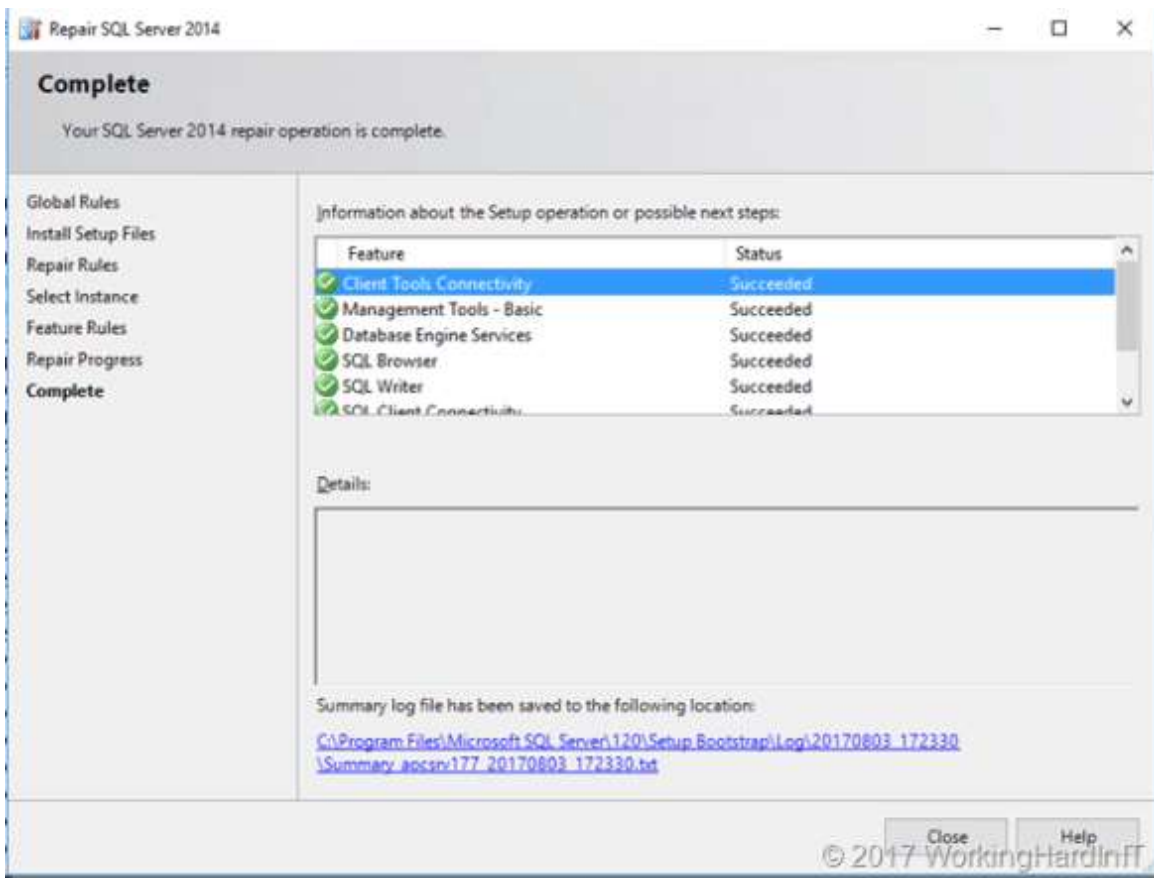
```
Writer name: 'SqlServerWriter'
Writer Id: {a65faa63-5ea8-4ebc-9dbd-a0c4db26912a}
Writer Instance Id: {db23d0fe-0fda-432d-b6d0-4b8b76437a3b}
State: [8] Failed
Last error: Non-retryable error
```

© 2017 WorkingHardInIT

It's very clear there is an issue with the SQL Server VSS Writer in this VM, and that cause the checkpoint to fail. You can search for manual fixes, but in the case of an otherwise functional SQL

Server, I chose to go for a repair install of SQL Server. The tooling for that is pretty good, and it's probably the fastest way to resolve the issues and any underlying ones we might otherwise still encounter.

After running a successful repair install of SQL Server, we get greeted by an all green result screen.



Recheck vssadmin list writers to make sure they are all healthy if not restart the SQL s or another relevant service if possible. Sometimes you can fix it by restarting a service, in that case, reboot the server. We did not need to do that. We just ran a new retry in Veeam Backup & Replication and were successful.

There you go. The storage live migration before the backup of that VM made me think we were dealing with an early Windows Server 2016 Hyper-V bug but that was not the case. Troubleshooting is also about avoiding tunnel vision.

## Chapter 6

# Install Veeam Backup and Replication using PowerShell

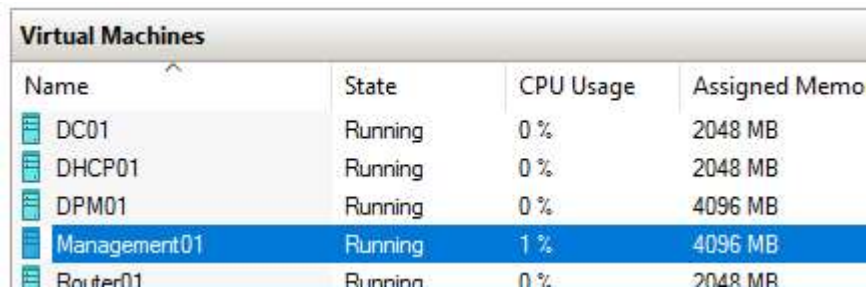
By:

Dave Kawula (Microsoft MVP / Veeam Vanguard)

Cristal Kawula (Microsoft MVP / Veeam Vanguard)

In this chapter, we will show you how you can deploy Veeam Backup and Replication 9.5 with Update Rollup 3 to a Hyper-V Virtual Machine. This is a convenient little script because it can help you get Veeam going in your lab in a snap. The necessary steps to accomplish this are as follows:

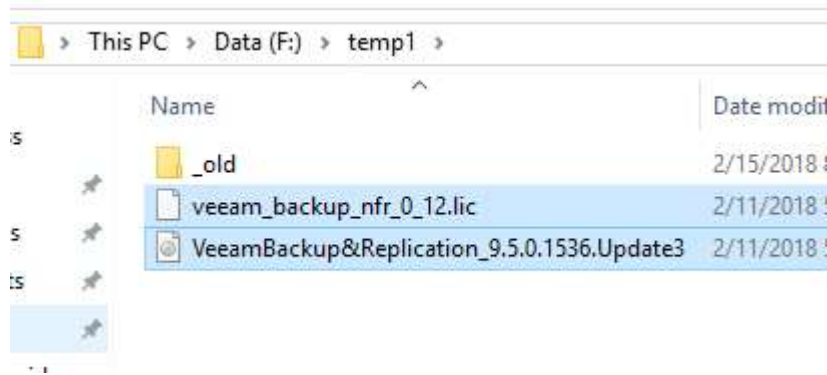
1. Create a Hyper-V Virtual Machine running Windows Server 2016. In my case, I have created a VM called Management01.



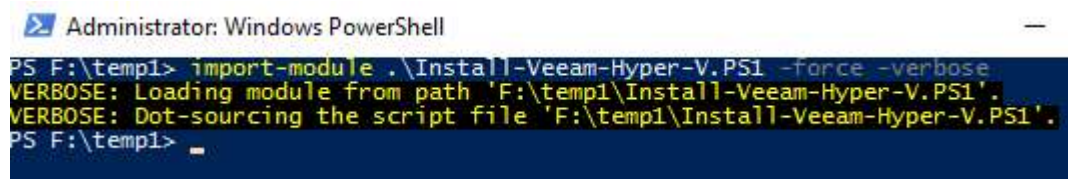
Virtual Machines			
Name	State	CPU Usage	Assigned Memo
DC01	Running	0 %	2048 MB
DHCP01	Running	0 %	2048 MB
DPM01	Running	0 %	4096 MB
Management01	Running	1 %	4096 MB
Router01	Running	0 %	2048 MB

2. Next, you will need to download a copy of Veeam Backup and Replication from <https://www.veeam.com/downloads.html>

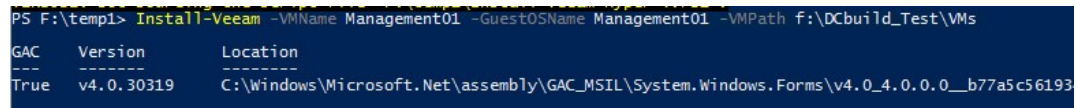
3. Next, you need to request a trial product key <VEEAM Can I have the proper link for this>
4. In our case, we have stored both files in F:\Temp1



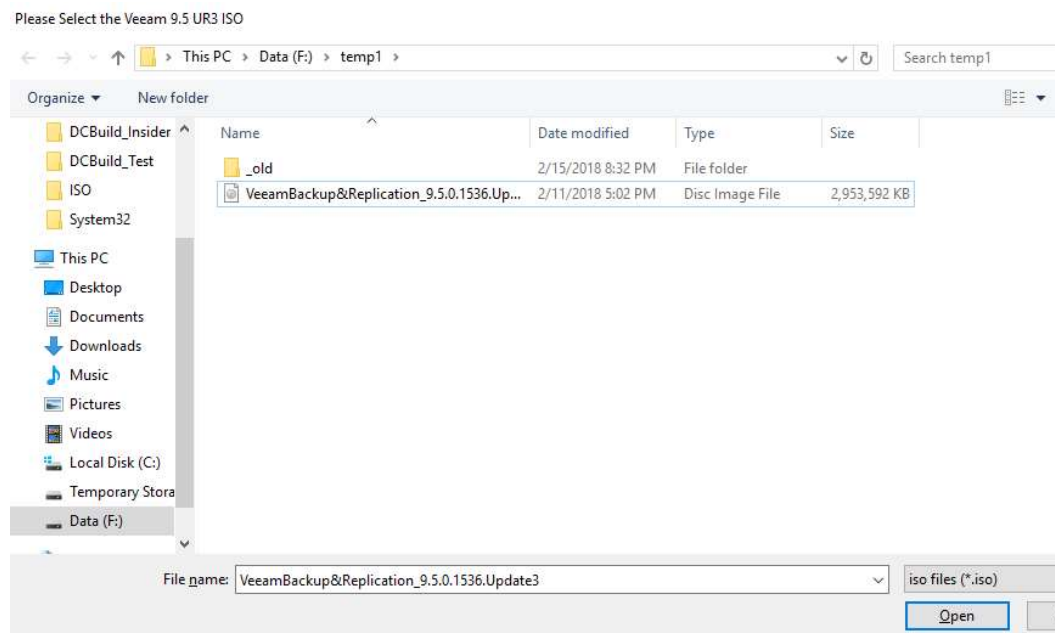
5. Next, grab a copy of the PowerShell script from my GitHub Repository.  
<https://raw.githubusercontent.com/dkawula/Operations/master/Veeam/Install-Veeam-HyperV.ps1> and save it to f:\Temp1
6. Open an Administrative PowerShell Prompt and change the path to f:\temp1. Then type `.\Import-Module .\Install-Veeam-Hyper-V.PS1 -Force -Verbose` and press Enter. This will load the Install-Veeam Module which we will use to deploy Veeam to our lab.



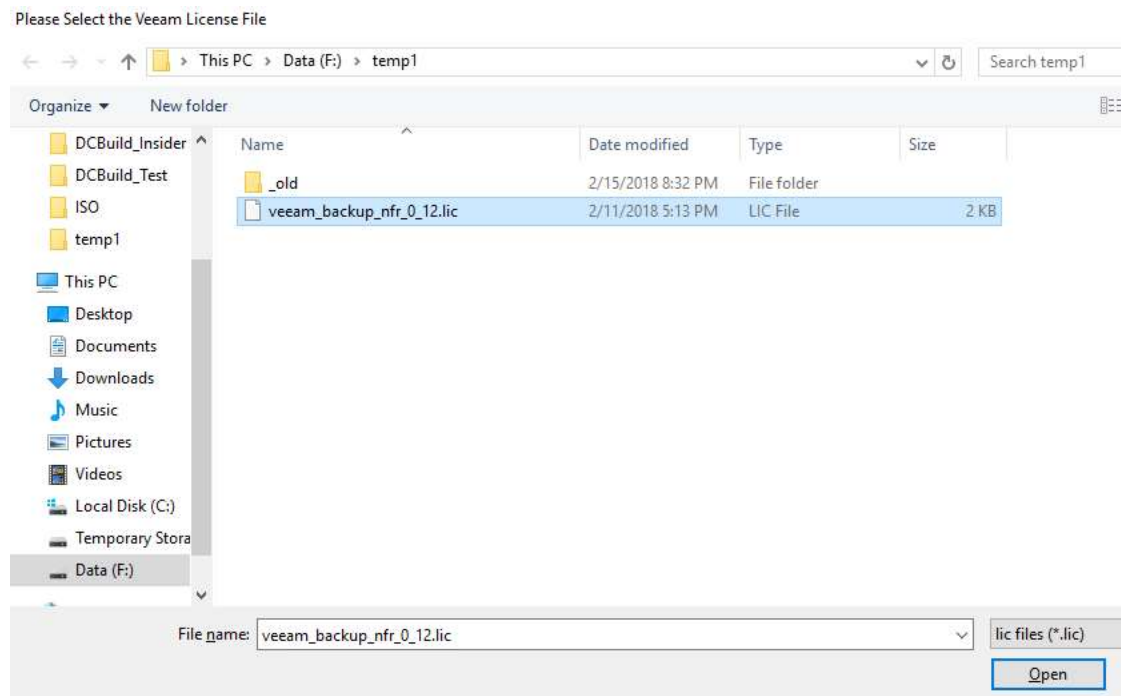
7. Next type `Install-Veeam -VMName Management01 -GuestOSName Management01 -VMPATH f:\dcbuild_test` and press enter.



8. On the Please Select, the Veeam 9.5 UR3 ISO window browse to f:\temp1\VeeamBackup&Replication\_9.5.0.1535.Update3.Iso



9. On the Please Select, the Veeam License File window browse to f:\temp1\Veeam\_Backup\_nfr\_0\_12.lic. In my case, I used my Veeam Vanguard Veeam NFR License.



10. Next input the credentials for your lab. In my case, I will install using my Domain Administrator Account MVPDays\Administrator. We will pipe this into the VM using PowerShell direct with the \$DomainCred variable.



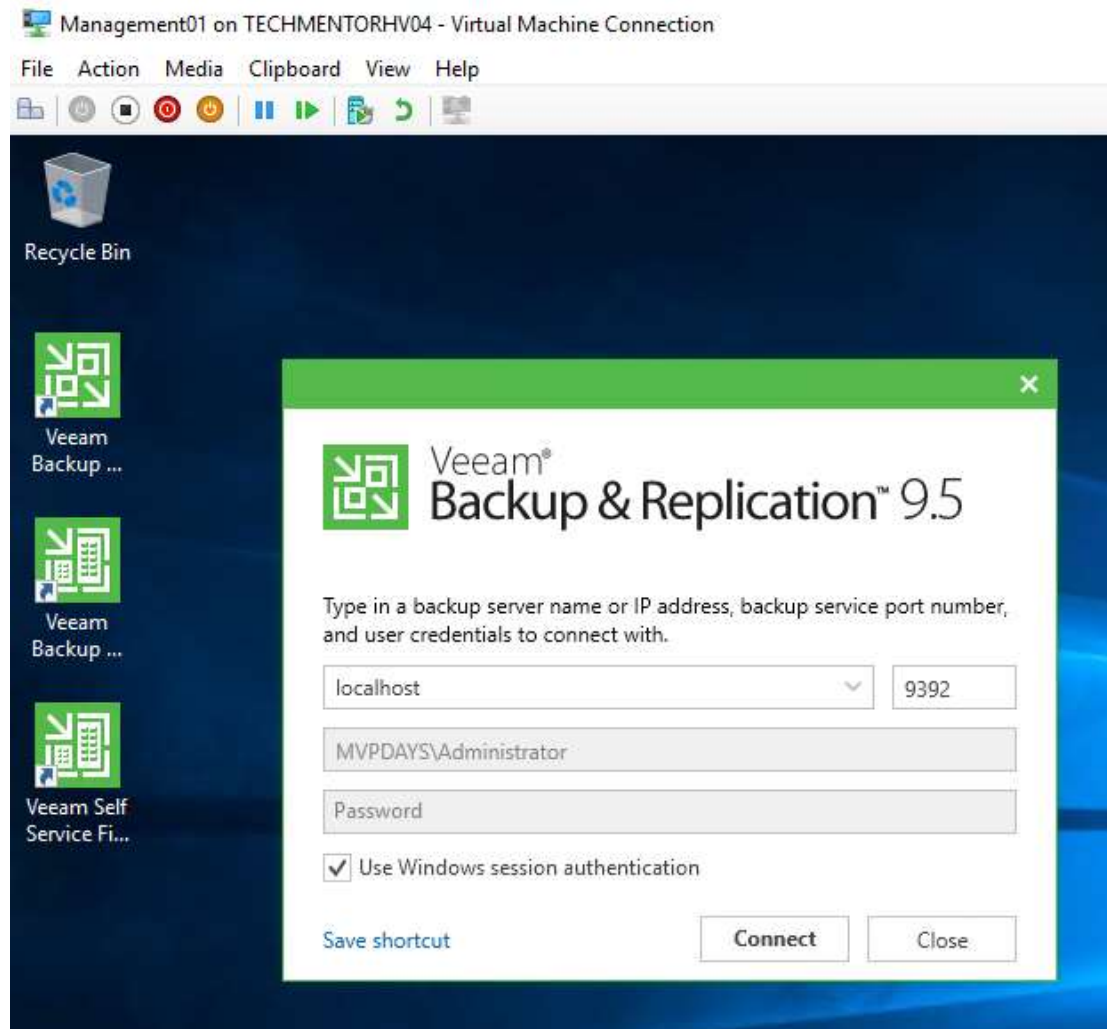
At this point it is time to sit back and relax as the next steps will happen automatically:

- Adding a New VHDx file from the Host
- Formatting the New VHDx file and giving it a volume label of Veeam
- Mounting the VHDx to the Host
- Copying the Veeam ISO specified earlier
- Copying the Veeam License file specified earlier
- Installs 2012 System CLR Types
- Installs SQL Express
- Installs the Veeam Backup and Replication 9.5 Server
- Installs the Veeam Backup Console
- Installs the Explorer for Active Directory
- Installs the Explorer for Exchange

- Installs the Explorer for SQL
- Installs the Explorer for Oracle
- Installs the Explorer for SharePoint
- Installs the Enterprise Manager web and Cloud Portal
- Finally updates to Rollup 3

```
[ ]:: Adding the new VHDx for the Veeam Install
E:
[ ]:: Mounting Veeam ISO
[ ]:: Installing Veeam Unattended
D:
Installing Global Prerequisites ...
  Installing 2012 System CLR Types ...
  Setup OK
  Installing 2012 Shared management objects ...
  Setup OK
  Installing SQL Express ...
Installing Veeam Backup & Replication ...
  Installing Backup Catalog ...
  Setup OK
  Installing Backup Server ...
  Setup OK
  Installing Backup Console ...
  Setup OK
  Installing Explorer For ActiveDirectory ...
  Setup OK
  Installing Explorer For Exchange ...
  Setup OK
  Installing Explorer For SQL ...
  Setup OK
  Installing Explorer For Oracle ...
  Setup OK
  Installing Explorer For SharePoint ...
  Setup OK
Installing Enterprise Manager ...
  Installing Enterprise Manager Prerequisites ...
  Setup OK
  Installing Enterprise Manager Web ...
  Setup OK
  Installing Enterprise Manager Cloud Portal ...
  Setup OK
Installing Update 3 ...
```

11. Your finished product will look just like this:



For your reference here is a copy of the script:

```
<#
Created:      2018-02-01
Version:      1.0
Author       Dave Kawula MVP
Homepage:     http://www.checkyourlogs.net

Disclaimer:
This script is provided "AS IS" with no warranties, confers no rights and
is not supported by the authors or CheckyourLogs or MVPDays Publishing

Author - Dave Kawula
Twitter: @DaveKawula
Blog    : http://www.checkyourlogs.net

.Synopsis
Deploys Veeam Backup and Replication 9.5 + UR3 to a Hyper-V Lab VM
.DESCRIPTION
This Script was part of my BIGDemo series, and I have broken it out into a
standalone function

You will need to have a Veeam Service Account Pre-Created, Veeam B&R ISO and
Product Key for this lab to work
The Script will prompt for the path of the ISO and .LIC files
The Script will prompt for an Admin Account which will be used in
$DomainCred
If your File names are different than mine adjust accordingly.

We will use PowerShell Direct to setup the Veeam Server in Hyper-V

The Source Hyper-V Virtual Machine needs to be Windows Server 2016

.EXAMPLE
TODO: Dave, add something more meaningful in here
.PARAMETER WorkingDir
Transactional directory for files to be staged and written
.PARAMETER VMname
The name of the Virtual Machine
.PARAMETER VMPath
The Path to the VM Working Folder - We create a new VHDx for the Veeam
Install
.PARAMETER GuestOSName
Name of the Guest Operating System Name

Usage: Install-Veeam -Vmname YOURVM -GuestOS VEEAMSERVER -VMpath
f:\VMS\Veeam -workingDir f:\Temp
#>
#Installs Veeam 9.5 and UR 3

Function Install-Veeam {
param
(
    [string] $VMName,
    [string] $GuestOSName,
    [string] $VMPath,
    [string] $WorkingDir
)
```

```

#Ask for Veeam ISO

[reflection.assembly]::loadwithpartialname("System.Windows.Forms")
$openFile = New-Object System.Windows.Forms.OpenFileDialog -Property @{
    Title="Please Select the Veeam 9.5 UR3 ISO"
}
$openFile.Filter = "iso files (*.iso)|*.iso|All files (*.*)|*.*"
If($openFile.ShowDialog() -eq "OK")
{
    Write-Host "File $($openfile.FileName) selected"
}
if (!$openFile.FileName){
    WriteErrorAndExit "Iso was not selected... Exiting"
}
$VeeamISO = $openfile.FileName
#$VeeamISO
#Ask for Veeam License

[reflection.assembly]::loadwithpartialname("System.Windows.Forms")
$openFile = New-Object System.Windows.Forms.OpenFileDialog -Property @{
    Title="Please Select the Veeam License File"
}
$openFile.Filter = "lic files (*.lic)|*.lic|All files (*.*)|*.*"
If($openFile.ShowDialog() -eq "OK")
{
    Write-Host "File $($openfile.FileName) selected"
}
if (!$openFile.FileName){
    WriteErrorAndExit "Iso was not selected... Exiting"
}
$VeeamLic = $openfile.FileName
#$VeeamLic

$DomainCred = Get-Credential
#$VMName = 'Management01'
#$GuestOSName = 'Management01'
#$VMPath = 'f:\dcbuild_Test\VMs'

Write-Output -InputObject "[$($VMName)]:: Adding Drive for Veeam Install"

New-VHD -Path "$($VMPath)\$($GuestOSName) - Veeam Data 4.vhdx" -Dynamic -
SizeBytes 60GB
Mount-VHD -Path "$($VMPath)\$($GuestOSName) - Veeam Data 4.vhdx"
$DiskNumber = (Get-Diskimage -ImagePath "$($VMPath)\$($GuestOSName) - Veeam
Data 4.vhdx").Number
Initialize-Disk -Number $DiskNumber -PartitionStyle GPT
Get-Disk -Number $DiskNumber | New-Partition -UseMaximumSize -
AssignDriveLetter | Format-Volume -FileSystem NTFS -NewFileSystemLabel "Veeam" -
Confirm:$False
$DriveLetter = get-wmiobject -class "win32_volume" -namespace "root\cimv2" |
where-object {$_.Label -like "Veeam*"}
$VeeamDriveLetter = $DriveLetter.DriveLetter

Write-Output -InputObject "[$($VMName)]:: Copying Veeam ISO and Rollups into
the new VHDx"

```

```
Copy-Item -Path $VeeamIso -Destination
"$($VeeamDriveLetter)\VeeamBackup&Replication_9.5.0.1536.Update3.iso" -Force
write-Output -InputObject "[$($VMName)]:: Copying Veeam license and Rollups
into the new VHDx"
Copy-Item -Path $VeeamLic -Destination
"$($VeeamDriveLetter)\veeam_backup_nfr_0_12.lic" -Force
Dismount-VHD -Path "$($VMPath)\$($GuestOSName) - Veeam Data 3.vhdx"
Add-VMHardDiskDrive -VMName $VMName -Path "$($VMPath)\$($GuestOSName) -
Veeam Data 4.vhdx" -ControllerType SCSI

icm -VMName $VMName -Credential $domainCred {

    write-Output -InputObject "[$($VMName)]:: Adding the new VHDx for the Veeam
Install"
    Get-Disk | Where OperationalStatus -EQ "Offline" | Set-Disk -IsOffline
$False
    Get-Disk | Where Number -NE "0" | Set-Disk -IsReadOnly $False
    $DriveLetter = get-wmiobject -class "win32_volume" -namespace "root\cimv2" |
where-object {$_.Label -like "Veeam*"}
    $VeeamDrive = $DriveLetter.DriveLetter
    $VeeamDrive

    write-Output -InputObject "[$($VMName)]:: Mounting Veeam ISO"

    $iso = Get-ChildItem -Path
"$($VeeamDrive)\VeeamBackup&Replication_9.5.0.1536.Update3.iso" #CHANGE THIS!

    Mount-DiskImage $iso.FullName

    write-Output -InputObject "[$($VMName)]:: Installing Veeam Unattended"

    $setup = $(Get-DiskImage -ImagePath $iso.FullName | Get-
Volume).DriveLetter + ':'
    $setup

    <#>

=====

    Original Source Created by: Markus Kraus

    Twitter: @VMarkus_K

    Private Blog: mycloudrevolution.com
    #Source PowerShell Code from
https://gist.github.com/usercontent/mycloudrevolution/b176f5ab987ff787ba4fce5c177780dc/raw/f20a78dc9b7c1085b1fe4d243de3fcb514970d70/VeeamBR95-Silent.ps1

    </#>

    # Requires PowerShell 5.1
    # Requires .Net 4.5.2 and Reboot
```

```

#region: Variables
$source = $setup
$licensefile = "$($VeeamDrive)\veeam_backup_nfr_0_12.lic"
$username = "svc_veeam"
$fulluser = "MVPDays\svc_Veeam"
$password = "P@ssw0rd"
$CatalogPath = "$($VeeamDrive)\vbrCatalog"
$vPowerPath = "$($VeeamDrive)\vPowerNfs"
#endregion

#region: logdir
$logdir = "$($VeeamDrive)\logdir"
$trash = New-Item -ItemType Directory -path $logdir -ErrorAction
silentlyContinue
#endregion

### Optional .Net 4.5.2
<#
Write-Host "    Installing .Net 4.5.2 ..." -ForegroundColor Yellow
$Arguments = "/quiet /norestart"
Start-Process "$source\Redistr\NDP452-KB2901907-x86-x64-AllOS-ENU.exe" -
ArgumentList $Arguments -Wait -NoNewWindow
Restart-Computer -Confirm:$true
#>

### Optional PowerShell 5.1
<#
Write-Host "    Installing PowerShell 5.1 ..." -ForegroundColor Yellow
$Arguments = "C:\_install\win8.1Andw2K12R2-KB3191564-x64.msu /quiet
/norestart"
Start-Process "wusa.exe" -ArgumentList $Arguments -Wait -NoNewWindow
Restart-Computer -Confirm:$true
#>

#region: Installation
# Info: https://www.veeam.com/unattended_installation_ds.pdf

## Global Prerequisites
Write-Host "Installing Global Prerequisites ..." -ForegroundColor
yellow

### 2012 System CLR Types
Write-Host "    Installing 2012 System CLR Types ..." -ForegroundColor
yellow

$MSIArguments = @(
    "/i"
    "$source\Redistr\x64\SQLSysClrTypes.msi"
    "/qn"
    "/norestart"
    "/L*v"
    "$logdir\01_CLR.txt"
)
Start-Process "msiexec.exe" -ArgumentList $MSIArguments -Wait -
NoNewWindow

if (Select-String -path "$logdir\01_CLR.txt" -pattern "Installation
success or error status: 0.") {
    Write-Host "    Setup OK" -ForegroundColor Green
}
else {

```

```

        throw "Setup Failed"
    }

    ### 2012 Shared management objects
    Write-Host "    Installing 2012 Shared management objects ..." -
ForegroundColors Yellow
    $MSIArguments = @(
        "/i"
        "$source\Redistr\x64\SharedManagementObjects.msi"
        "/qn"
        "/norestart"
        "/L*v"
        "$logdir\02_Shared.txt"
    )
    Start-Process "msiexec.exe" -ArgumentList $MSIArguments -wait -
NoNewWindow

    if (Select-String -path "$logdir\02_Shared.txt" -pattern "Installation
success or error status: 0.") {
        Write-Host "    Setup OK" -ForegroundColor Green
    }
    else {
        throw "Setup Failed"
    }

    ### SQL Express
    ### Info: https://msdn.microsoft.com/en-us/library/ms144259.aspx
    Write-Host "    Installing SQL Express ..." -ForegroundColor Yellow
    $Arguments = "/HIDECONSOLE /Q /IACCEPTSQLSERVERLICENSETERMS
/ACTION=install /FEATURES=SQLEngine,SNAC_SDK /INSTANCENAME=VEEAMSQL2012
/SQLSVCAccount=`"NT AUTHORITY\SYSTEM`" /SQLSYSADMINACCOUNTS=`"$fulluser`"
"Builtin\Administrators`" /TCPENABLED=1 /NPENABLED=1 /UpdateEnabled=0"
    Start-Process "$source\Redistr\x64\SQLEXPR_x64_ENU.exe" -ArgumentList
$Arguments -Wait -NoNewWindow

    ## Veeam Backup & Replication
    Write-Host "Installing Veeam Backup & Replication ..." -ForegroundColor
Yellow
    ### Backup Catalog
    Write-Host "    Installing Backup Catalog ..." -ForegroundColor Yellow
    $trash = New-Item -ItemType Directory -path $CatalogPath -ErrorAction
SilentlyContinue
    $MSIArguments = @(
        "/i"
        "$source\Catalog\VeeamBackupCatalog64.msi"
        "/qn"
        "/L*v"
        "$logdir\04_Catalog.txt"
        "VM_CATALOGPATH=$CatalogPath"
        "VBRC_SERVICE_USER=$fulluser"
        "VBRC_SERVICE_PASSWORD=$password"
    )
    Start-Process "msiexec.exe" -ArgumentList $MSIArguments -wait -
NoNewWindow

    if (Select-String -path "$logdir\04_Catalog.txt" -pattern "Installation
success or error status: 0.") {
        Write-Host "    Setup OK" -ForegroundColor Green
    }
}

```

```

        else {
            throw "Setup Failed"
        }

    ### Backup Server
    Write-Host "    Installing Backup Server ..." -ForegroundColor Yellow
    $trash = New-Item -ItemType Directory -path $vPowerPath -ErrorAction
silentlyContinue
    $MSIArguments = @(
        "/i"
        "$source\Backup\Server.x64.msi"
        "/qn"
        "/L*v"
        "$logdir\05_Backup.txt"
        "ACCEPTTEULA=YES"
        "VBR_LICENSE_FILE=$licensefile"
        "VBR_SERVICE_USER=$fulluser"
        "VBR_SERVICE_PASSWORD=$password"
        "PF_AD_NFSDATASTORE=$vPowerPath"
        "VBR_SQLSERVER_SERVER=$env:COMPUTERNAME\VEEAMSQL2012"
    )
    Start-Process "msiexec.exe" -ArgumentList $MSIArguments -Wait -
NoNewWindow

    if (Select-String -path "$logdir\05_Backup.txt" -pattern "Installation
success or error status: 0.") {
        Write-Host "    Setup OK" -ForegroundColor Green
    }
    else {
        throw "Setup Failed"
    }

    ### Backup Console
    Write-Host "    Installing Backup Console ..." -ForegroundColor Yellow
    $MSIArguments = @(
        "/i"
        "$source\Backup\Shell.x64.msi"
        "/qn"
        "/L*v"
        "$logdir\06_Console.txt"
        "ACCEPTTEULA=YES"
    )
    Start-Process "msiexec.exe" -ArgumentList $MSIArguments -Wait -
NoNewWindow

    if (Select-String -path "$logdir\06_Console.txt" -pattern "Installation
success or error status: 0.") {
        Write-Host "    Setup OK" -ForegroundColor Green
    }
    else {
        throw "Setup Failed"
    }

    ### Explorers
    Write-Host "    Installing Explorer For ActiveDirectory ..." -
ForegroundColor Yellow
    $MSIArguments = @(
        "/i"
        "$source\Explorers\VeeamExplorerForActiveDirectory.msi"
    )

```

```

        "/qn"
        "/L*v"
        "$logdir\07_ExplorerForActiveDirectory.txt"
        "ACCEPTTEULA=YES"
    )
    Start-Process "msiexec.exe" -ArgumentList $MSIArguments -Wait -
NoNewWindow
    if (Select-String -path "$logdir\07_ExplorerForActiveDirectory.txt" -
pattern "Installation success or error status: 0.") {
        Write-Host "    Setup OK" -ForegroundColor Green
    }
    else {
        throw "Setup Failed"
    }

yellow Write-Host "    Installing Explorer For Exchange ..." -ForegroundColor
$MSIArguments = @(
    "/i"
    "$source\Explorers\VeeamExplorerForExchange.msi"
    "/qn"
    "/L*v"
    "$logdir\08_VeeamExplorerForExchange.txt"
    "ACCEPTTEULA=YES"
)
    Start-Process "msiexec.exe" -ArgumentList $MSIArguments -Wait -
NoNewWindow
    if (Select-String -path "$logdir\08_VeeamExplorerForExchange.txt" -
pattern "Installation success or error status: 0.") {
        Write-Host "    Setup OK" -ForegroundColor Green
    }
    else {
        throw "Setup Failed"
    }

    Write-Host "    Installing Explorer For SQL ..." -ForegroundColor Yellow
    $MSIArguments = @(
        "/i"
        "$source\Explorers\VeeamExplorerForSQL.msi"
        "/qn"
        "/L*v"
        "$logdir\09_VeeamExplorerForSQL.txt"
        "ACCEPTTEULA=YES"
    )
    Start-Process "msiexec.exe" -ArgumentList $MSIArguments -Wait -
NoNewWindow
    if (Select-String -path "$logdir\09_VeeamExplorerForSQL.txt" -pattern
"Installation success or error status: 0.") {
        Write-Host "    Setup OK" -ForegroundColor Green
    }
    else {
        throw "Setup Failed"
    }

yellow Write-Host "    Installing Explorer For Oracle ..." -ForegroundColor

```

```

$MSIArguments = @(
    "/i"
    "$source\Explorers\VeeamExplorerForOracle.msi"
    "/qn"
    "/L*v"
    "$logdir\10_VeeamExplorerForOracle.txt"
    "ACCEPTTEULA=YES"
)
Start-Process "msiexec.exe" -ArgumentList $MSIArguments -Wait -
NoNewWindow

if (Select-String -path "$logdir\10_VeeamExplorerForOracle.txt" -pattern
"Installation success or error status: 0.") {
    Write-Host "    Setup OK" -ForegroundColor Green
}
else {
    throw "Setup Failed"
}

Write-Host "    Installing Explorer For SharePoint ..." -ForegroundColor
Yellow
$MSIArguments = @(
    "/i"
    "$source\Explorers\VeeamExplorerForSharePoint.msi"
    "/qn"
    "/L*v"
    "$logdir\11_VeeamExplorerForSharePoint.txt"
    "ACCEPTTEULA=YES"
)
Start-Process "msiexec.exe" -ArgumentList $MSIArguments -Wait -
NoNewWindow

if (Select-String -path "$logdir\11_VeeamExplorerForSharePoint.txt" -
pattern "Installation success or error status: 0.") {
    Write-Host "    Setup OK" -ForegroundColor Green
}
else {
    throw "Setup Failed"
}

## Enterprise Manager
Write-Host "Installing Enterprise Manager ..." -ForegroundColor Yellow
### Enterprise Manager Prerequisites
Write-Host "    Installing Enterprise Manager Prerequisites ..." -
ForegroundColor Yellow
$trash = Install-WindowsFeature Web-Default-Doc,Web-Dir-Browsing,Web-
Http-Errors,Web-Static-Content,Web-Windows-Auth -Restart:$false -WarningAction
SilentlyContinue
$trash = Install-WindowsFeature Web-Http-Logging,Web-Stat-
Compression,Web-Filtering,Web-Net-Ext45,Web-Asp-Net45,Web-ISAPI-Ext,Web-ISAPI-
Filter,Web-Mgmt-Console -Restart:$false -WarningAction SilentlyContinue

$MSIArguments = @(
    "/i"
    "$source\Redistr\x64\rewrite_amd64.msi"
    "/qn"
    "/norestart"
    "/L*v"
    "$logdir\12_Rewrite.txt"
)

```

```

    )
    Start-Process "msiexec.exe" -ArgumentList $MSIArguments -Wait -
NoNewWindow
    if (Select-String -path "$logdir\12_Rewrite.txt" -pattern "Installation
success or error status: 0.") {
        Write-Host "    Setup OK" -ForegroundColor Green
    }
    else {
        throw "Setup Failed"
    }

    ### Enterprise Manager Web
    Write-Host "    Installing Enterprise Manager Web ..." -ForegroundColor
Yellow
    $MSIArguments = @(
        "/i"
        "$source\EnterpriseManager\Backupweb_x64.msi"
        "/qn"
        "/L*v"
        "$logdir\13_Entweb.txt"
        "ACCEPTTEULA=YES"
        "VBREM_LICENSE_FILE=$licensefile"
        "VBREM_SERVICE_USER=$fulluser"
        "VBREM_SERVICE_PASSWORD=$password"
        "VBREM_SQLSERVER_SERVER=$env:COMPUTERNAME\VEEAMSQL2012"
    )
    Start-Process "msiexec.exe" -ArgumentList $MSIArguments -Wait -
NoNewWindow
    if (Select-String -path "$logdir\13_Entweb.txt" -pattern "Installation
success or error status: 0.") {
        Write-Host "    Setup OK" -ForegroundColor Green
    }
    else {
        throw "Setup Failed"
    }

    ### Enterprise Manager Cloud Portal
    Write-Host "    Installing Enterprise Manager Cloud Portal ..." -
ForegroundColor Yellow
    <#
    $MSIArguments = @(
        "/i"
        "$source\Cloud Portal\BackupCloudPortal_x64.msi"
        "/L*v"
        "$logdir\14_EntCloudPortal.txt"
        "/qn"
        "ACCEPTTEULA=YES"
    )
    Start-Process "msiexec.exe" -ArgumentList $MSIArguments -Wait -
NoNewWindow
    #>
    Start-Process "msiexec.exe" -ArgumentList "/i `"$source\Cloud
Portal\BackupCloudPortal_x64.msi`" /l*v $logdir\14_EntCloudPortal.txt /qn
ACCEPTTEULA="YES" "" -Wait -NoNewWindow

    if (Select-String -path "$logdir\14_EntCloudPortal.txt" -pattern
"Installation success or error status: 0.") {

```

```
        Write-Host "    Setup OK" -ForegroundColor Green
    }
    else {
        throw "Setup Failed"
    }

    ### Update 3
    Write-Host "Installing Update 3 ..." -ForegroundColor Yellow
    $Arguments = "/silent /noreboot /log $logdir\15_update.txt
VBR_AUTO_UPGRADE=1"
    Start-Process
"$source\Updates\veeam_backup_9.5.0.1536.update3_setup.exe" -ArgumentList
$Arguments -wait -NoNewWindow
    #endregion
}
}
```

I hope you enjoyed this post and please feel free to update my code, use it in your lab, and as always happy learning.

## Chapter 7

# Veeam for Nutanix AHV

By: Ian Sanderson (VMWare vEXPERT / Veeam Vanguard)

## Background

Veeam announced their upcoming Nutanix Acropolis Hypervisor support a few months ago, and since then there have been various internal betas available for testing, one of which I demoed in this online presentation. (<https://www.youtube.com/watch?v=sRP3URXyZqI>) The latest Beta is near feature complete, so I thought it would be great to share what we can expect from the full product. This blog post will run through the initial configuration of the Veeam Nutanix appliance and creation of a backup job.

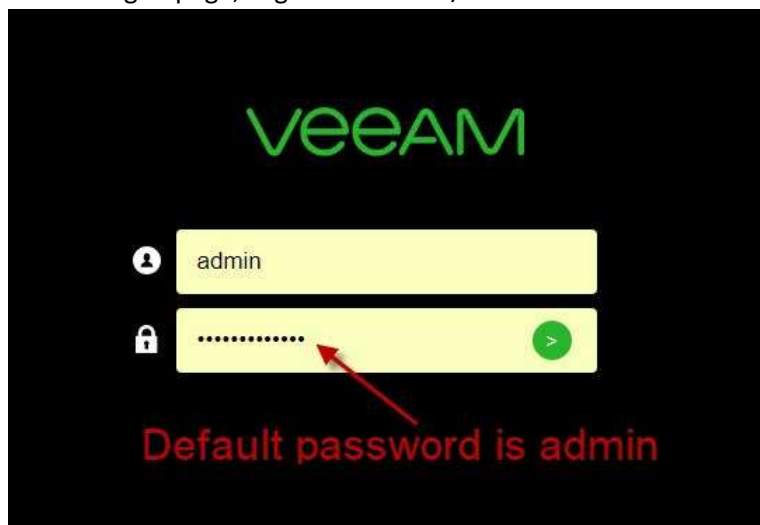
## Before you begin

There are some prerequisites for the Nutanix appliance to work correctly. You need to have a Veeam Backup and Replication server running version 9.5U3a or later to connect the appliance to. In my run through below, I have already deployed the virtual appliance machine. Screenshots start at initial appliance configuration.

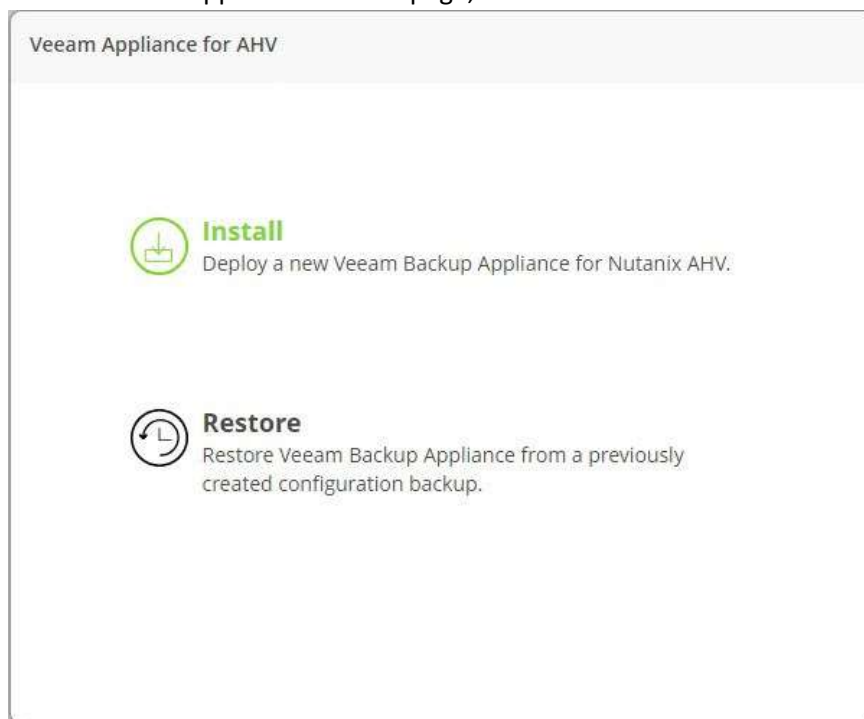
## Appliance Setup

First things first, connect to the IP address of the appliance on port 8100. The console of the virtual machine should show you the IP address.

1. On the Logon page, Login with admin/admin



2. On the Veeam Appliance for AHV page, click Install



3. On the Veeam Appliance for AHV Installation Page, read and accept the EULA and click Next.

The screenshot shows a window titled "Veeam Appliance for AHV: installation". On the left is a vertical progress bar with four steps: 1. End User License Agreement (highlighted with a green circle), 2. Credentials, 3. Network Settings, and 4. Summary. The main area is titled "End User License Agreement" and contains the text: "Please read and accept the following license agreement." Below this is a scrollable box containing the EULA text. The text starts with "Veeam Software ('Veeam') End User Software License Agreement ('EULA')", followed by "IMPORTANT - READ CAREFULLY" and a paragraph explaining the legal binding nature of the agreement. It then lists "1.0 DEFINITIONS" and "1.1 'Fee(s)'" which defines fees as license, maintenance, professional services, and consulting. At the bottom of the scrollable box is a checkbox labeled "I accept the terms of license agreement", which is currently unchecked. A green "Next" button is located at the bottom right of the window.

Veeam Appliance for AHV: installation

1 End User License Agreement

2 Credentials

3 Network Settings

4 Summary

**End User License Agreement**

Please read and accept the following license agreement.

**Veeam Software ("Veeam")**  
**End User Software License Agreement ("EULA")**

**IMPORTANT - READ CAREFULLY**

This EULA is a legally binding agreement between licensee end user ("End User") and Veeam setting forth the terms and conditions governing the use and operation of Veeam's proprietary computer software products (the "Software") and the written technical specifications for the use and operation of the Software (the "Documentation"). Where the sense and context permit, references in this EULA to the Software include the Documentation. By downloading and installing, copying or otherwise using the Software, and/or otherwise accepting this EULA, End User agrees to be bound by the terms and conditions of this EULA. If End User does not agree to or accept the terms of this EULA, End User may not access or use the Software.

**1.0 DEFINITIONS**

1.1 "**Fee(s)**" means any License, Maintenance, professional services, consulting

☐ I accept the terms of license agreement

Next

4. On the Veeam Appliance for AHV Credentials Page, define a new password and click Next.

The screenshot shows a window titled "Veeam Appliance for AHV: installation". On the left, a vertical progress bar has four steps: 1. End User License Agreement, 2. Credentials (highlighted in green), 3. Network Settings, and 4. Summary. The main area is titled "Credentials" and contains the instruction: "Specify the credentials for the Veeam Availability for Nutanix AHV Appliance web console." Below this, there are four input fields: "Login:" with the text "admin", "Old Password:" with five asterisks, "New Password:" with ten asterisks, and "Confirm New Password:" with ten asterisks. At the bottom right, there are two buttons: "Previous" and "Next".

5. On the Veeam Appliance for AHV Network Settings Page, give the appliance a name and defined static IP address settings and click Next.

Veeam Appliance for AHV: installation

1

End User License Agreement

2

Credentials

3

Network Settings

4

Summary

**Network settings**

Specify the network for settings for the Veeam Availability for Nutanix AHV Appliance.

Appliance host name:

SnurfVeeamNTX01

☐ Obtain an IP address automatically

IP address:

10 . 0 . 1 . 10

Subnet mask:

255 . 255 . 0 . 0

Default gateway:

10 . 0 . 0 . 2

Preferred DNS server:

10 . 0 . 0 . 3

Previous

Next

6. On the Veeam Appliance for AHV Summary Page, click Finish.

Veeam Appliance for AHV: installation

1

End User License Agreement

2

Credentials

3

Network Settings

4

Summary

Summary

Please review the Appliance settings before continuing.  
Specified settings will be applied immediately after completing this wizard.

Hostname:	SnurfVeeamNTX01
IP settings	IPv4, Static
IP address:	10.0.1.10
Subnet mask	255.255.0.0
Default gateway	10.0.0.2
DNS server	10.0.0.3

Previous

Finish

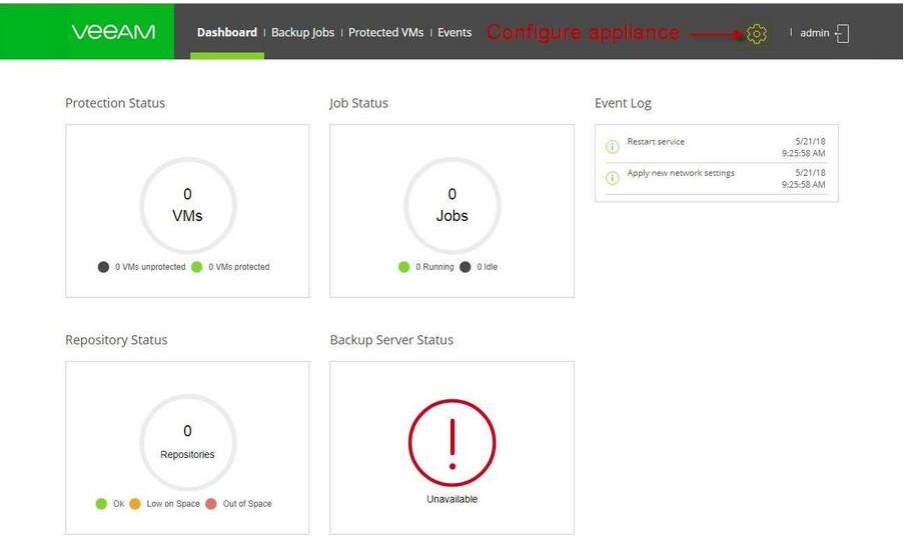
## Veeam and Nutanix Integration

1. On the Logon Screen type, the credentials defined earlier.

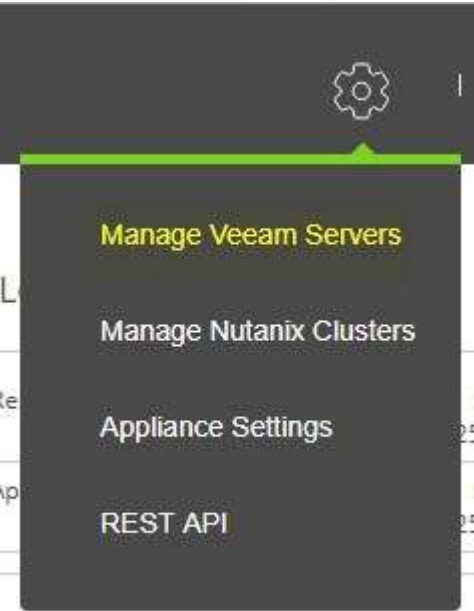


**Note:** that there is an error for the backup server. This is because the appliance has not been linked to a Veeam Backup and Replication server yet.

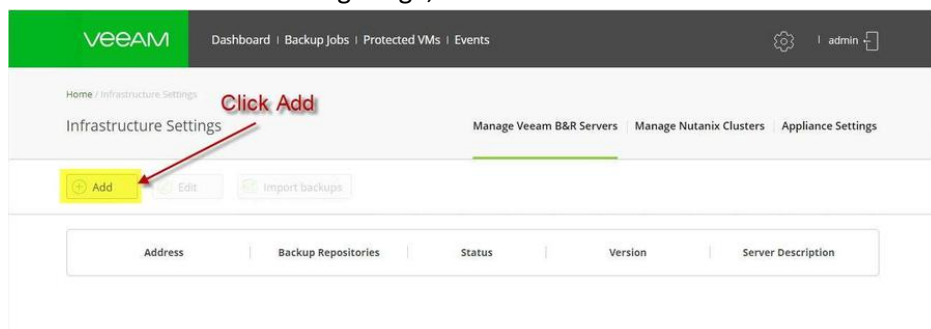
2. On the Dashboard, click the settings cog to configure integration.



3. Choose Manage Veeam Servers.



4. On the Infrastructure Settings Page, click add to add a Veeam server.



5. On the Add Veeam Server Page, type the DNS Name, Port, Description, UserName, and Password and click Add.

Add Veeam Server

DNS Name or IP:

snurfveeam01.snurf.local

Port:

10006

Description:

SNurf Veeam

User Name:

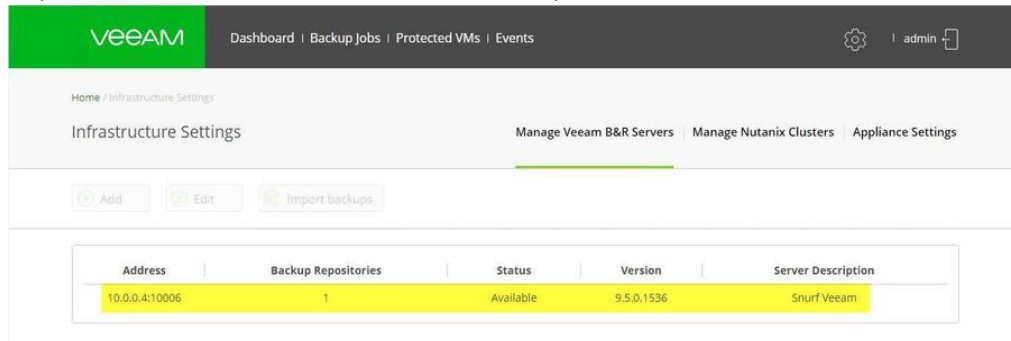
snurf\svc\_veeam

Password:

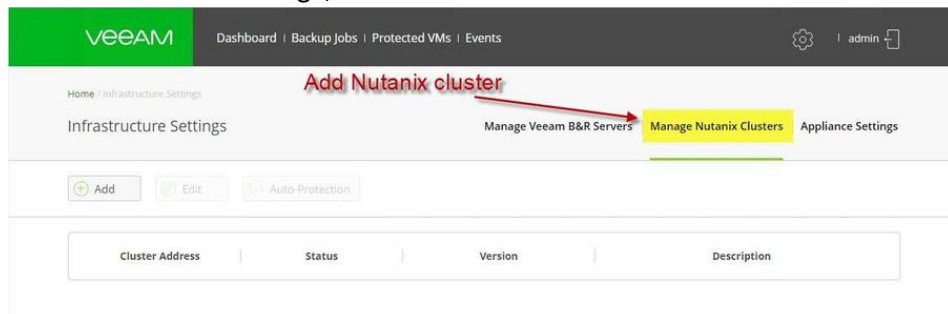
.....

Add

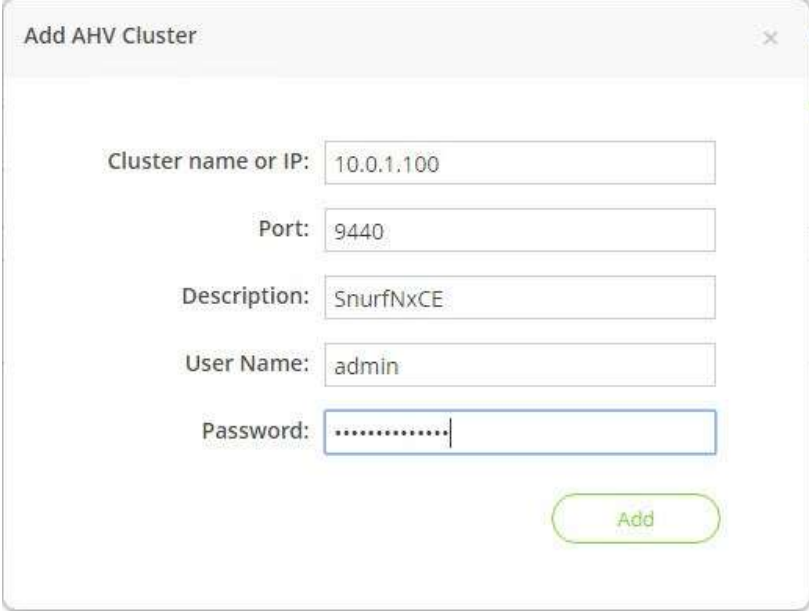
6. On the Infrastructure Page, Review the settings and ensure the Veeam Backup and Replication server has been added successfully.



7. On the Infrastructure Page, click as below to add a Nutanix cluster.



- On the Add AHV Cluster Page, type Cluster IP, Port, Description, User Name, Password and click Add.



**Add AHV Cluster**

Cluster name or IP: 10.0.1.100

Port: 9440

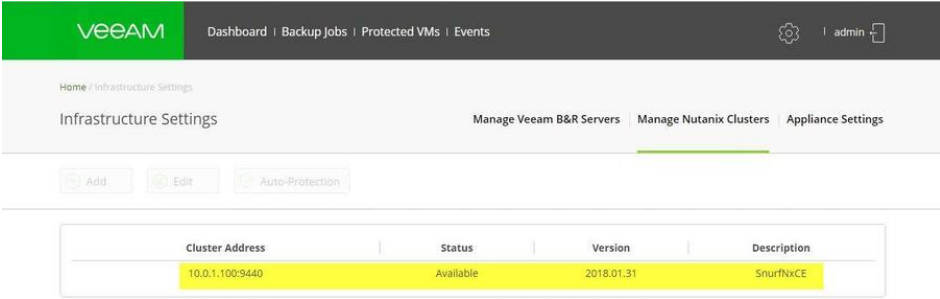
Description: SnurfNxCE

User Name: admin

Password: .....

Add

- On the Infrastructure Setup page, ensure the Nutanix cluster has been added to the appliance.



VEEAM | Dashboard | Backup Jobs | Protected VMs | Events | admin

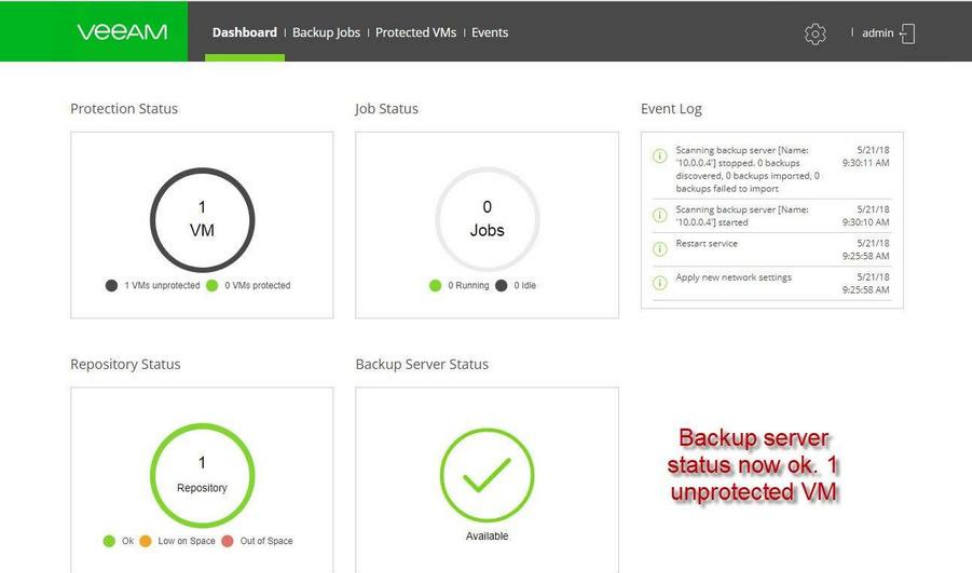
Home / Infrastructure Settings

Infrastructure Settings | Manage Veeam B&R Servers | Manage Nutanix Clusters | Appliance Settings

Add Edit Auto-Protection

Cluster Address	Status	Version	Description
10.0.1.100:9440	Available	2018.01.31	SnurfNxCE

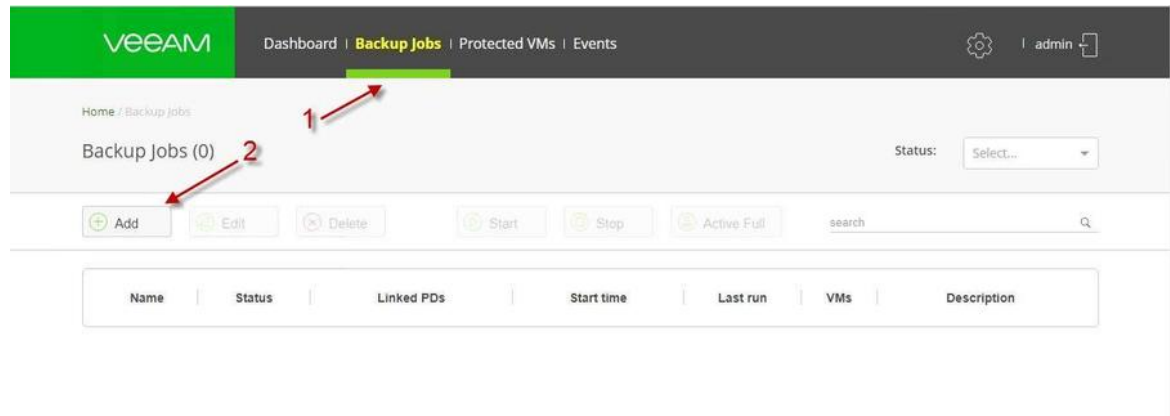
10. Now we have no errors for backup server status, and the appliance has discovered 1 unprotected virtual machine.



## Creating a Backup Job

Now that everything is configured, the next step is to create a backup job. I think it's a nice touch that out the box, it is easy to see if there are unprotected virtual machines. Veeam ONE is usually required for this information in a VMware or Hyper-V estate.

1. On the Backup Jobs Page, click add.



2. On the Add new Backup Job Page General Settings, type the backup job name and click Next.

Add new backup job

1 General Settings

2 Assign VMs

3 Backup Destination

4 Configure Schedule

5 Review Summary

**General Settings**

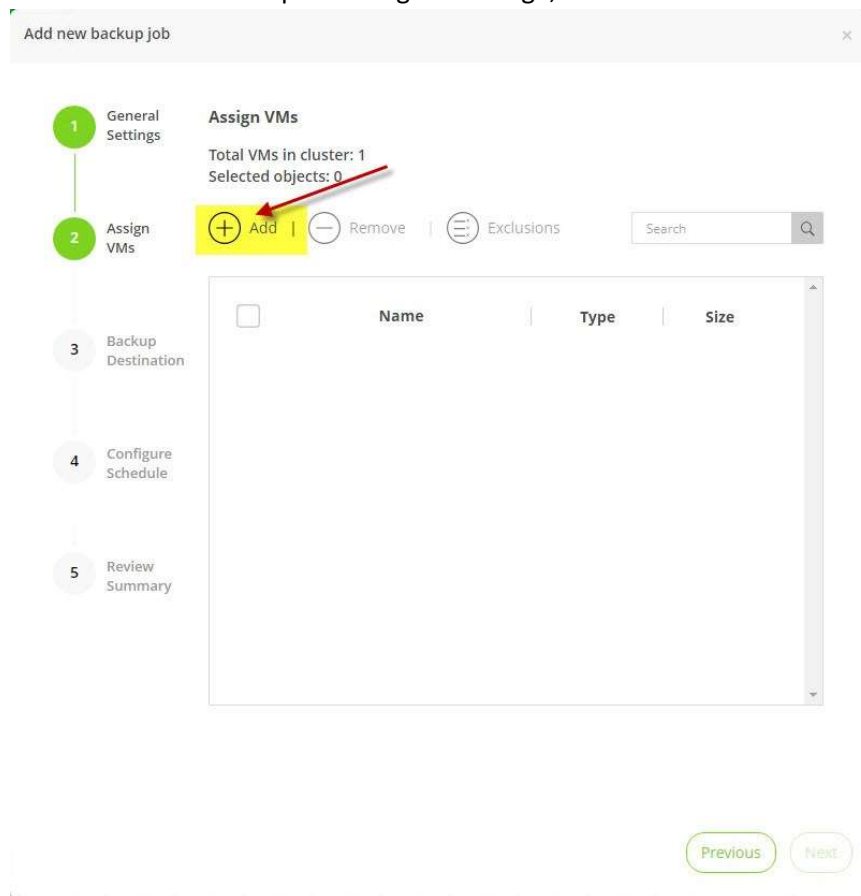
Job Name:

All VM's Backup

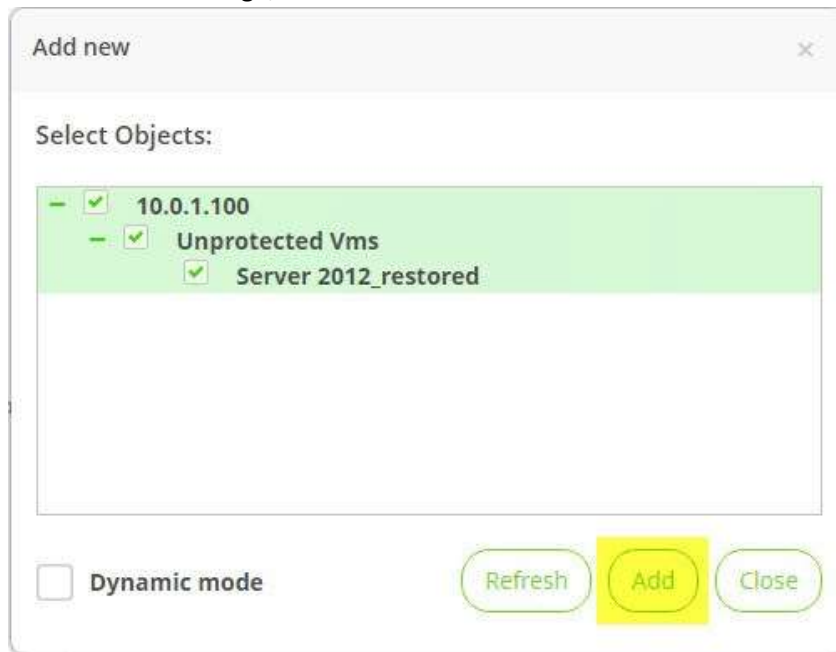
Job Description:

Next

3. On the Add new Backup Job Assign VMs Page, click add.



4. On the Add New Page, choose the VM's and click Add.



5. On the Add new Backup Job Assign VMs Page, Confirm the VM selection and click Add.

The screenshot shows the 'Add new backup job' wizard with five steps: 1. General Settings, 2. Assign VMs (current step), 3. Backup Destination, 4. Configure Schedule, and 5. Review Summary. The 'Assign VMs' step displays a table of VMs in the cluster. The table has columns for Name, Type, and Size. One VM is listed: 'Server 2012\_restored' of type 'VirtualMachine' and size '45 GB'. There are checkboxes for selecting VMs. Above the table are buttons for '+ Add', '- Remove', and 'Exclusions', along with a search bar. At the bottom right are 'Previous' and 'Next' buttons.

Name	Type	Size
Server 2012_restored	VirtualMachine	45 GB

6. On the Add new Backup Job Backup Destination Page, choose the backup target. This is any backup repository that has been defined in Veeam Backup and Replication and click

Next.

Add new backup job

1

General Settings

2

Assign VMs

3

Backup Destination

4

Configure Schedule

5

Review Summary

Backup Destination

Selected VMs: 1  
Approximate Backup size: 45 GB

Backup repository:  
SnurfVeeam01 - Local Repo - D: (--- Get size from Veeam SDK ---)

Capacity: ?

Advanced job settings including compression, deduplication and other settings.

Advanced

Above storage enumeration issue will be fixed in Veeam 9.5 U3a

Previous

Next

114

Technet24

7. On the Add new Backup Job Configure Schedule Page, Define the job schedule and click Next.

Add new backup job

1

General Settings

2

Assign VMs

3

Backup Destination

4

Configure Schedule

5

Review Summary

Configure Schedule

Specify the job scheduling. If a schedule is not set, the job will need to be controlled manually.

☐ Run the job automatically

☒ Daily at this time:

☐ Monthly at this time:

☐ Periodically every:

12:00

On week days

Restore points to keep on disk:

14

☐ Automatic retry

Retry failed VMs processing:

3

times

Wait before each retry attempt for

10

minutes

☐ Backup window

Terminate job if it exceeds allowed backup window

Windows

If the job does not complete within allocated backup window, it will be terminated

Previous

Next

115

8. On the Add new Backup Job Review Summary Page, click Finish.

Add new backup job

1

General Settings

2

Assign VMs

3

Backup Destination

4

Configure Schedule

5

Review Summary

Review Summary

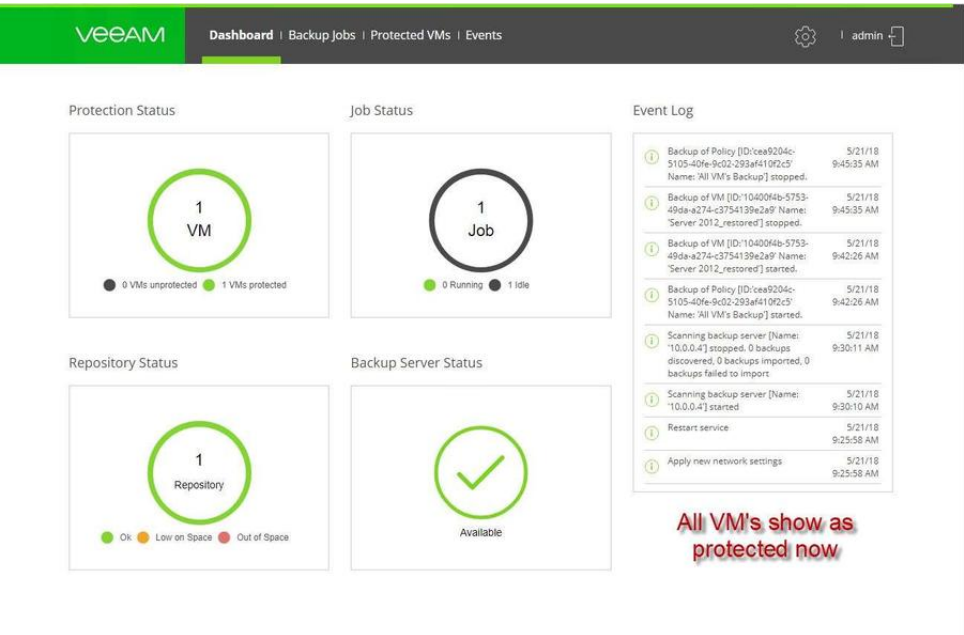
Name	All VM's Backup
VMs	1
Target repository	SnurfVeeam01 - Local Repo - D: (--- Get size fro...
Retention	14

☒ Run backup job when I click Finish

Previous

Finish

9. Everything is now green on the dashboard, which means happy days.



## Chapter 8

# Veeam Backup and Replication and SMB MultiChannel

By: Didier Van Hoya (Microsoft MVP / Veeam Vanguard)

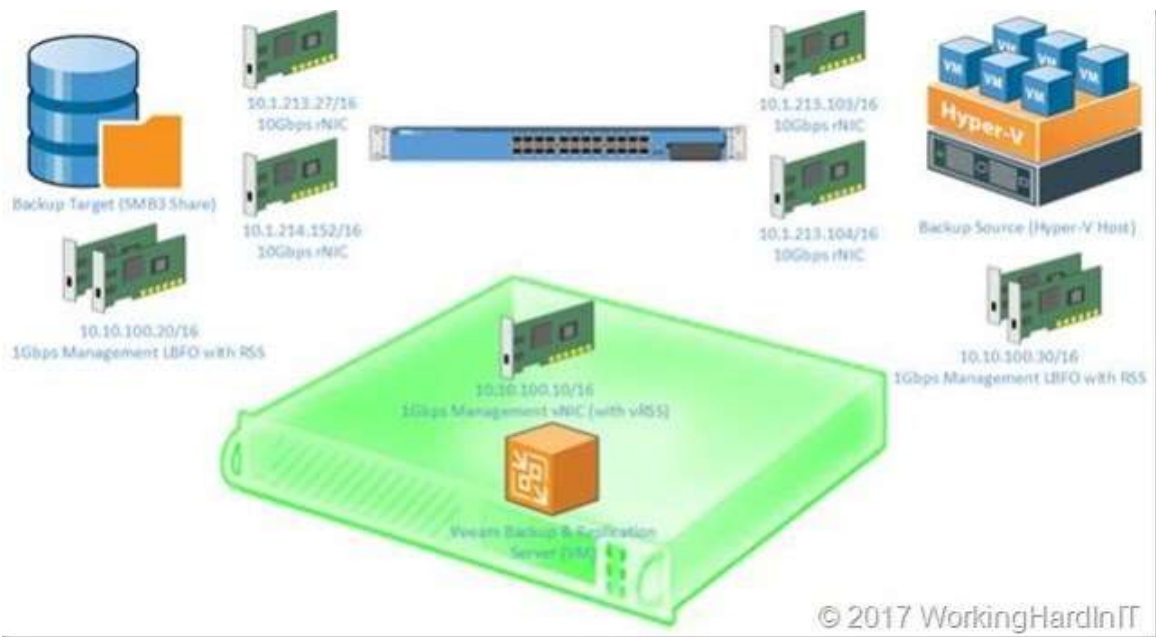
## Introduction

Is it true that Veeam Backup & Replication leverages SMB Multichannel? That is a question that I was asked recently. The answer is yes when you have a backup design and configuration that allows for this. If that's the case, it will even happen automatically when possible. That's how SMB 3 works. That means it's a good idea to pay attention to the network design so that you're not surprised by the route your backup traffic flows. Mind you that this could be a pleasant surprise, but you might want to plan for it.

I'll share a quick lab setup where SMB 3 Multichannel kicks in. Please don't consider this a reference guide for your backup architectural design but as a demo of how SMB multichannel can be leveraged to your advantage.

## Proving Veeam Backup and Replication leverages SMB Multichannel

Here's a figure of a quick lab setup I threw together.



There are a couple of significant things to note here when it comes to the automatic selection of the best possible network path.

SMB 3 Multichannel picks the best solution based on its logic. You can read more about that here: <https://blogs.technet.microsoft.com/josebda/2012/06/28/the-basics-of-smb-multichannel-a-feature-of-windows-server-2012-and-smb-3-0/>

I've included the figure with the overview below.

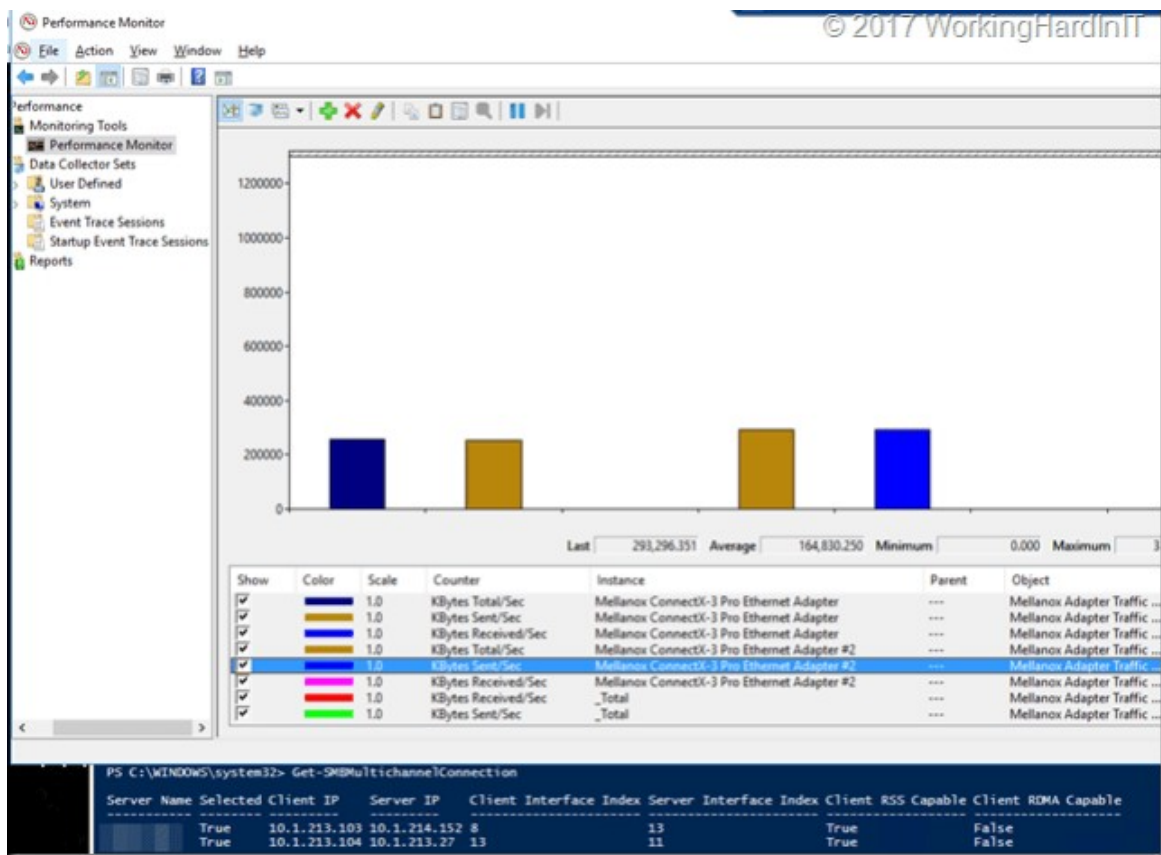
	Throughput	Fault Tolerance for SMB	Fault Tolerance for non-SMB	Lower CPU utilization
Single NIC (no RSS)	▲			
Multiple NICs (no RSS)	▲▲	▲		
Multiple NICs (no RSS) + NIC Teaming	▲▲	▲▲	▲	
Single NIC (with RSS)	▲▲			
Multiple NICs (with RSS)	▲▲▲	▲		
Multiple NICs (with RSS) + NIC Teaming	▲▲▲	▲▲	▲	
Single NIC (with RDMA)	▲▲▲			▲
Multiple NICs (with RDMA)	▲▲▲▲	▲		▲

© 2017 WorkingHardInIT

The figure nicely shows the capabilities of the NIC situation. To select the best possible network path SMB 3 uses the following logic:

1. RDMA capable NICs (rNICs) are preferred and chosen first. rNICs combine the highest throughput, the lowest latency and bring CPU offloading. on the processor when pushing through massive amounts of data.
2. RSS capable NICs: NICs with Receive Side Scaling (RSS) improve scalability by not being limited to core zero on the server. Configured correctly RSS offers the second-best capabilities.
3. The speed of the NICs is the 3rd evaluation criteria: a 10 Gbps NIC offers way more throughput than a 1 Gbps NIC.

Following this logic, it is clear that Multichannel will select our 2 RDMA capable 10Gbps NICs over the management LBFO interface which does not support RDMA and while supporting RSS can only deliver 2Gbps throughput at best. As seen in the screenshot below.



## Conclusion

So yes, Veeam Backup & Replication leverages SMB Multichannel! Please note that this did not require us to set SMB 3 Multichannel constraints or a preferred network for backups in Veeam Backup & Replication. It's possible to do so when needed, but ideally, you design your solution to have no need for this and let automatic detection chose the best network path correctly. This is the case in our little lab setup. The backup traffic flows over 10.1.0.0/16 network even when our Veeam Backup & Replication VM, the Hyper-V host, and the backup target have 10.10.0.0/16 as their management subnet. That's the one they exist on the Active Directory domain they belong to for standard functionality. However, as both the source and the target can be reached via 2\*10Gbps RDMA capable NICs on the 10.1.0.0/16 subnet SMB3 will select those according to its selection criteria. No intervention needed.

## SMB Direct Support

Now that we have shown that Veeam Backup & Replication backups in specific configurations can and will leverage SMB Multichannel to your benefit another question pops up. Can and does Veeam Backup & Replication leverage SMB Direct? The answer to that is also, yes. If SMB Direct is correctly configured on all the hosts and switches their networks paths in between, it will. Multichannel is the mechanism used to detect SMB Direct capabilities, so if multichannel works and sees SMB Direct is possible, it will leverage that. That's why when SMB Direct or RDMA is enabled on your NICs it's essential that it is configured correctly throughout the entire network path used. Poorly configured SMB Direct leads to horrible experiences.

Now think about that. High throughput, low latency, and CPU offloading, minimizing the CPU impact on your Hyper-V hosts, SOFS nodes, S2D nodes, and backup targets. Not bad at all, especially not since you're probably already implementing SMB Direct in many of these deployments. It's certainly something that could and should be considered when design solutions or optimizing existing ones.

## More SMB3 and Windows Server 2016 Goodness

When you put your SMB3 file share continuously available on a Windows 2012 (R2) or Windows Server 2016 cluster (it doesn't need to be on a CSV disk). You will gain high availability through transparent failover with SMB3, and except for a short pause, your backups will keep running even when the backup target node reboots or crashes after the File Server role has failed over. Now, start combining that with ReFSv3 in Windows Server 2016 and the Veeam Backup & Replication v9.5 support of this and you can see a lot of potential here to optimize many aspects of your backup design delivering effective and efficient solutions.

## Items to investigate further

One question that pops up in my mind is what happens if we configure a preferred backup network in Veeam Backup & Replication. Will this affect the operation of SMB multichannel at all? By that, I mean would enable a preferred network in Veeam to prevent multichannel from using more than one NIC?

In my opinion it should allow for multiple scenarios. When you have equally capable NICs that are on different subnets, you might want to make sure it uses only one. After all, Veeam uses the subnet to configure a preferred path or multiple subnets for that matter. Now multichannel will kick in with multiple equally capable NICs whether they are on the same subnet or not and if they are on the same subnet you might want them both to be leveraged even when setting a preferred path in Veeam. Remember that 1 IP / NIC is used to set up an SMB session and then it detects capabilities available, i.e., multiple paths, SMB Direct, RSS, speed, within 1 or across multiple subnets.

## Chapter 9

# Veeam Backup and Replication Preferred Subnet and SMB Multichannel

By: Didier Van Hoya (Microsoft MVP / Veeam Vanguard)

## Introduction

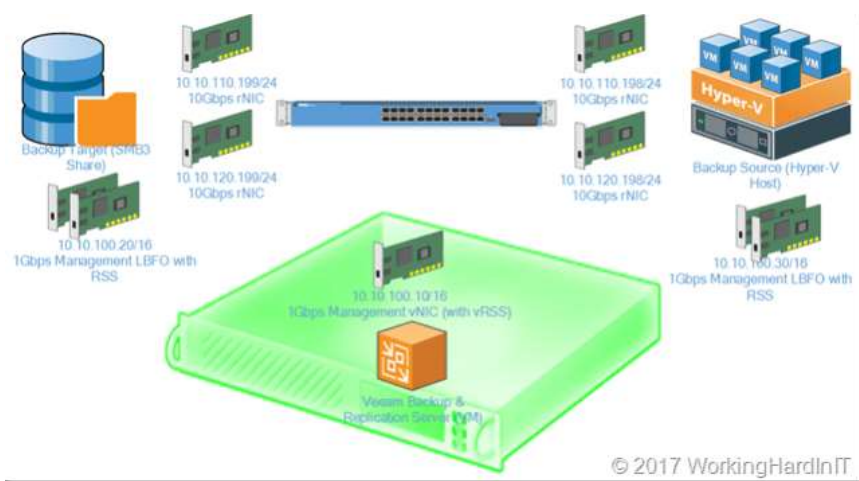
In the previous chapter, we showed that Veeam Backup & Replication leverages SMB multichannel when possible.

However, what about Veeam Backup & Replication Preferred Subnet & SMB Multichannel, does that work? We mentioned that we wanted to answer the question what happens if we configure a preferred backup network in Veeam Backup & Replication. Would this affect the operation of SMB multichannel at all? By that I means, would enable a preferred network in Veeam to prevent multichannel from using more than one NIC?

In this chapter, we dive into that question and some scenarios. We need to be able to deal with multiple scenarios. When you have equally capable NICs that are on different subnets, you might want to make sure it uses only one. Likewise, you want both to be used whether they are or are not on the same subnet even if you set a preferred subnet in Veeam. The good news is that the nature of SMB Multichannel and how Veeam preferred networks work do allow for flexibility to achieve this. However, it might not work like you would expect unless you understand SMB Multichannel.

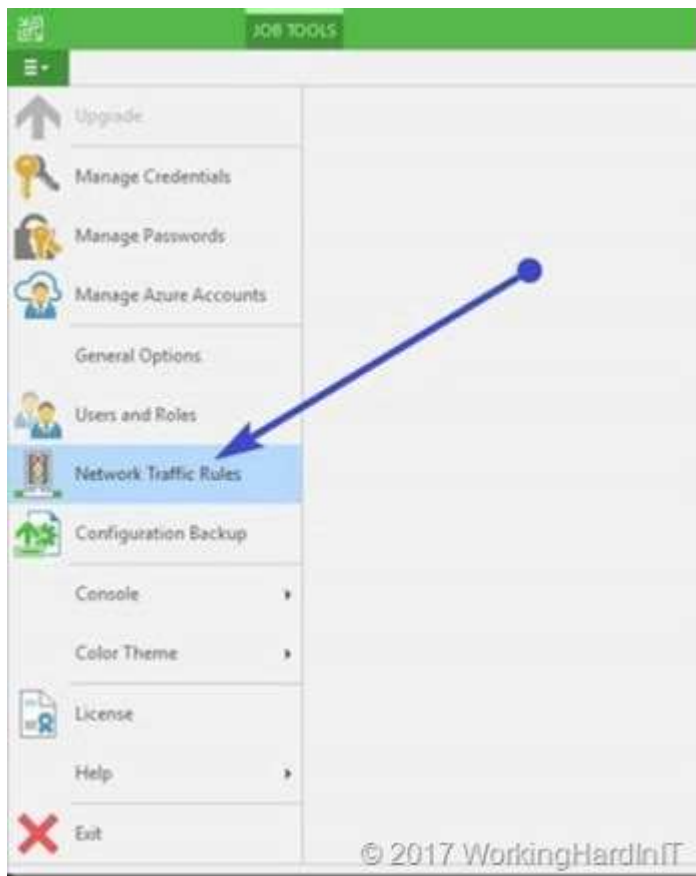
## Veeam Backup and Replication Preferred Subnet & SMB Multichannel

For this chapter, we adapt our lab networking a bit so that our non-management 10Gbps rNICs are on different subnets. We have subnet 10.10.110.0/24 for one set of NICs and 10.10.120.0/24 for the second set of NICs. This is shown in the figure below.

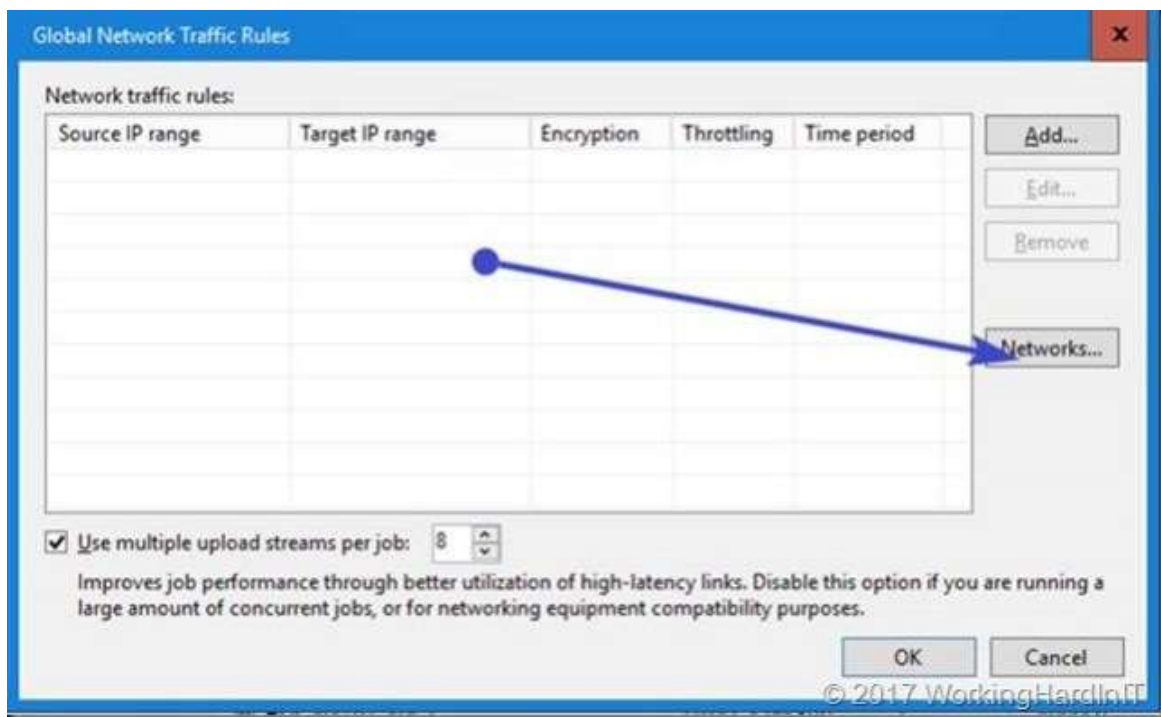


These networks can live in a separate VLAN or not, that doesn't matter. It does matter if to have a tagged VLAN or VLANs if you want to use RDMA because you need it to have the priority set.

We now need to configure our preferred network in Veeam Backup & Replication. We go to the main menu and select Network Traffic Rules



In the Global Network Traffic Rules window, click Networks.



In the Preferred Networks window, select the Prefer the following networks for backup and replication traffic checkbox.



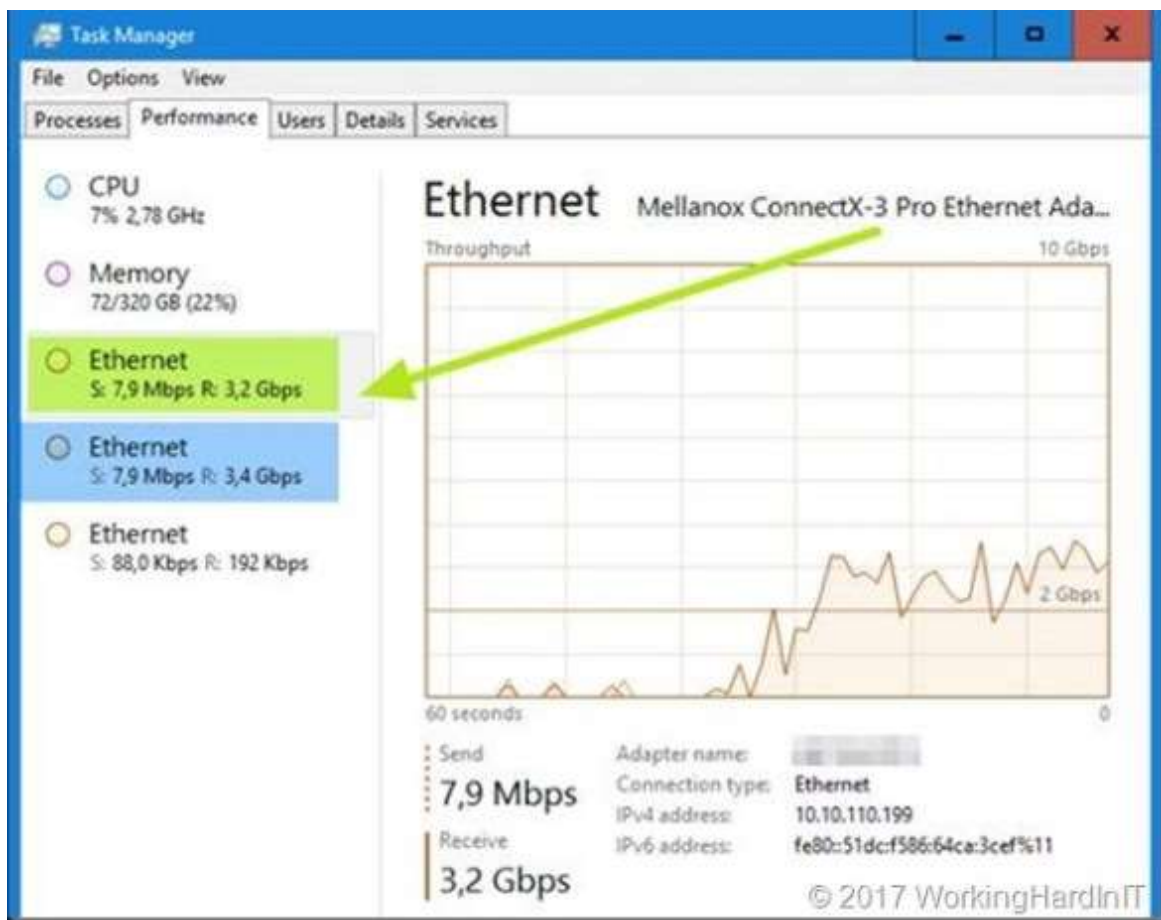
Click Add. We use the CIDR notation to fill out our preferred network, or you can use the network mask and click OK.

To prove a point in regards to how Multichannel works isn't affected by what you fill out here, we add only one of our two subnets here. SMB will see where it can leverage SMB multichannel and it will kick in. Veeam isn't blocking any of its logic.

So now we kick off a backup of our Hyper-V host to our SMB hare target backup repository. We can see multichannel work just fine.



Below is a screenshot on the backup target of the backup running over SMB multichannel, leveraging 2 subnets, while having set only one of those as the preferred network in Veeam Backup & Replication

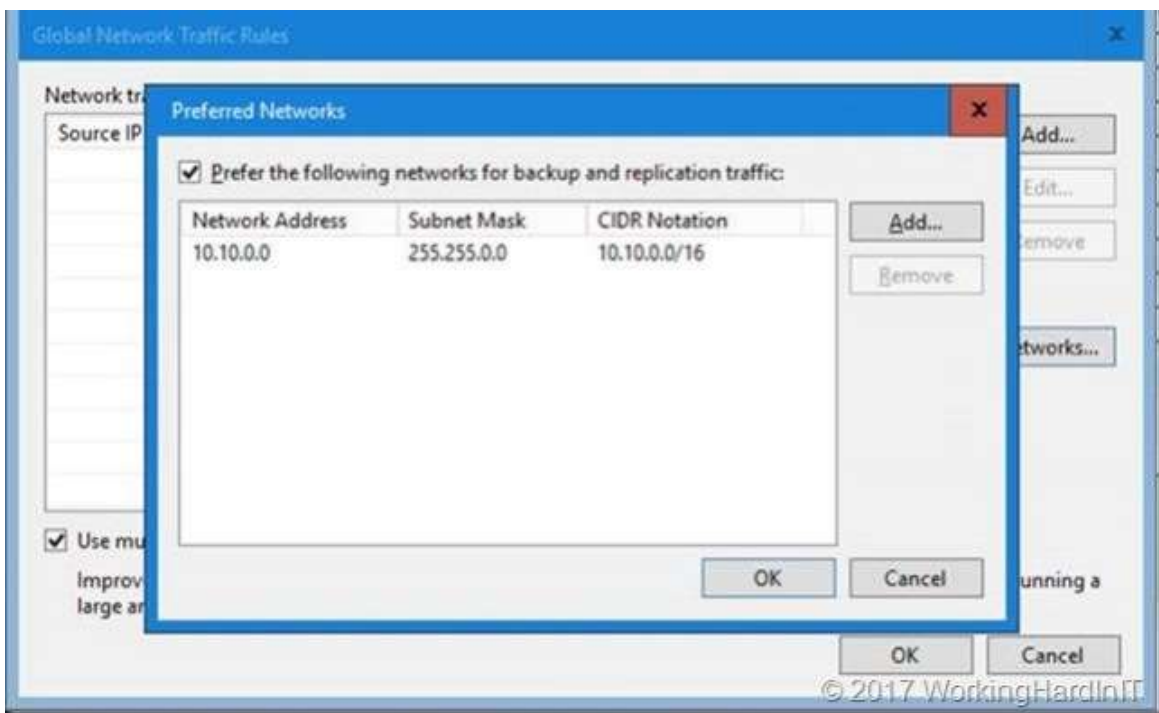


Look at my backup fly ... moreover, this is only one host that's being backup (4 VMs actually). Have I told you how much I love flash storage? Moreover, why I'm so interested in getting ReFS hybrid volumes with SSD/SATA disks to work as a backup target? I bet you do!

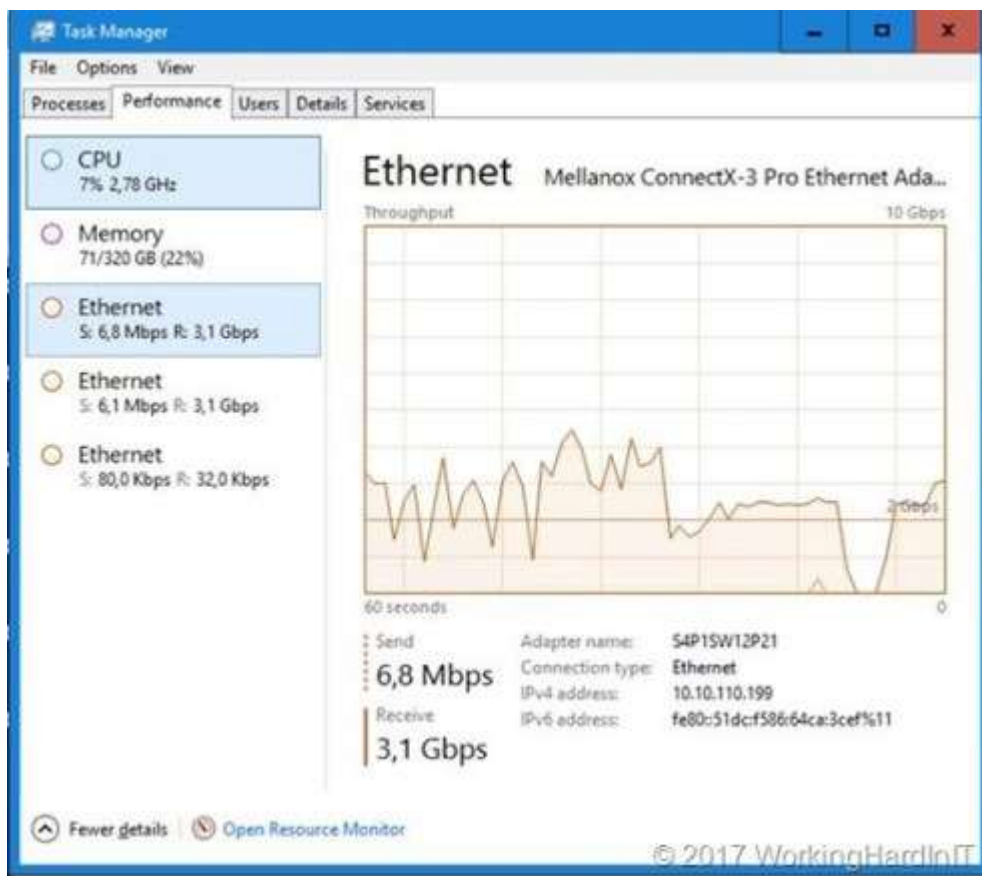
Looking good and it's easy, right? Well not so fast!

Veeam does not control SMB Multichannel

Before you think you're golden here and in control via Veeam lets, do another demo. In the preferred network, we enter a subnet available to both the source and the target server, but that is an LBFO (teamed) NIC with to 1Gbps members (RSS is enabled).



Now, let's see what happens when we kick off a backup.



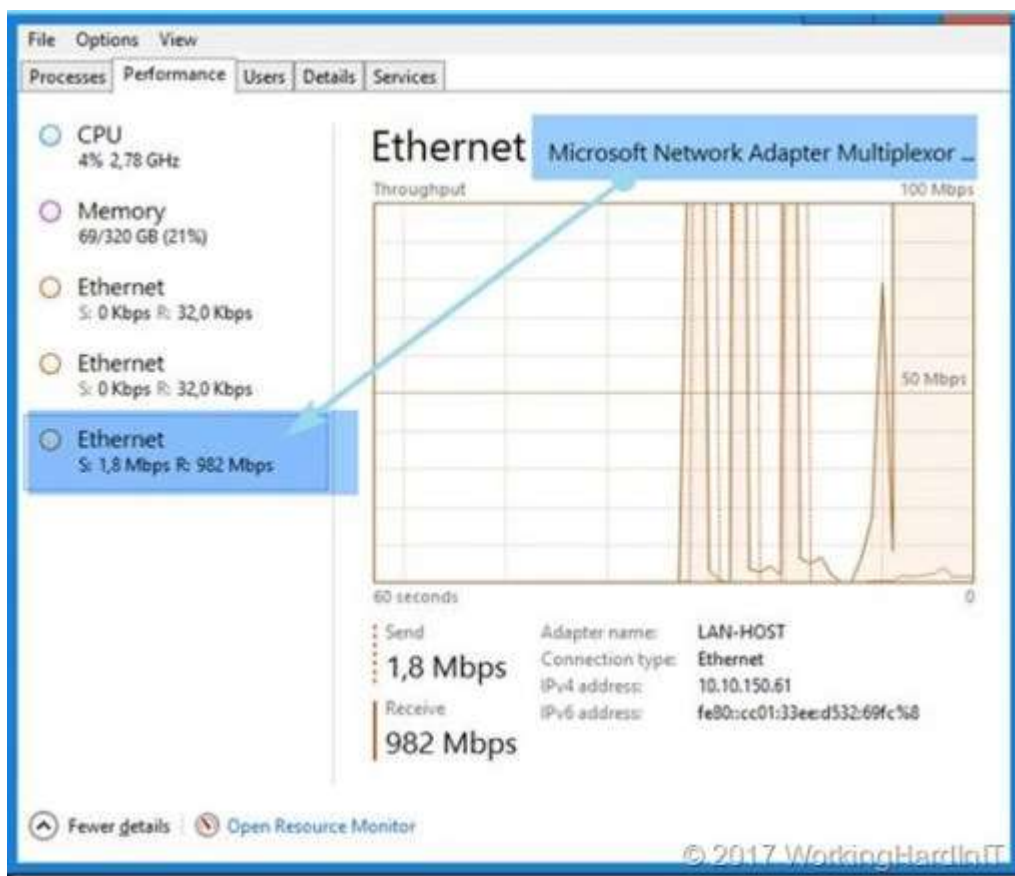
Well SMB multichannel went through its rules and decided to take the two best, equally capable NICs. These are still our two 10Gbps rNICs. Whatever you put in the preferred network is ignored.

This is neither good or bad, but you need to be aware of this to arrange for backups to leverage the network path(s) you had in mind. This is to avoid surprises. The way to do that the same as you plan and design for all SMB multichannel traffic.

As stated in the previous blog post you can control what NICs SMB multichannel will use by designing around the NIC capabilities or if needed disabling or enabling some of these or by disabling SMB multichannel on a NIC. This isn't always possible or can lead to issues for other

workloads, so the easiest way to go is using SMB Multichannel Constraints. Do note however that you need to take into consideration what other workloads on your server leverage SMB Multichannel when you go that route to avoid possible issues.

As an example, I disabled multichannel on my hosts. Awful idea but it's to prove a point. Moreover, still, with our 10.10.0.0/16 subnet set as a preferred subnet, I reran a backup.



As you can see the 2\*1Gbps, LBFO NIC is doing all the lifting on both hosts as it switches independent and not LACP load balancing mode we're limited to 1Gbps.

## **So how do we control the NICs used with SMB Multichannel?**

Well SMB Multichannel rules apply. You use your physical design, the capabilities of the NICs and SMB constraints. In reality, you're better off using your design, and if needed SMB multichannel constraints to limit SMB to the NICs, you want it to use. Do not that disabling SMB Multichannel (client and or server) is global for the host. Consider this as it affects all NICs on the host, not just the ones you have in mind for your backups.

In most cases, these NICs will be the same. Messing around with disabling multichannel or NIC capabilities (RSS, RDMA) isn't a great solution. However, it's good to know the options and behavior.

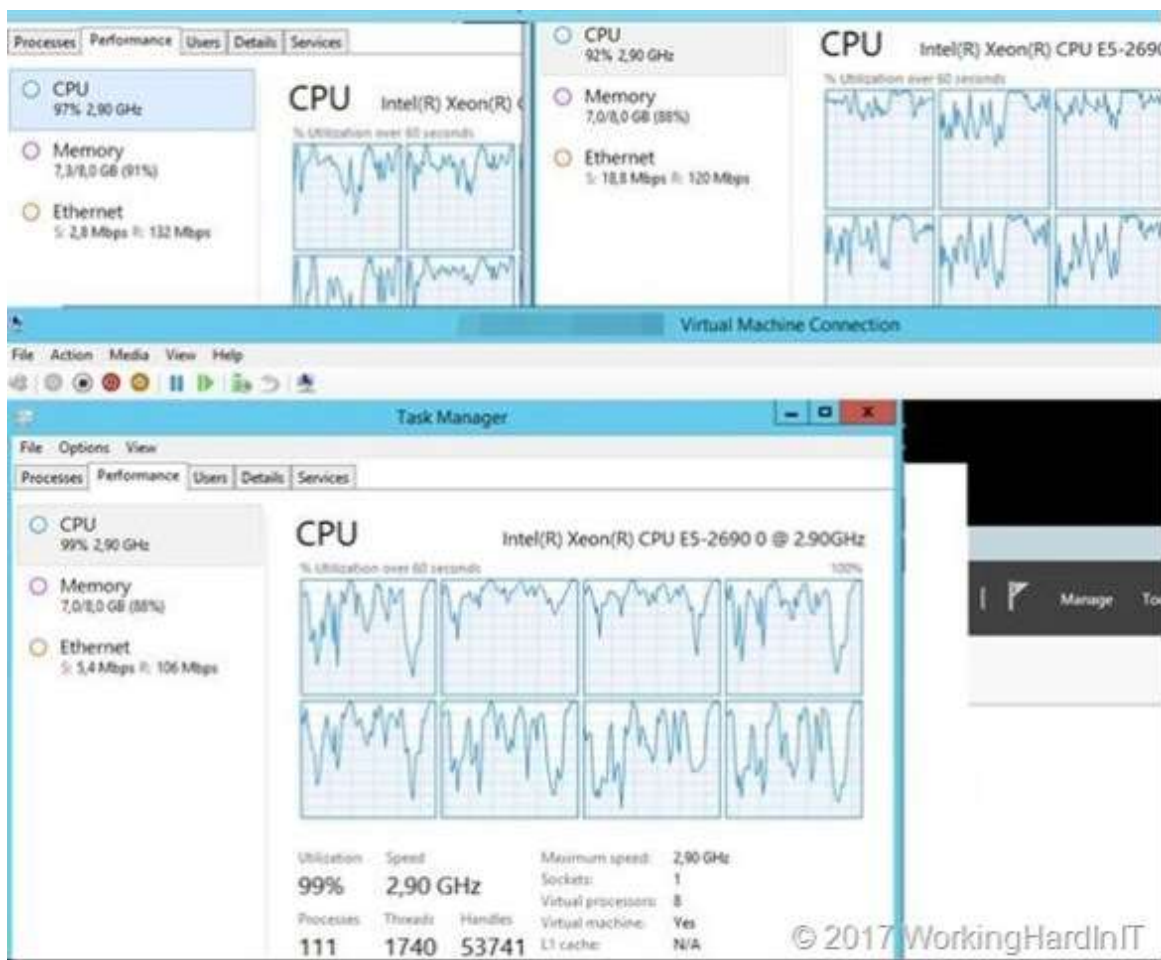
## **Some additional things to note**

Realize you don't even have to set both subnets in the preferred subnets if they are different. SMB kicks off over one see it can leverage both and just does so. The only thing you manipulated here SMB multichannel wise is which subnet is used first.

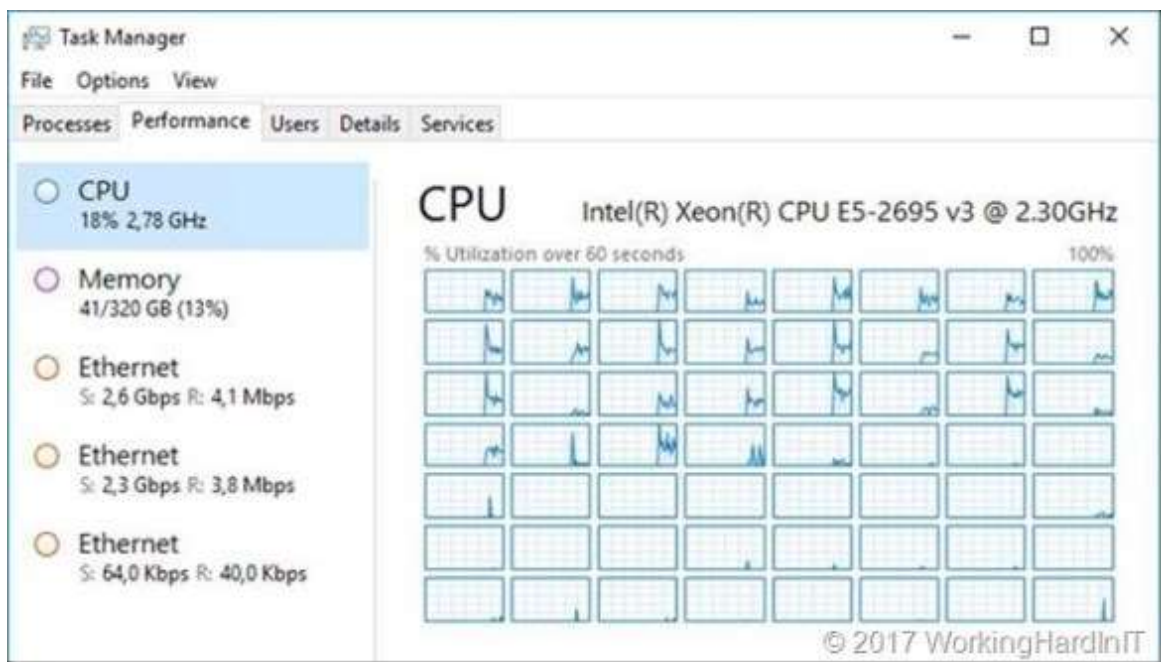
If both of our rNICs had been on the same subnet, you would not even have manipulated this.

Another thing that's worth pointing out that this doesn't require your Veeam Backup & Replication VM to have an IP address in any of the SMB multichannel subnets. So as long as the source Hyper-V hosts and the backup target are connected you're good to go.

Last but not least, and already mentioned in the previous blog post, this also leverages RDMA capabilities when available to help you get the best throughput, lowest latency and leave those CPU cycles for other needs. Scalability baby! No, I realize that you might think that the CPU offload benefit is not a huge deal on your Hyper-V host but consider the backup target being hammered by several simultaneous backups. Moreover, also consider that some people their virtual machines look like below in regards to CPU usage, in ever more need of more vCPU and CPU time slices.



Moreover, that's what the Hyper-V host looks like during a backup without SMB Direct (with idle VMs mind you).



All I'm saying here is don't dismiss RDMA too fast, everything you can leverage to help out and that is available for free in the box is worth considering.

Note: I have gotten the feedback that Veeam doesn't support SMB Direct and that this was confirmed by Veeam Support. Well, Veeam Backup & Replication leverages SMB 3, but that's an OS feature. Veeam Backup & Replication will work with SMB Multichannel, Direct, Signing, Transparent Failover ... It's out of the Veeam Backup & Replication scope of responsibilities as we have seen here. You feel free to leverage SMB Direct whether that is using iWarp/RoCE or Infiniband. This information was confirmed by Veeam and bears the "Anton Gostev seal of approval". So if SMB Direct cause issues you have a configuration problem with that feature, it's not Veeam not being able to support it, it doesn't know or care actually.

## Conclusion

The elegance and simplicity of the Veeam Backup & Replication GUI are deceiving. Veeam is extremely powerful and is surprisingly flexible in how you can leverage and configure it. I hope both my previous blog post and this one have given you some food for thought and ideas. There's more Veeam goodness to come in the coming months when times allows. Many years ago, when SMB 3 was introduced, I demonstrated the high availability capabilities this offered for Veeam backups.

## Chapter 10

# Using StarWind Cloud VTL for AWS and Veeam

**By: Karl Widmer (VMWare vEXPERT / Veeam Vanguard)**

Most of you know of Tape as a favorite type of backup media type. Moreover, some of you will not have good memories of them. Tapes we're once used for backups. Daily incremental and weekly full backups. It took ages to complete depending on the amount of data you had. You would be happy if a full backup fit into an 8-hour window throughout the night.

You thought tapes were dead but, tapes are coming into play more and more again. For example, you have heard about the WannaCry Virus, right? With that type of virus, you are going to want to ensure you have air gapped your backups. Even large companies like Google have all of its backups on tapes. Moreover, there are reasons why some companies still invent and develop new tape technologies like the new LTO standard, new drives, and libraries.

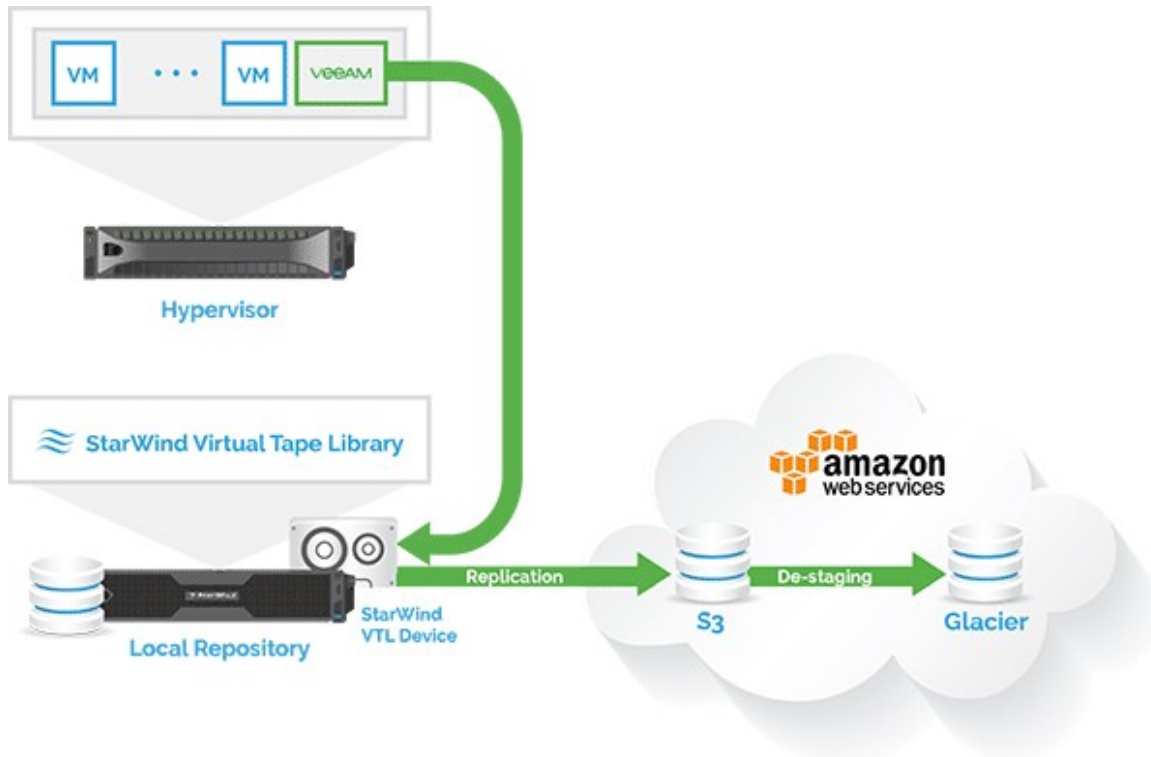
Most of us running labs at home don't have a tape drive nor a tape library. Moreover, this is where a Virtual Tape Library comes in very handy. You can emulate a complete tape library with software and use this emulation in an availability software like Veeam Backup & Replication or Veeam Availability Suite.

To make it clear, I'm not only talking about a solution for our labs at home. This emulation software is not a beta or alpha. It's a stable production version which is available on the market. In this chapter, you learn how to install and configure StarWind VTL for AWS and Veeam.

With StarWind Cloud VTL for AWS and Veeam, you can leverage your backup data for the usage with a virtual tape library to archive your backups to Amazon S3 storage in the cloud. This solution can also be extended to move your archive data from Amazon S3 buckets to Amazon Glacier.

## How does the Solution Work?

The solution we're building is shown in the image below:



- We have our hypervisor (in my case it's ESXi / VMware vSphere, but it can also be Microsoft Hyper-V) in place.
- Our virtual machines are running on this hypervisor.
- With Veeam Backup & Replication we create the backups of these virtual machines.
- StarWind Cloud VTL for AWS and Veeam creates a virtual tape library, which we use then in Veeam to have a "tape" backup.
- It's not just a VTL, but also StarWind replicates the backed up data to Amazon S3 storage, and optional to Amazon Glacier.

## Download the Software

First, we have to download all the needed software components. We need StarWind Cloud VTL for AWS and Veeam, also needed are the HPE StoreEver Tape Drivers, and, last but not least, Veeam Backup & Replication.

You can get the software here:

- StarWind Cloud VTL for AWS and Veeam

<https://www.starwindsoftware.com/starwind-cloud-vtl-for-veeam>

- HPE StoreEver Tape Drivers (Version 4.2.0.0 does support Windows Server 2016)

[https://h20566.www2.hp.com/hpsc/swd/public/detail?swItemId=MTX\\_7e9f343865d1445e92cfbaf0b1](https://h20566.www2.hp.com/hpsc/swd/public/detail?swItemId=MTX_7e9f343865d1445e92cfbaf0b1)

- Veeam Backup & Replication (if you haven't it already in use)

<https://www.veeam.com/downloads.html>

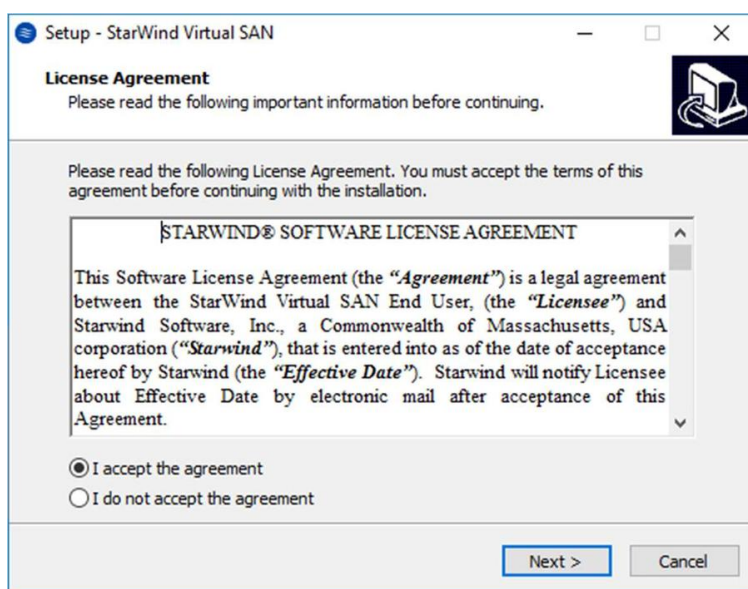
Save these components somewhere on your backup server. It doesn't matter if it's a virtual or physical server.

## Installation of StarWind VTL

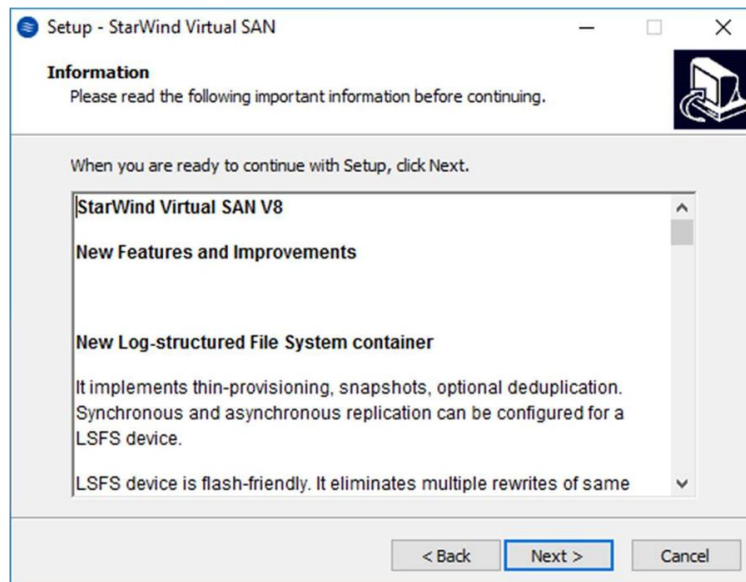
I assume that you've installed Veeam Backup & Replication or Veeam Availability Suite already, and you're now ready for the next steps.

The first component we install is StarWind Cloud VTL for AWS and Veeam. You've downloaded the software already and got your license file ready. Let's start.

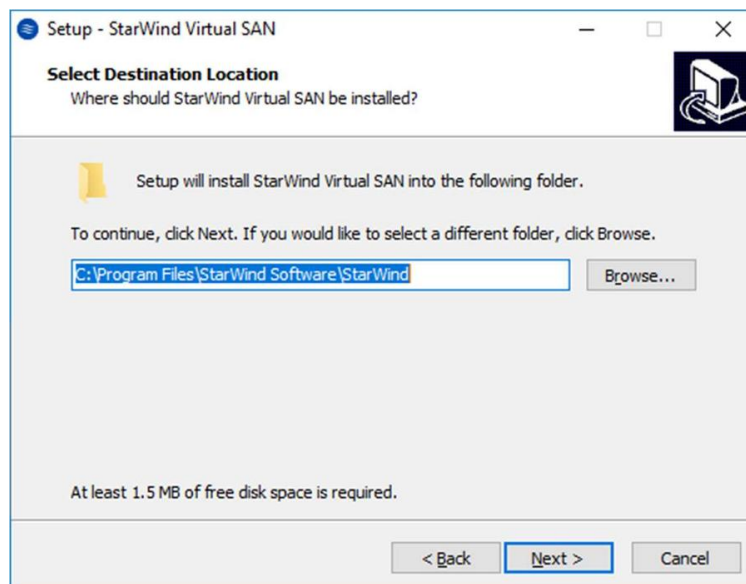
1. Double-click the installer file you've downloaded before to start the setup process.
2. On the License Agreement page, Accept the license agreement and click Next.



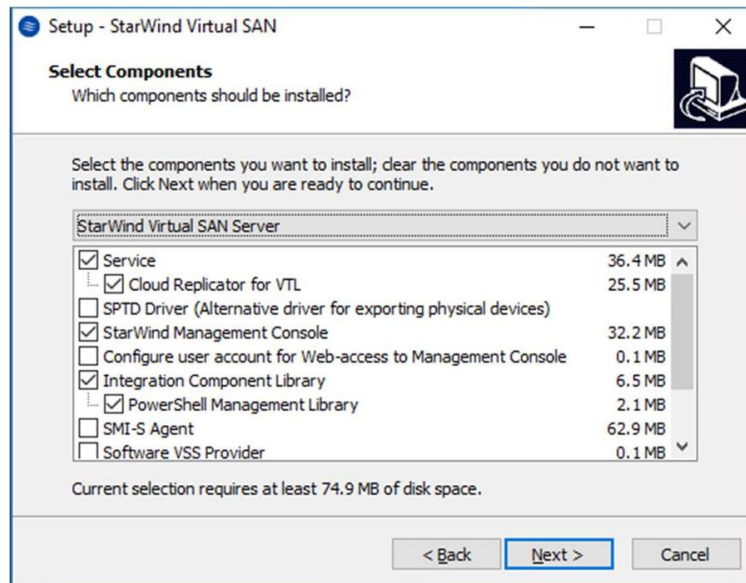
3. On the information Page, Click Next.



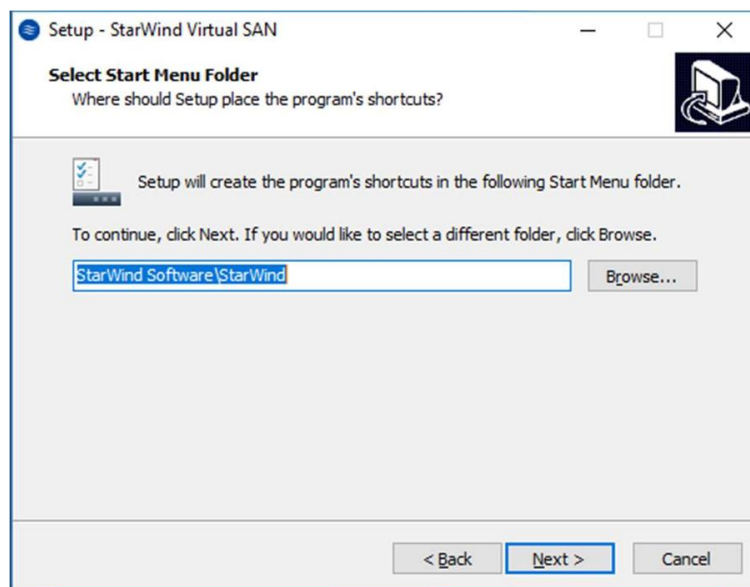
4. On the Set Destination Location Page, Set or Select the Installation folder and click next.



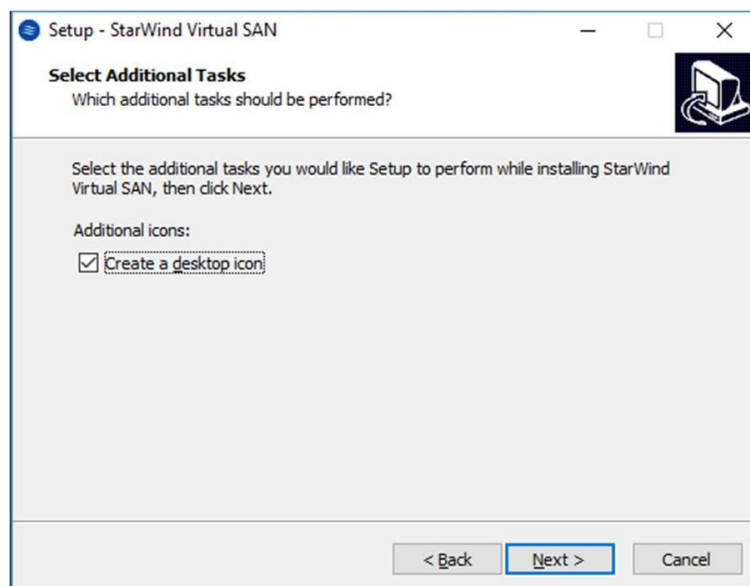
5. On the Select Components Page, Select the components you'd like to install and click Next.



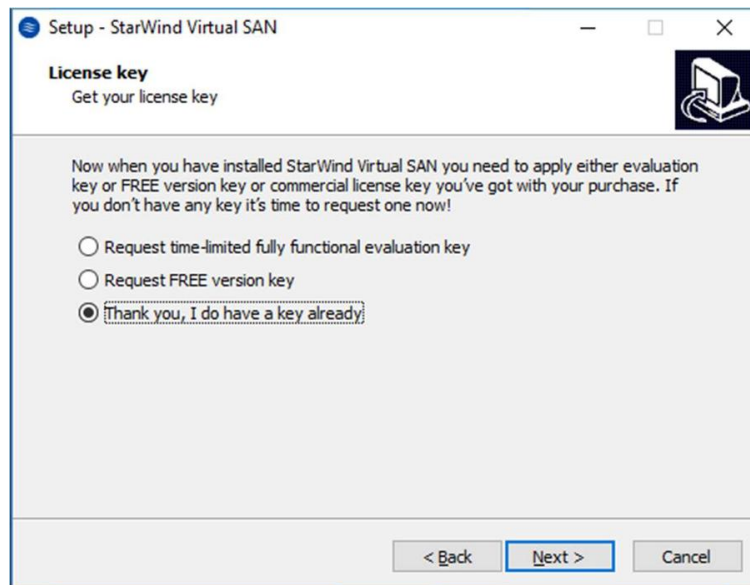
6. On the Select Start Menu Folder Page, check the location of the application shortcut in the start menu and click Next.



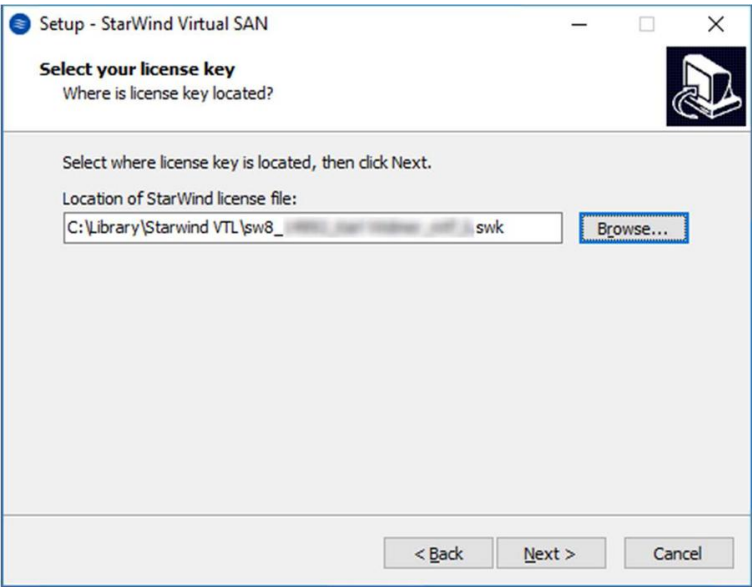
7. On the Select Additional Tasks Page, Select Create a desktop icon and click Next.



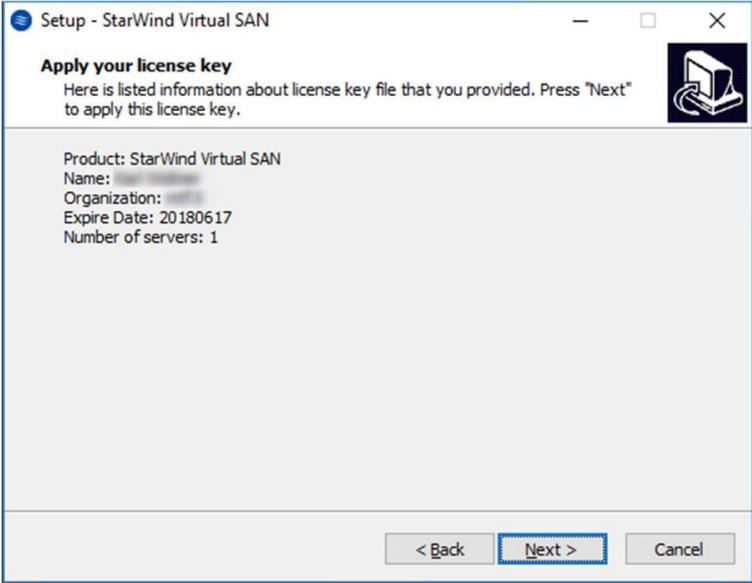
8. On the License Key page, select Thank you, I do have a key already and click Next.



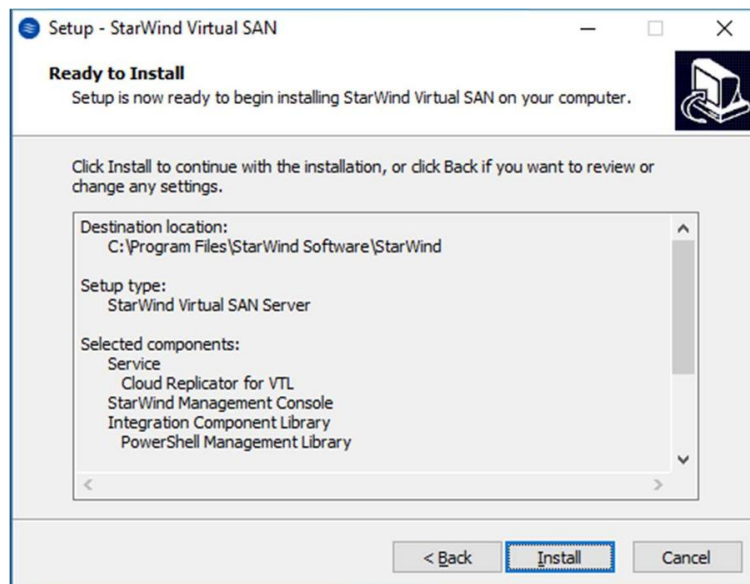
9. On the Select your license key page, browse to the path of your license file and click Next.



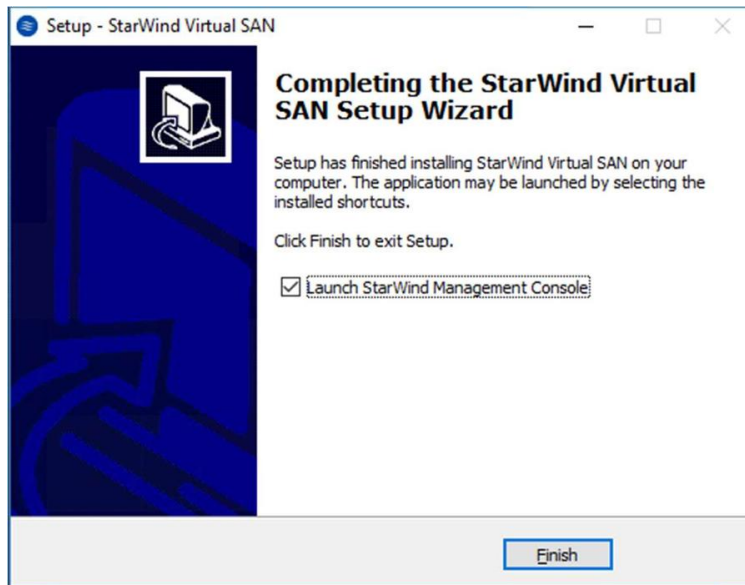
10. On the Apply your license key page, review your license details and click Next.



11. On the Ready to Install page, click Install.



12. On the Completing the StarWind Virtual SAN Setup Wizard page, select Launch StarWind Management Console and click Finish.



Congratulations you've installed StarWind VTL now on your server. Don't be worried about the name "StarWind Virtual SAN." The StarWind Virtual SAN is the tool that is used to configure the newly installed Virtual Tape Library that is used later in this chapter.

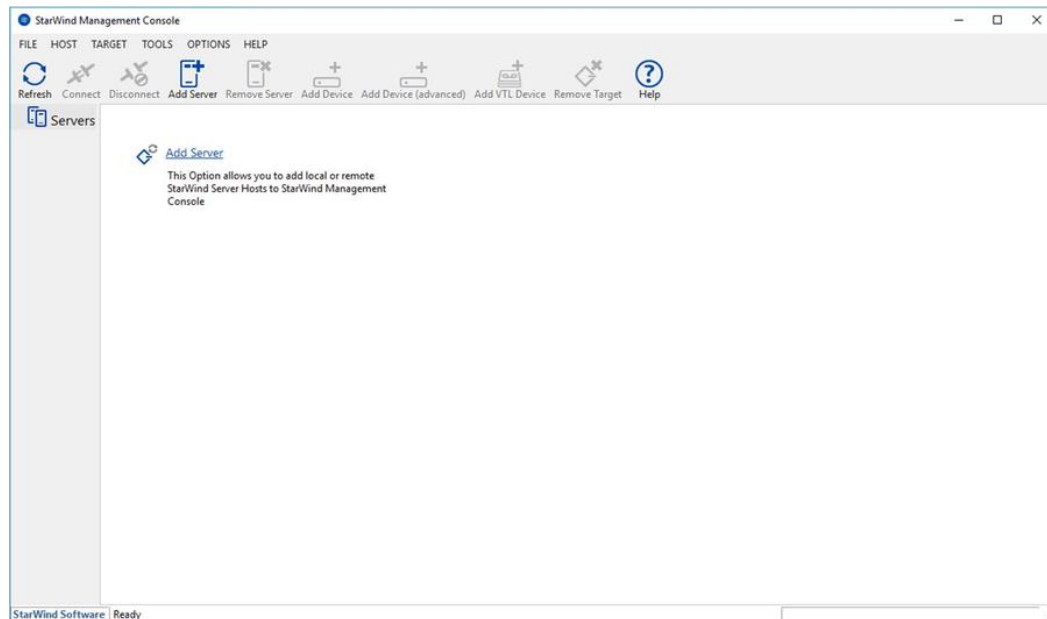
## The configuration of StarWind VTL

The next step after installation is the configuration of this software. Now, we will setup a virtual tape library and configure it for the lab.

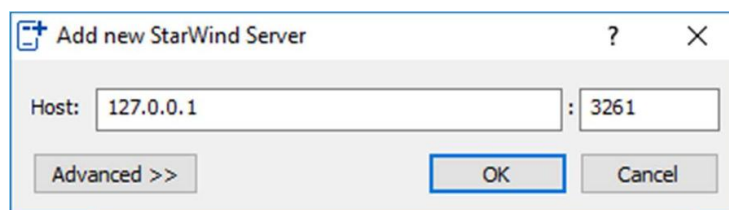
1. Double-click on the StarWind Management Console icon on your desktop.



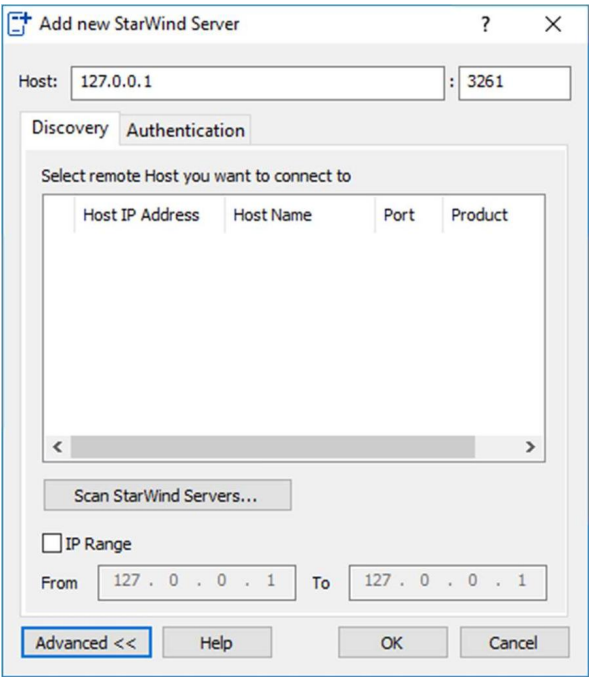
2. Now you're in the console. Let's get our hands on the next configuration steps. Click Add Server.



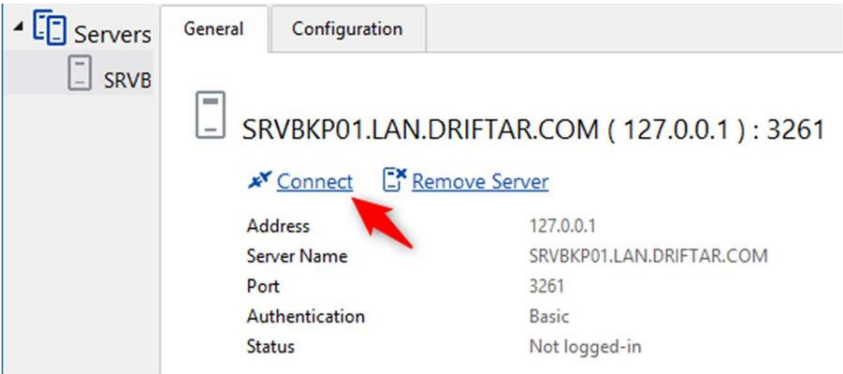
3. Since we're using the management console on the same server as we've installed it, you can add the localhost address.



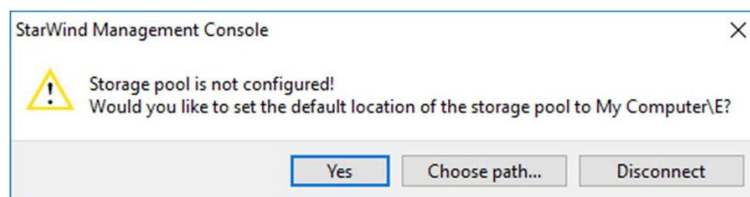
4. Now click Advanced and then Scan StarWind Servers. Your local server should show up in the list. Select it and click OK.



5. You'll see your server now in the management console. Click Connect.



6. Since we don't have a storage pool created, we have to configure one. Click Yes to use the default setting.



Note: This storage pool is a repository where the virtual tape files are stored. Depending on the configuration we set, later on, these virtual tape files will be stored or deleted after replication.

7. To add a virtual tape library, click Add VTL Device in the management console.



8. Since we don't have a VTL in use at the moment, we create a new one. Set a name and click Next. As you can see on the screenshot, the location is "E:\VTL" of my server. It will use the storage pool you've created before. You can also create another storage pool if

needed.

The screenshot shows a Windows-style dialog box titled "Add Device Wizard" with a back arrow, a question mark, and a close button (X). The main heading is "Virtual Tape Library Location". There are two radio button options: "Create a New Virtual Tape Library" (which is selected) and "Use an Existing Virtual Tape Library". Under the selected option, there are two text input fields: "Name:" with the text "VTL-Veeam-AWS" and "Location:" with the text "My Computer\E|VTL\" and a browse button (three dots). Under the unselected option, there is a "Location:" field with a dropdown arrow and a browse button. At the bottom right, there are "Next" and "Cancel" buttons.

← Add Device Wizard ? X

Virtual Tape Library Location

☒ Create a New Virtual Tape Library

Name: VTL-Veeam-AWS

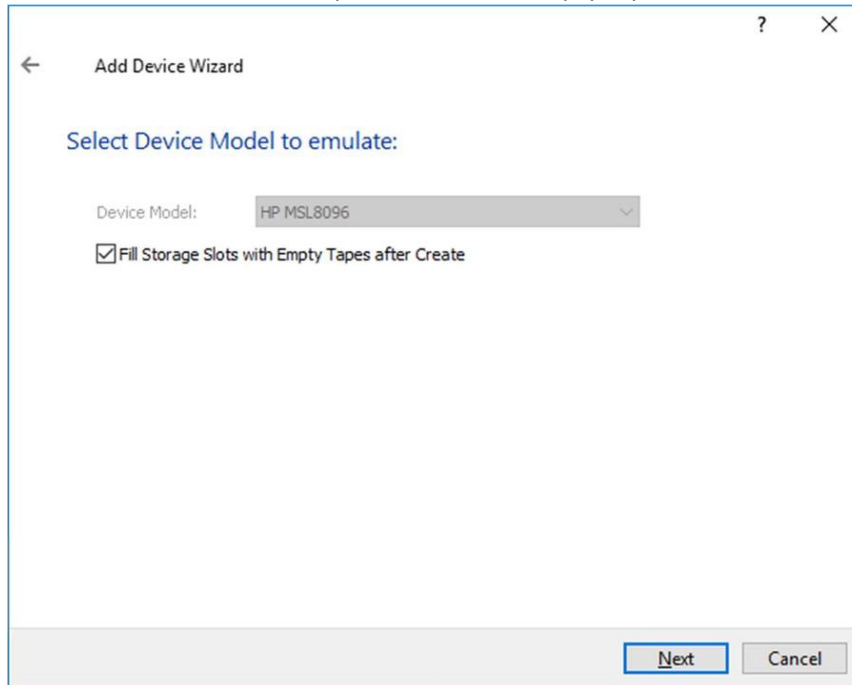
Location: My Computer\E|VTL\ ...

☐ Use an Existing Virtual Tape Library

Location: ...

Next Cancel

9. Next step is to set the tape library model. In our case it's an HP MSL8096. Select the checkbox to fill up the slots with empty tapes. Click Next.



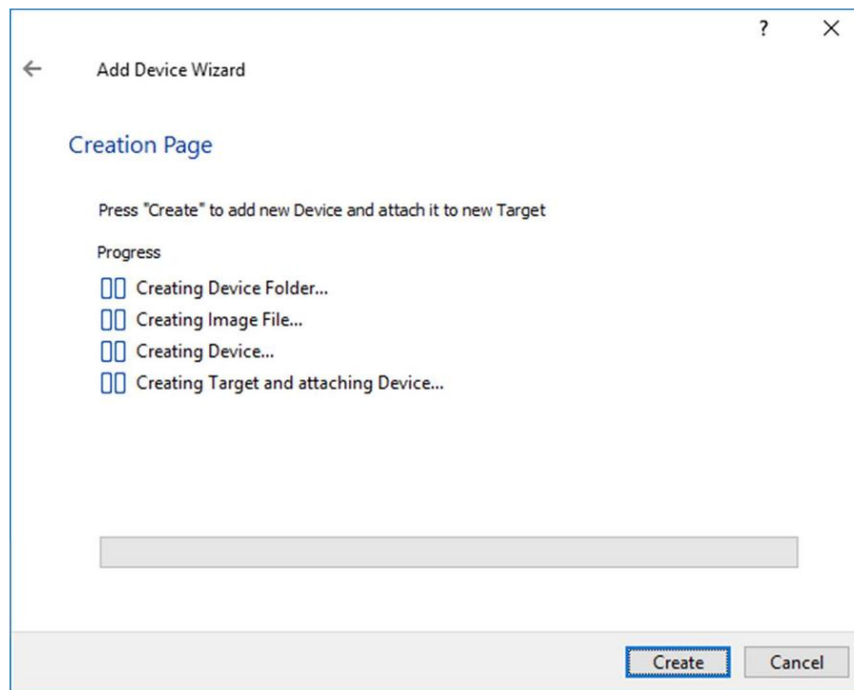
10. The VTL will be attached through iSCSI. Create a new target alias. Mark the checkbox to allow multiple concurrent iSCSI connections. Click Next.

The screenshot shows a window titled "Add Device Wizard" with a back arrow and help/cancel icons. The "Target Parameters" section contains the following fields and options:

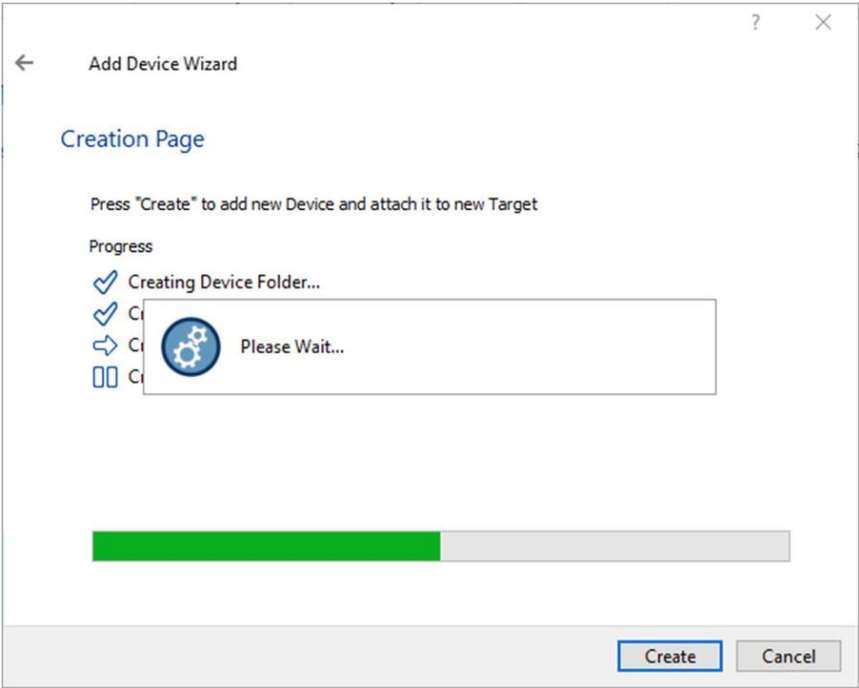
- Choose a Target Attachment Method:** A dropdown menu with "Create new Target" selected.
- Target Alias:** A text box containing "VTL-Veeam-AWS".
- Target Name:** An unchecked checkbox followed by a text box containing the iSCSI Qualified Name: "iqn.2008-08.com.starwindsoftware:svb01.lan.driftar.com-vtl-veeam-aws".
- Allow multiple concurrent iSCSI Connections:** A checked checkbox.

At the bottom right, there are "Next" and "Cancel" buttons.

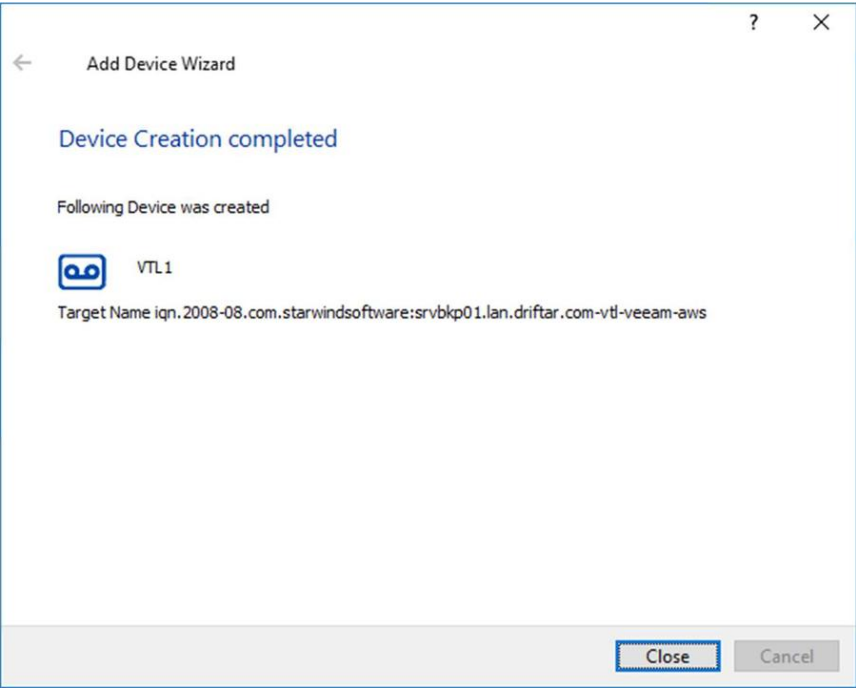
11. Click Create. The assistant will now create the environment.



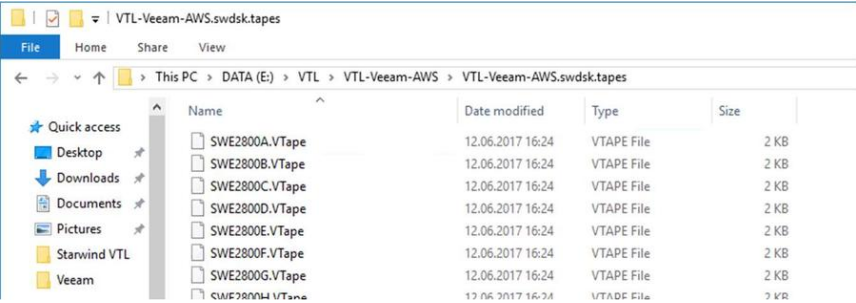
12. It will take a moment to complete.



13. After the wizard has completed the setup, click Close.



14. Let's have a quick look in Windows Explorer how this looks on our hard disk.



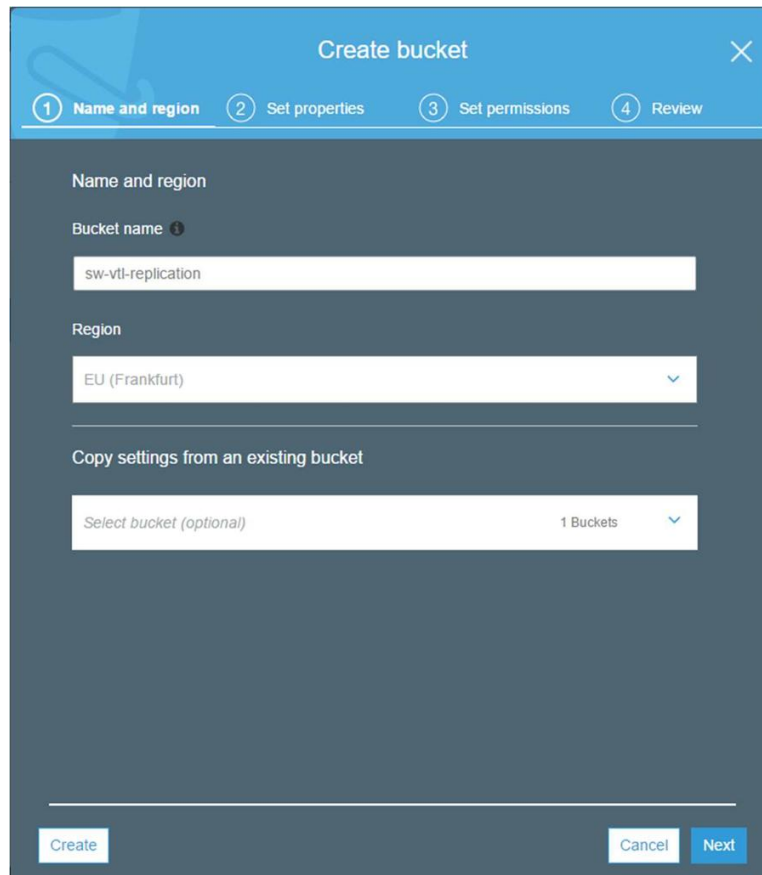
You see all the VTape files, about 96 in total.

We're done with the basic installation and configuration of our StarWind Cloud VTL for AWS and Veeam.

## Creating an Amazon S3 Bucket

I assume that you're already an Amazon Web Services user or customer and that you are familiar with creating buckets.

1. Log in to your Amazon AWS account. Make sure you have your Access Key ID and Secret Access Key ready. We need that later on.
2. Create a new S3 bucket and give it a name. The name has to be unique. Choose also the region where this bucket should be created. The closer the region is to your location, the better. Click Next.



The screenshot shows the 'Create bucket' wizard in the AWS Management Console. The wizard has a blue header with the title 'Create bucket' and a close button (X). Below the header is a progress bar with four steps: 1. Name and region (active), 2. Set properties, 3. Set permissions, and 4. Review. The main content area is dark gray and contains the following fields:

- Name and region**
  - Bucket name**: A text input field containing 'sw-vtl-replication'.
  - Region**: A dropdown menu showing 'EU (Frankfurt)'.
- Copy settings from an existing bucket**: A section with a dropdown menu showing 'Select bucket (optional)' and '1 Buckets'.

At the bottom of the wizard are three buttons: 'Create' (white), 'Cancel' (white), and 'Next' (blue).

3. If you like, you can set versioning, logging and tags to this bucket. I didn't because i'm only testing it so click Next.

The screenshot shows the 'Create bucket' wizard in the AWS Management Console, specifically the 'Set properties' step. The wizard has four steps: 1. Name and region (completed), 2. Set properties (current), 3. Set permissions, and 4. Review. The 'Set properties' step contains three optional settings, each with a 'Learn more' link and a radio button to enable or disable the feature. All three settings (Versioning, Logging, and Tags) are currently disabled. At the bottom right, there are 'Previous' and 'Next' buttons, with 'Next' being the active button.

Property	Description	Learn more	Status
Versioning	Keep multiple versions of an object in the same bucket.	<a href="#">Learn more</a>	Disabled
Logging	Set up access log records that provide details about access requests.	<a href="#">Learn more</a>	Disabled
Tags	Use tags to track your cost against projects or other criteria.	<a href="#">Learn more</a>	0 Tags

4. Manage users and permissions if needed and click Next.

The screenshot shows a 'Create bucket' wizard with four steps: 1. Name and region, 2. Set properties, 3. Set permissions, and 4. Review. Step 3 is currently active. The main area of the wizard is dark blue and contains three expandable sections: 'Manage users', 'Manage public permissions', and 'Manage system permissions'. At the bottom right, there are two buttons: 'Previous' and 'Next'.

Create bucket

1 Name and region 2 Set properties 3 Set permissions 4 Review

- Manage users
- Manage public permissions
- Manage system permissions

Previous Next

5. Review the settings and click Create.

Create bucket

✓ Name and region

✓ Set properties

✓ Set permissions

4 Review

Name and region

Bucket name

sw-vtl-replication

Region

EU (Frankfurt)

Edit

Properties

Versioning

Disabled

Logging

Disabled

Tagging

0 Tags

Edit

Permissions

Users

1

Public permissions

Disabled

System permissions

Disabled

Edit

Previous

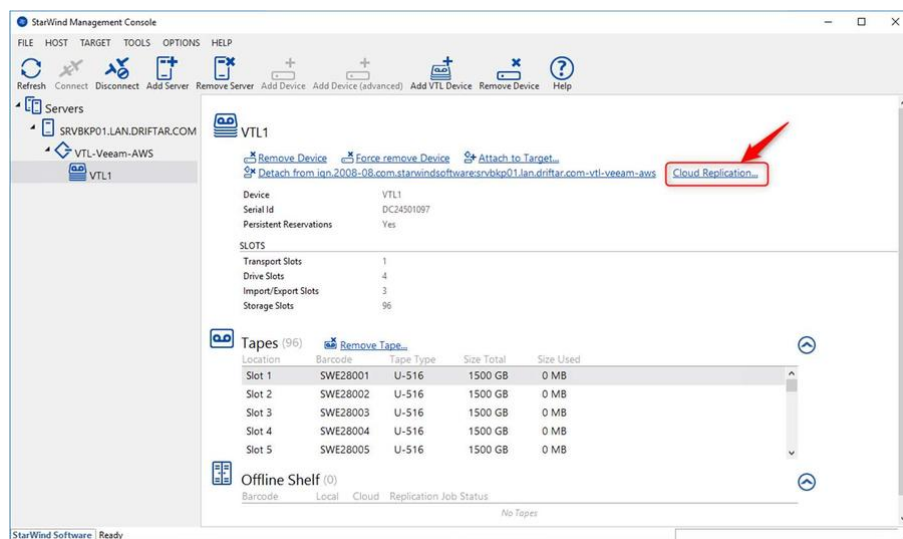
Create bucket

This Amazon S3 bucket is now ready for usage.

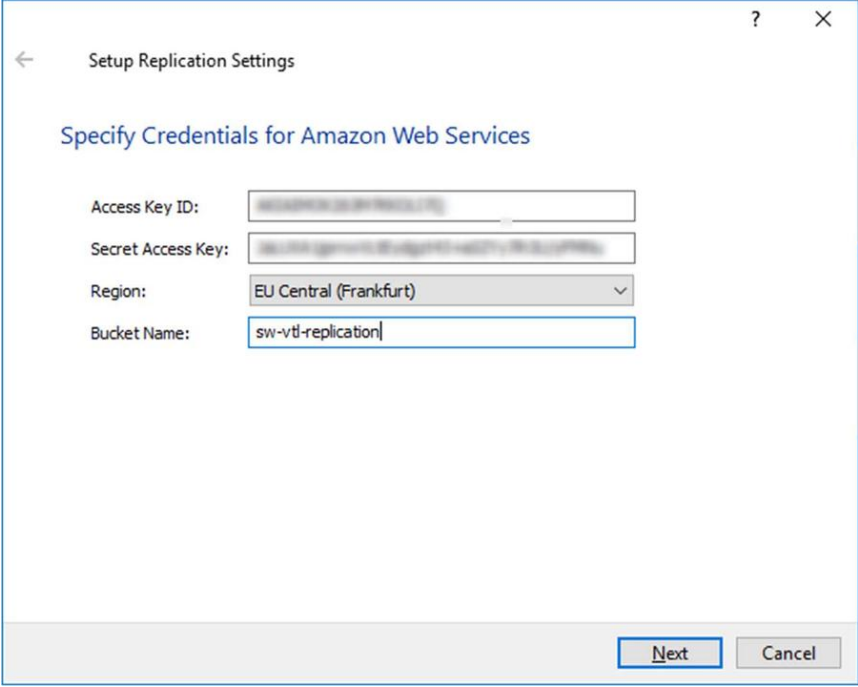
## Setup Cloud Replication

Now that our Amazon S3 bucket, we have to setup cloud replication within StarWind Cloud VTL. Follow the steps below to complete this task.

1. In the StarWind management console click Cloud Replication.



2. Enter your Amazon Access Key ID and the Secret Access Key. Specify the region and enter the bucket name and click Next.



Setup Replication Settings

Specify Credentials for Amazon Web Services

Access Key ID:

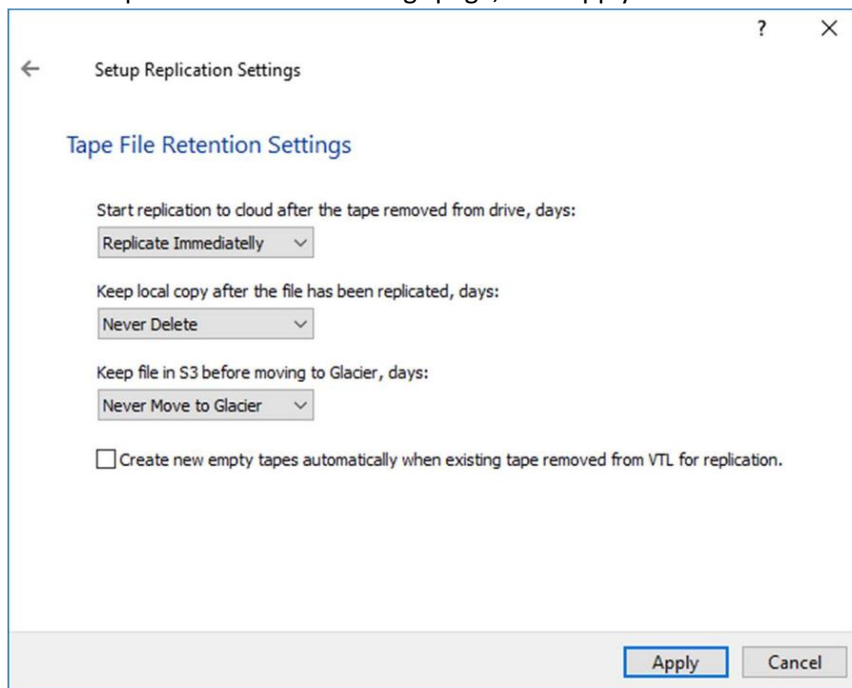
Secret Access Key:

Region:

Bucket Name:

Next Cancel

3. On the Tape File Retention Settings page, click Apply.



Setup Replication Settings

Tape File Retention Settings

Start replication to cloud after the tape removed from drive, days:  
Replicate Immediately

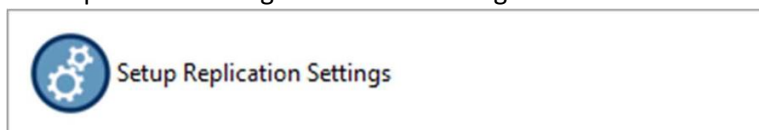
Keep local copy after the file has been replicated, days:  
Never Delete

Keep file in S3 before moving to Glacier, days:  
Never Move to Glacier

☐ Create new empty tapes automatically when existing tape removed from VTL for replication.

Apply Cancel

4. The replication settings will now be configured.

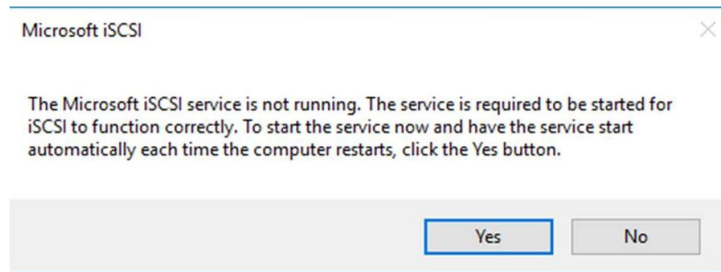


Now we've got:

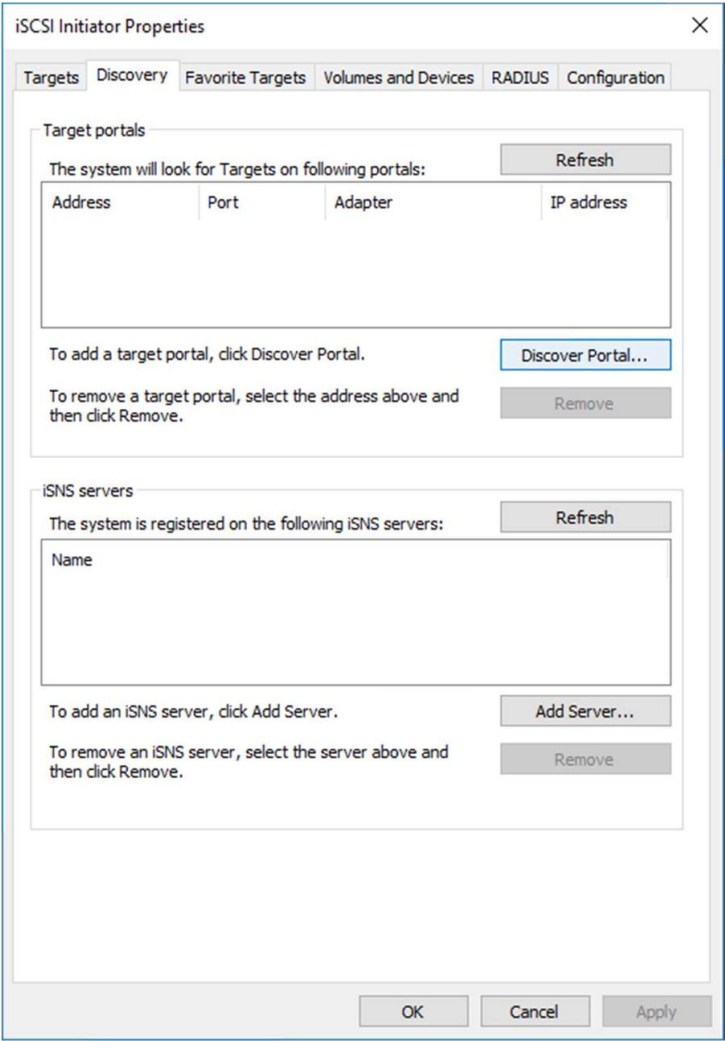
- Veeam Backup & Replication installed
- StarWind Cloud VTL installed
- Amazon S3 bucket created
- Cloud Replication configured

## Configure Microsoft iSCSI Initiator

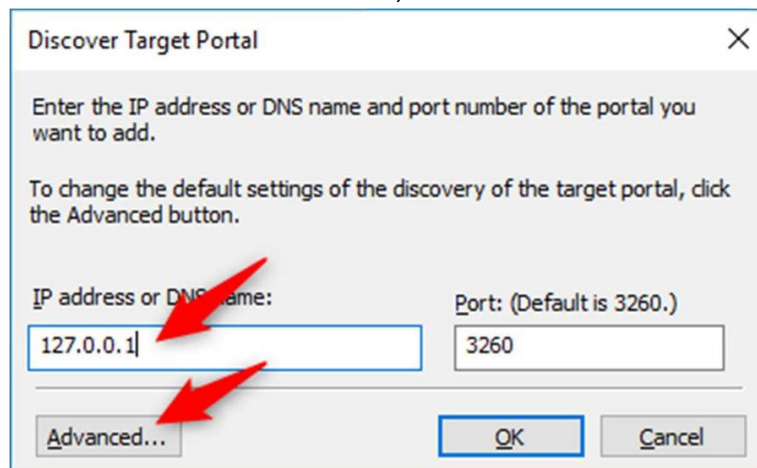
1. As I mentioned before, this VTL connects through iSCSI. Open Microsoft iSCSI Initiator. If it's not already in use, Windows tells you to start the service.



2. Navigate to the Discovery tab and click the Discover Portal button.



3. As we're still on the same server, enter the localhost address 127.0.0.1 and then click OK.



4. Select Microsoft iSCSI initiator for Local Adapter and click OK.

The screenshot shows the 'Advanced Settings' dialog box with the 'IPsec' tab selected. The 'Connect using' section is highlighted with a red rectangle, showing 'Local adapter:' set to 'Microsoft iSCSI Initiator'. Below this, 'Initiator IP:' is set to 'Default' and 'Target portal IP:' is empty. The 'CRC / Checksum' section has 'Data digest' and 'Header digest' checkboxes, both unchecked. The 'Enable CHAP log on' checkbox is also unchecked. The 'CHAP Log on information' section contains a text box for 'Name:' with the value 'iqn.1991-05.com.microsoft:svbtkp01.lan.driftnet.com' and an empty 'Target secret:' text box. At the bottom, there are three checkboxes: 'Perform mutual authentication' (unchecked), 'Use RADIUS to generate user authentication credentials' (unchecked), and 'Use RADIUS to authenticate target credentials' (unchecked). The 'OK', 'Cancel', and 'Apply' buttons are at the bottom right.

Advanced Settings

General IPsec

Connect using

Local adapter: Microsoft iSCSI Initiator

Initiator IP: Default

Target portal IP:

CRC / Checksum

☐ Data digest ☐ Header digest

☐ Enable CHAP log on

CHAP Log on information

CHAP helps ensure connection security by providing authentication between a target and an initiator.

To use, specify the same name and CHAP secret that was configured on the target for this initiator. The name will default to the Initiator Name of the system unless another name is specified.

Name: iqn.1991-05.com.microsoft:svbtkp01.lan.driftnet.com

Target secret:

☐ Perform mutual authentication

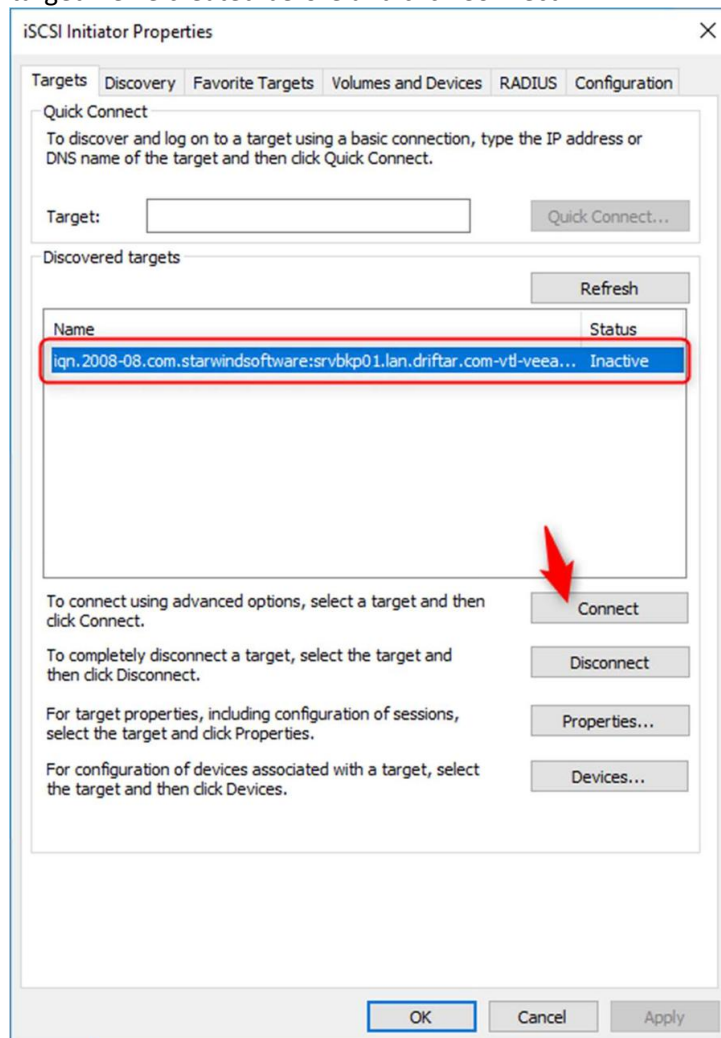
To use mutual CHAP, either specify an initiator secret on the Configuration page or use RADIUS.

☐ Use RADIUS to generate user authentication credentials

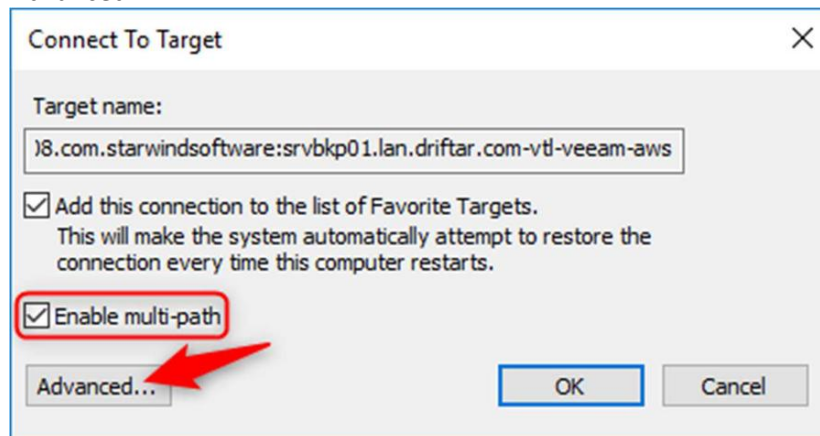
☐ Use RADIUS to authenticate target credentials

OK Cancel Apply

5. On the iSCSI Initiator Page, navigate to the Targets tab. You should see here the iSCSI target we've created before and click Connect.



6. On the Connect to Target Page, Select the Enable multi-path checkbox and click Advanced.



7. On the Advanced Page, Set the local adapter to Microsoft iSCSI initiator and set 127.0.0.1 / 3260 as target portal and click OK.

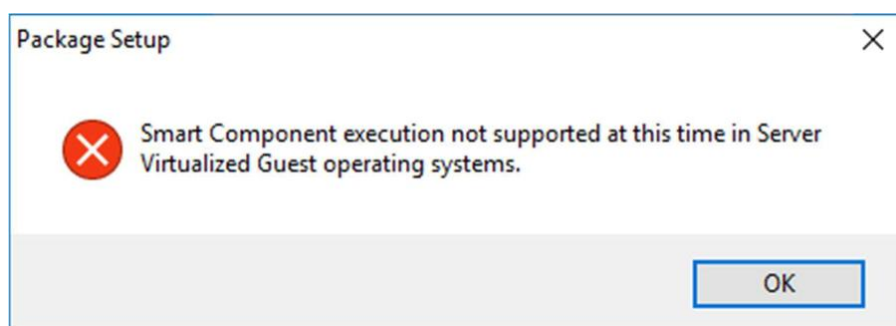
The screenshot shows the 'Advanced Settings' dialog box for the Microsoft iSCSI Initiator. The 'General' tab is selected. Under the 'Connect using' section, three fields are highlighted with red rectangles: 'Local adapter' set to 'Microsoft iSCSI Initiator', 'Initiator IP' set to 'Default', and 'Target portal IP' set to '127.0.0.1 / 3260'. Below this, there are checkboxes for 'CRC / Checksum' (Data digest and Header digest), 'Enable CHAP log on', and 'CHAP Log on information'. The 'Name' field is filled with 'iqn.1991-05.com.microsoft:svb01.lan.drifstar.com'. There are also checkboxes for 'Perform mutual authentication', 'Use RADIUS to generate user authentication credentials', and 'Use RADIUS to authenticate target credentials'. At the bottom, there are 'OK', 'Cancel', and 'Apply' buttons.

8. After completing these steps, You should notice that VTL iSCSI target is listed as "Connected" in the list. Close the Microsoft iSCSI initiator assistant.

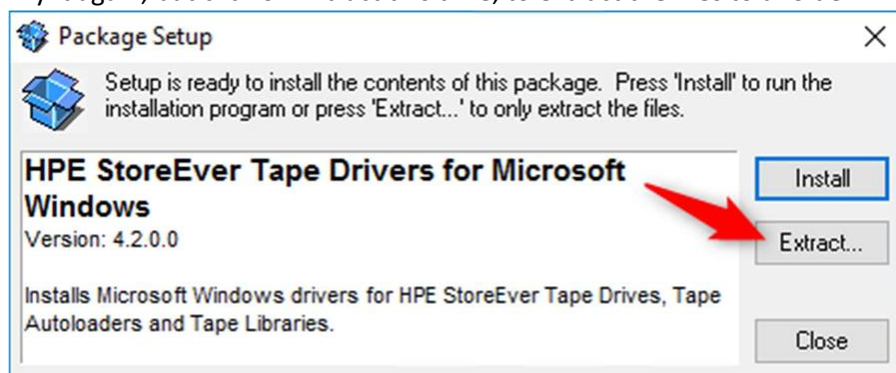
## Installing HPE StoreEver Tape Drivers

It is necessary to install the HPE StoreEver tape drives to make this solution work. If you do not complete these steps, it is not possible to get this VTL to connect to Veeam. It is an easy setup and let's do that now.

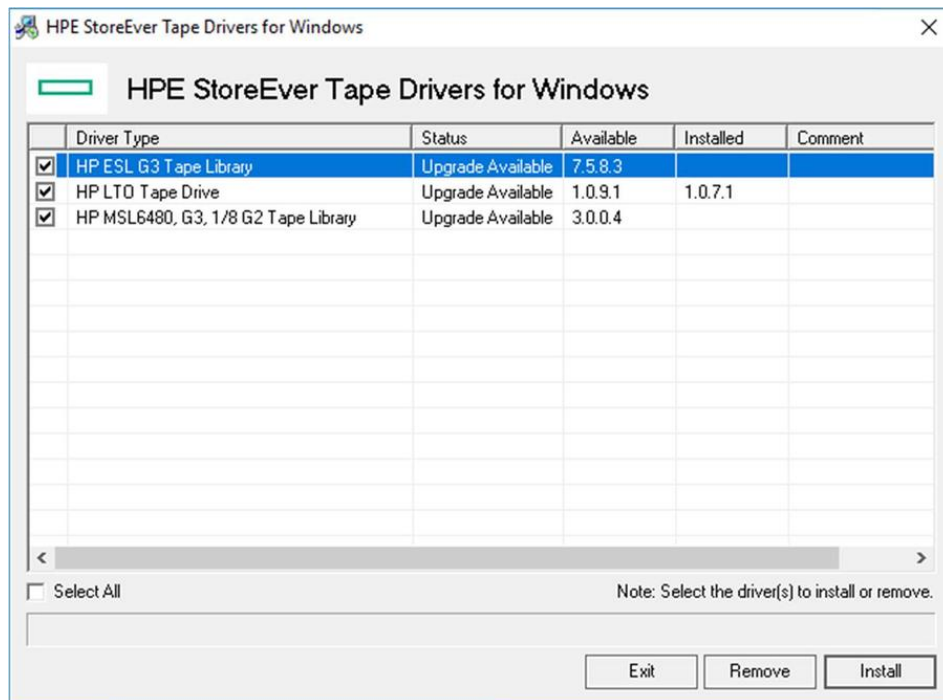
1. If you try to install the HPE StoreEver Tape Drivers on your virtual server by double-clicking the software package, you will notice an error. It's not possible to install the drivers on this way on a virtual server. To work around this error continue with step 2.



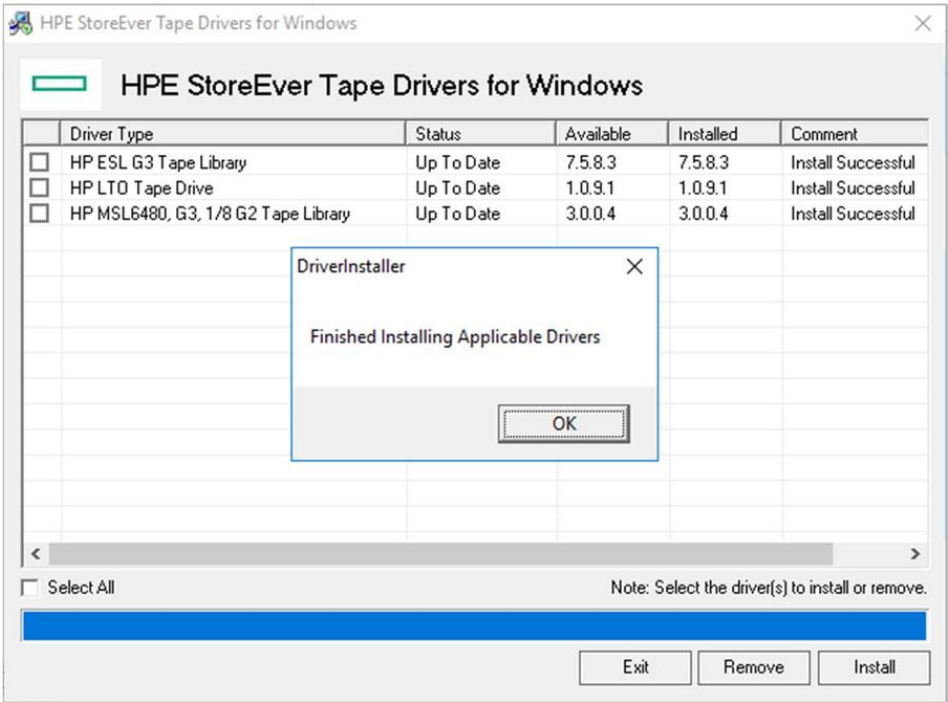
2. Try it again, but click on Extract this time, to extract the files to a folder.



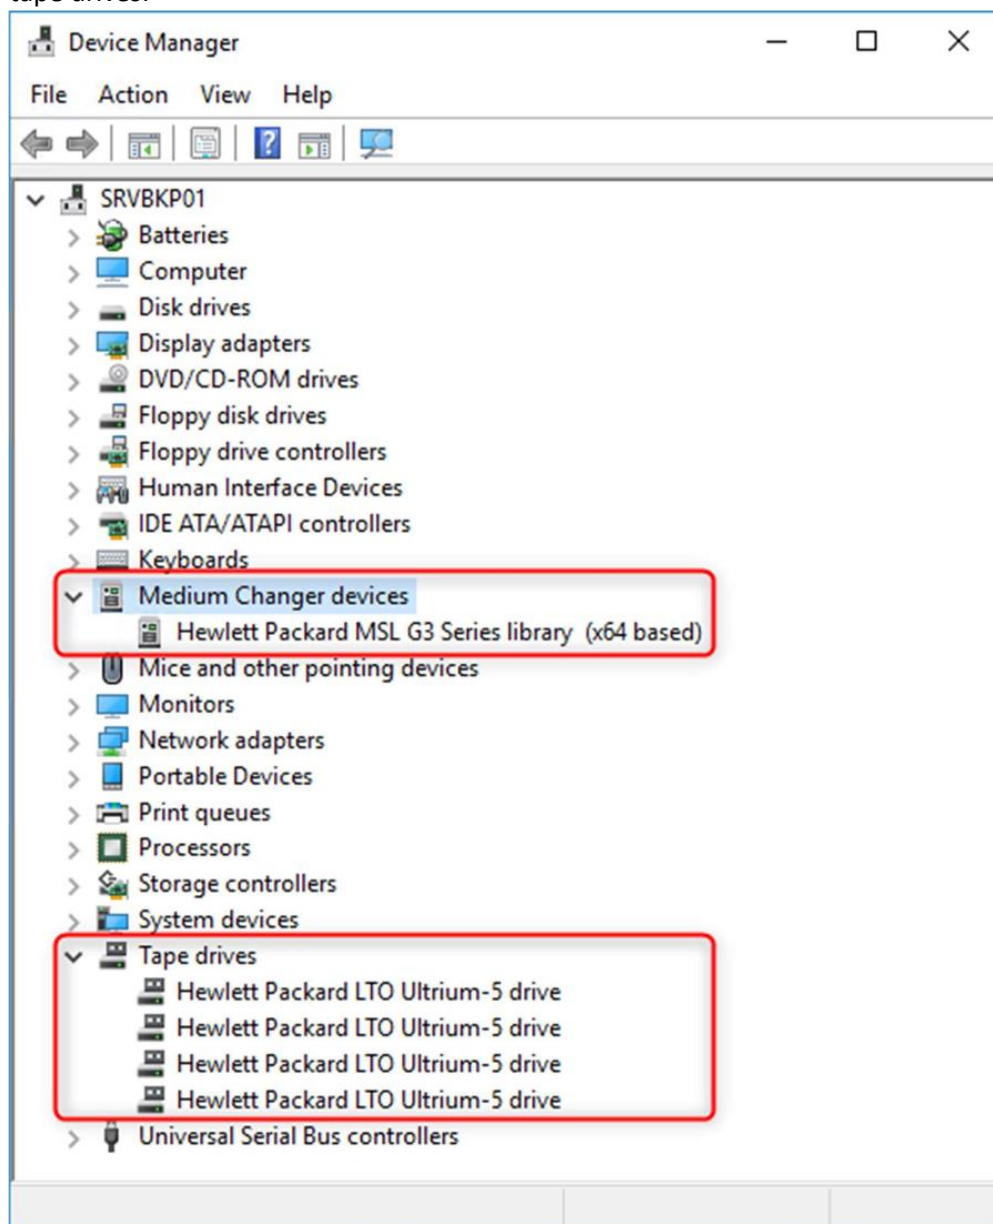
3. Navigate to this folder and start the setup. Make sure all checkboxes are selected and click Install.



4. On the HP StoreEver Tape Drivers for Windows page, click OK to close the confirmation popup. Then click Close to close the setup assistant.R



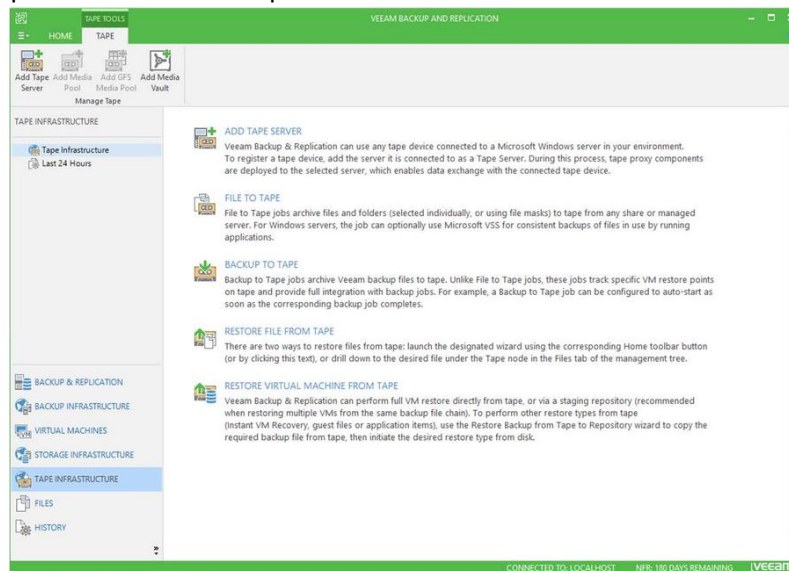
5. Open Device Manager and review the settings. Notice the HPE tape library and some tape drives.



The next step is to setup the tape library in Veeam Backup & Replication and to configure a tape job.


## Add the Virtual Tape Library to Veeam

1. Open the Veeam Backup & Replication console. Click Tape Infrastructure on the left pane and click Add Tape Server.



2. On the New Tape Server Page, Click Next.

**New Tape Server**

 **Server**

Choose a server to install tape server components on. You can only select between Microsoft Windows servers added to the managed servers tree in the console.

**Server**

Traffic

Review

Apply

Summary

Choose server:

SRVBKP01.lan.drifstar.com

Add New...

Description:

Starwind VTL Veeam to AWS

< Previous

**Next >**

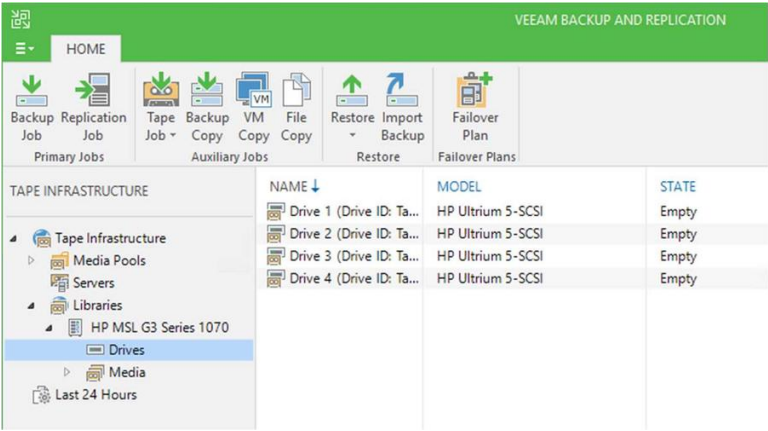
Finish

Cancel

3. Follow through the setup assistant. I didn't mention the traffic rules. You can setup traffic rules if you need to. Click Finish to close the assistant.

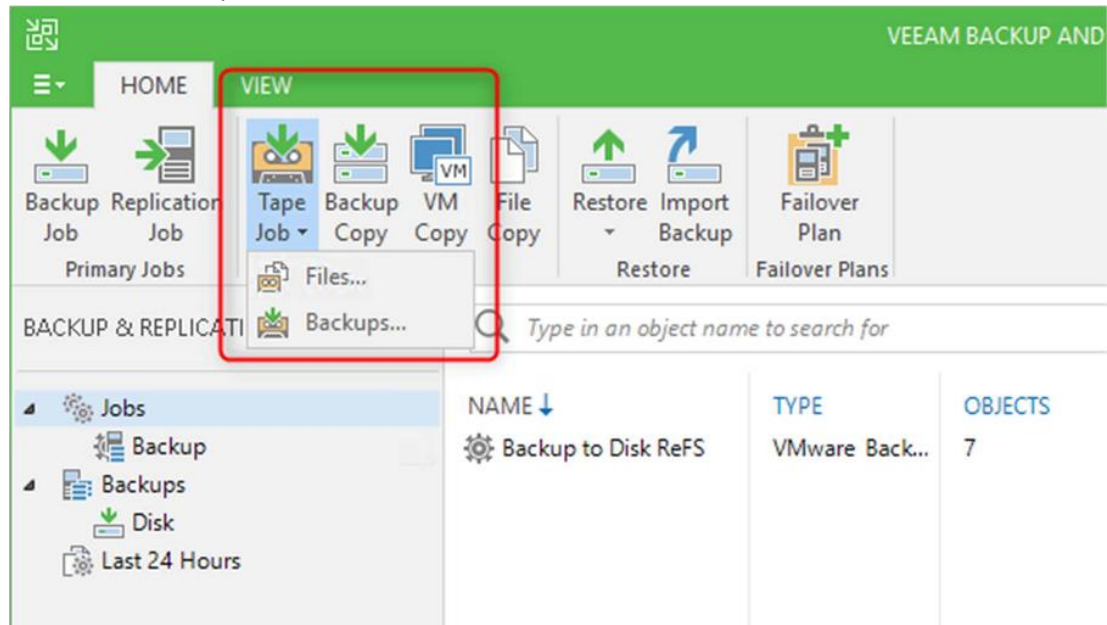
[illegible]

4. The setup should look like this. You see the tape library with it's drives.




## Configure a Tape backup Job

1. The next step is to configure a Tape Job for Backups. In the Ribbon select Tape Job, and then select Backups.



2. On the New File to Tape Job Page, set a name and optionally a description and click Next.

New File to Tape Job

 **Name**  
Type in a name and description for this job. File to tape job performs scheduled backup of selected files to tape.

**Name**

Files and Folders

Full Backup

Incremental Backup

Options

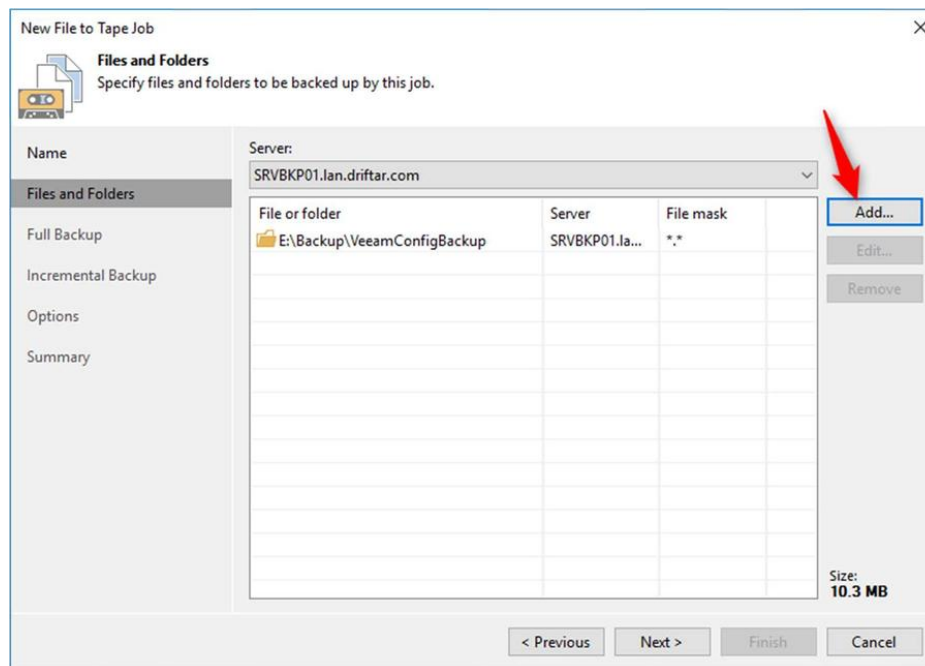
Summary

**Name:**  
Config Backup to Tape

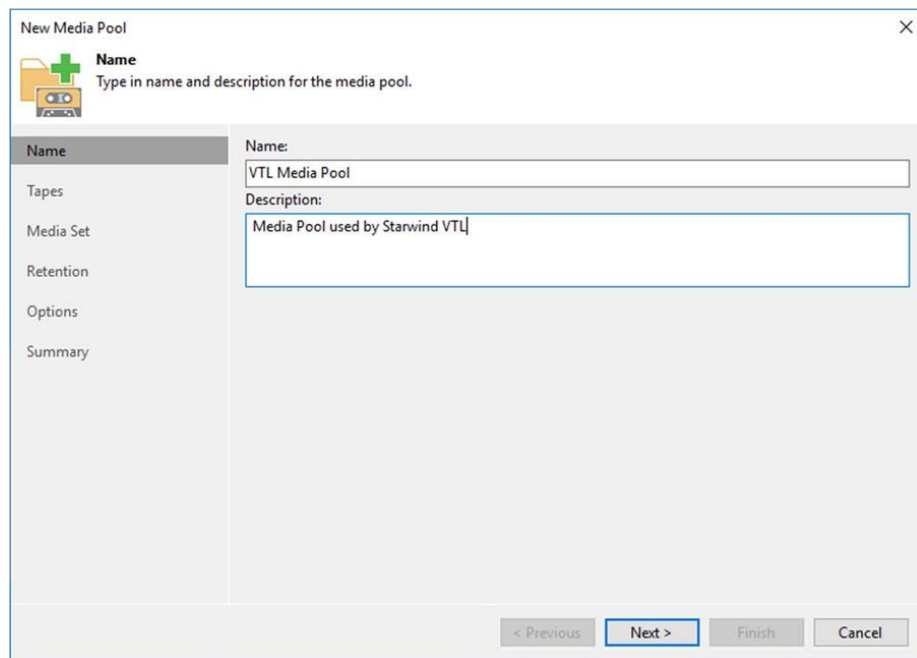
**Description:**  
Copies Veeam Config Backup to Tape

< Previous **Next >** Finish Cancel

3. On the Files and Folders Page, specify the files and folders you want to have on tape and click Next.




4. On the new Media Pool Page, Type a name and click Next.



The screenshot shows a 'New Media Pool' wizard window. The window has a title bar with a close button (X). Below the title bar is a header area with a folder icon and a green plus sign, followed by the word 'Name' and the instruction 'Type in name and description for the media pool.' On the left side, there is a vertical list of steps: 'Name', 'Tapes', 'Media Set', 'Retention', 'Options', and 'Summary'. The 'Name' step is currently selected and highlighted. The main area of the window contains two text input fields. The first field is labeled 'Name:' and contains the text 'VTL Media Pool'. The second field is labeled 'Description:' and contains the text 'Media Pool used by Starwind VTL'. At the bottom of the window, there are four buttons: '< Previous', 'Next >', 'Finish', and 'Cancel'. The 'Next >' button is highlighted with a blue border.

5. On the Tapes Page, click Next.

New Media Pool

 **Tapes**  
Add tapes to the media pool.

**Name**

**Tapes**

Media Set

Retention

Options

Summary

Tape library:  
HP MSL G3 Series 1070 (SRVBKP01) Manage...

Tapes:

Name	Capacity	Remaining
------	----------	-----------

Add...  
Remove

Capacity:  
**0.0 B**  
Remaining:  
**0.0 B**

☒ Add tapes from Free media pool automatically when more tapes are required

< Previous Next > Finish Cancel

6. On the Media Set Page, leave the defaults and click Next.

New Media Pool

**Media Set**  
Specify media set name and how often a new media set should be automatically created.

**Name**  
Media set name:

**Tapes**

**Media Set**

**Retention**

**Options**

**Summary**

Automatically create new media set

☒ Do not create, always continue using current media set

☐ Create new media set for every backup session

☐ Daily at

< Previous   Next >   Finish   Cancel

7. On the Retention Page, leave the defaults and click Next.

New Media Pool

Retention

Specify the tape retention settings for this media pool.

Name

Tapes

Media Set

Retention

Options

Summary

Data retention policy

☐ Do not protect data (cyclically overwrite tapes as required)

☐ Protect data for 

1

Weeks

☒ Never overwrite data

Offline media tracking

☐ Move all offline tapes from this media pool into the following media vault:

Vault:

Add New...

< Previous

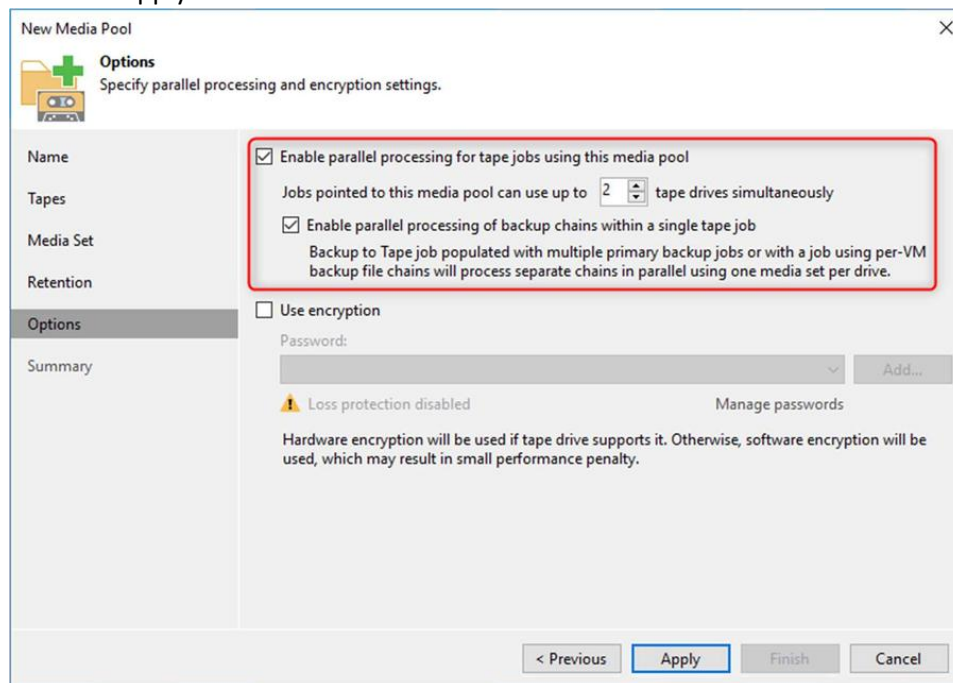
Next >

Finish

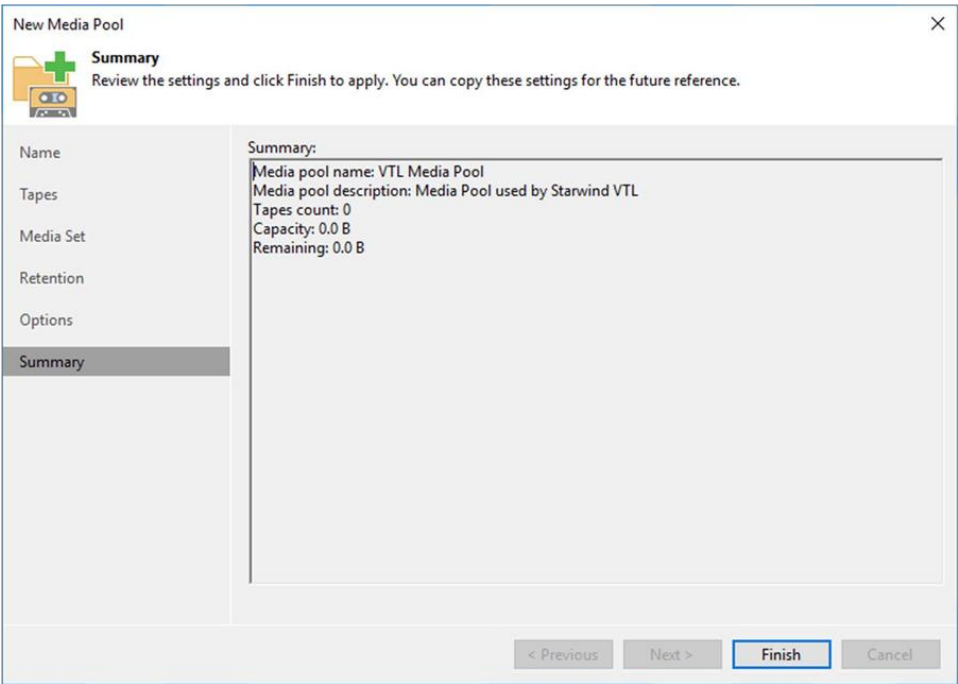
Cancel

185

8. On the Options Page, enable parallel processing (we've got more than one tape drive) and click Apply.



9. On the Summary Page, review the settings and click Finish.



10. On the Files to Tape Job Full Backup Page, define your schedule and click Next.

New File to Tape Job

**Full Backup**  
Choose media pool and set schedule for full backups.

**Name**

**Files and Folders**

**Full Backup**

Incremental Backup

Options

Summary

Media pool for full backups:  
VTL Media Pool (HP MSL G3 Series 1070) Add New...

Tapes: 0  
Capacity: 0.0 B  
Remaining: 0.0 B

☒ Run the full backup automatically

☒ Daily at this time: 02:00 On these days Days...

☐ Monthly at this time: 22:00 Fourth Saturday Months...

< Previous Next > Finish Cancel

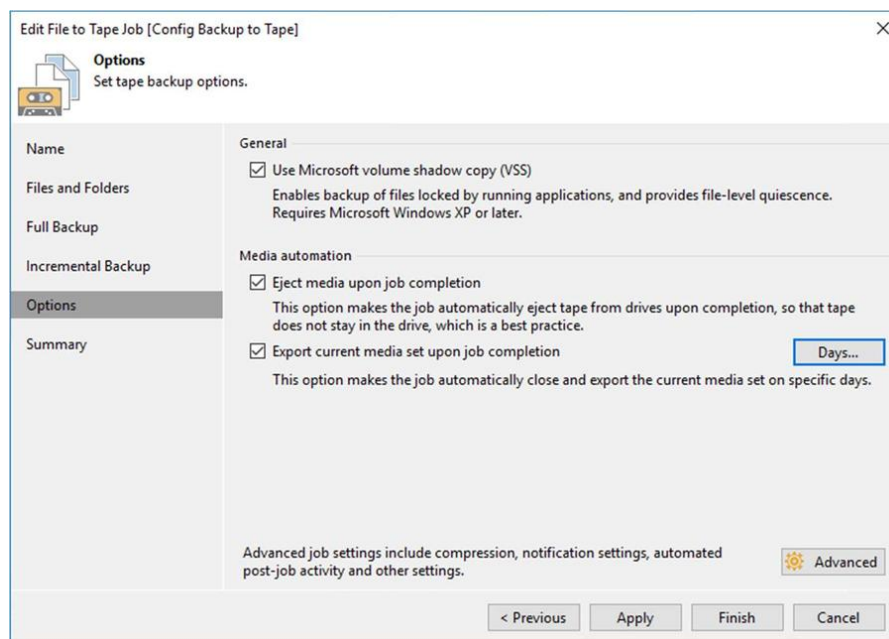
11. On the Files to Tape Job Incremental Backup Page, leave the defaults and click Next.

The screenshot shows the 'New File to Tape Job' dialog box with the 'Incremental Backup' tab selected. The dialog has a sidebar on the left with options: Name, Files and Folders, Full Backup, Incremental Backup (selected), Options, and Summary. The main area is titled 'Incremental Backup' and contains the following settings:

- Media pool for incremental backup:** A dropdown menu showing 'VTL Media Pool (HP MSL G3 Series 1070)' with an 'Add New...' button to its right.
- Tapes:** 0
- Capacity:** 0.0 B
- Remaining:** 0.0 B
- Run incremental backup automatically:** An unchecked checkbox.
- Daily at this time:** A radio button selected, with a time dropdown set to '18:00' and a frequency dropdown set to 'On week-days' with a 'Days...' button.
- Monthly at this time:** An unselected radio button, with a time dropdown set to '22:00', a frequency dropdown set to 'Fourth', a day dropdown set to 'Saturday', and a 'Months...' button.

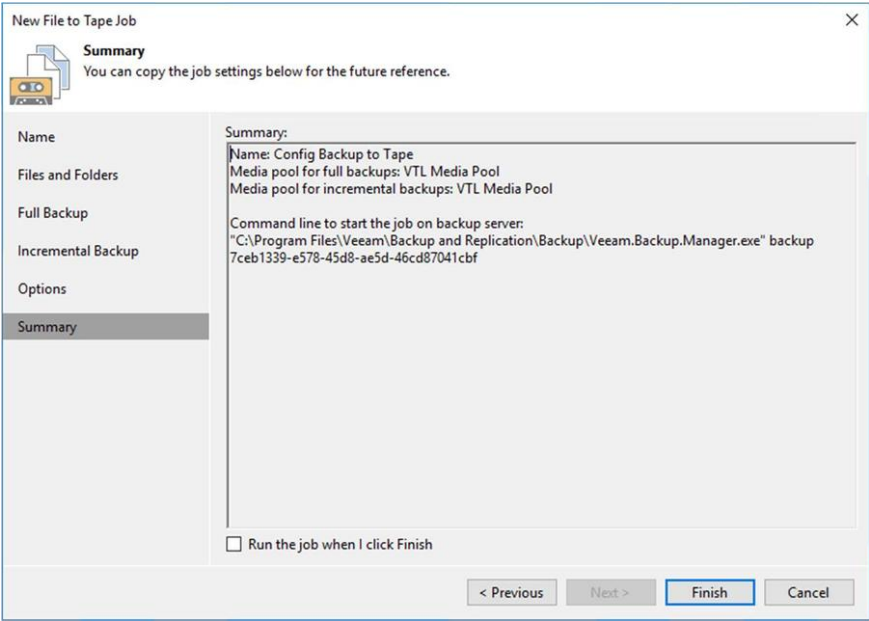
At the bottom of the dialog are four buttons: '< Previous', 'Next >' (highlighted with a blue border), 'Finish', and 'Cancel'.

12. On the Options Page, Select Use Microsoft Volume Shadow Copy (VSS), Select Eject media upon job completion, Select Export current media set upon job completion and then click Apply.



**Real World Note:** Enable to export current media set upon job completion. Otherwise, StarWind Cloud VTL will not replicate to AWS.

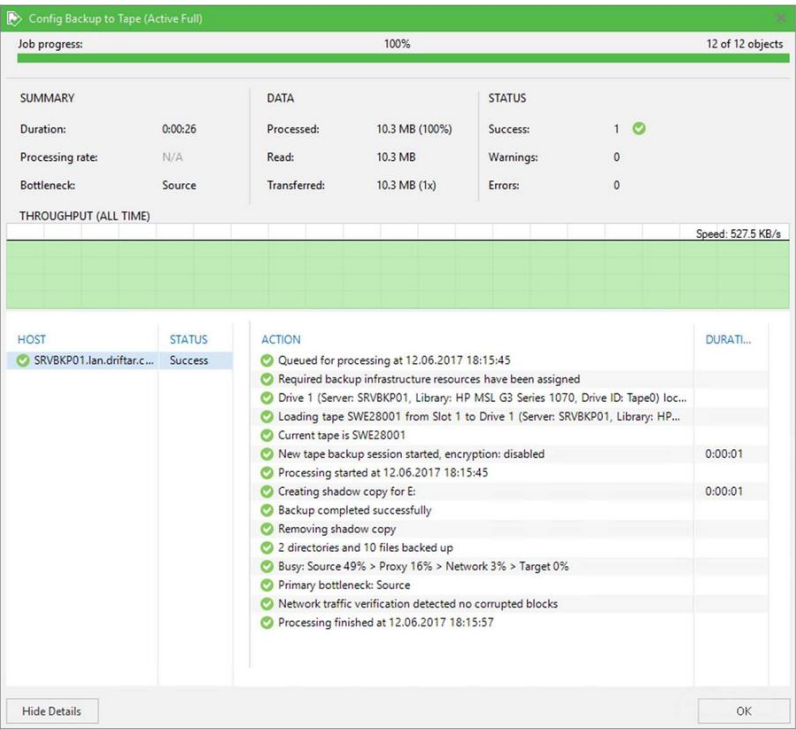
13. Review the job settings and click Finish.



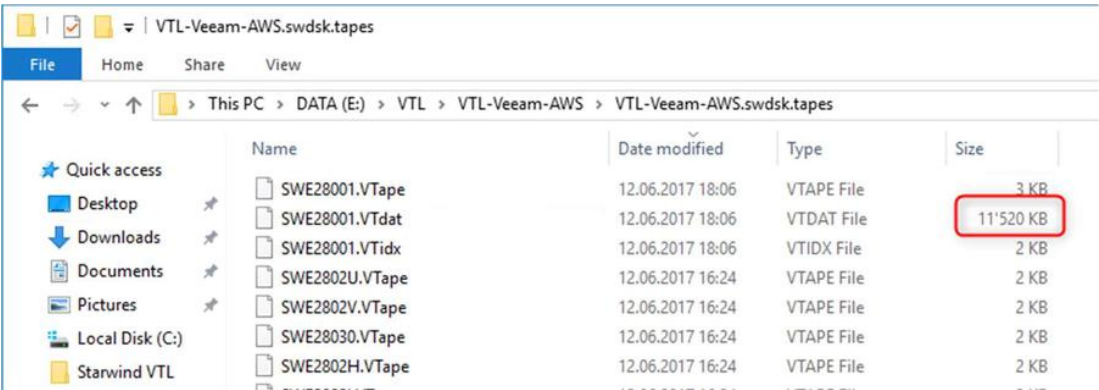
That’s it! You finally made it! Congratulations

As the last thing, let’s look into the running job.

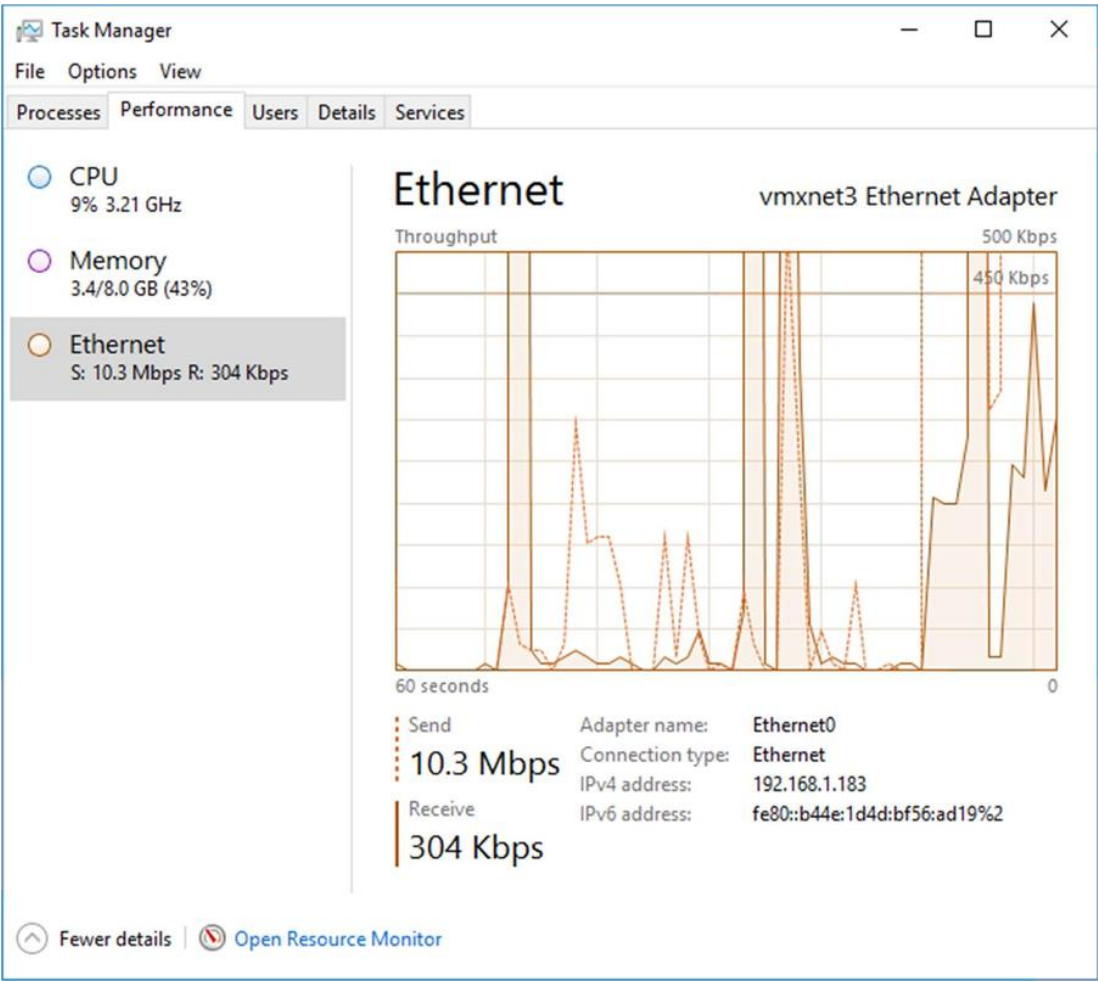
You can see that Veeam assigns our virtual tape library and loads the tape.



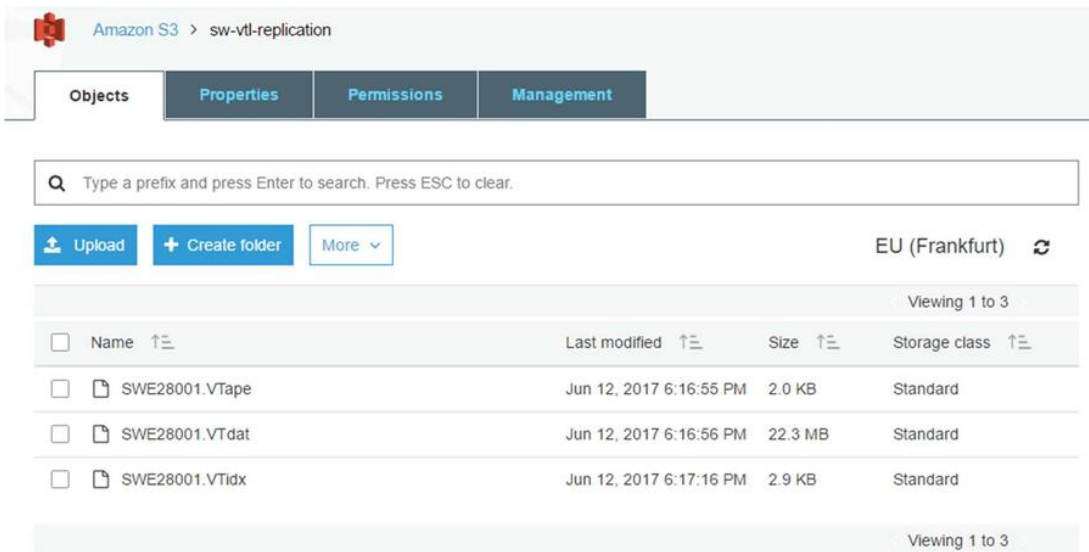
If we look at the VTL folder in Windows Explorer, we can see the size has increased.



When we look at the task manager and select network performance, we can see some spikes. I don't have a lot of internet bandwidth, and my lab is connected through DLAN (powerline), so again a loss of bandwidth.



When we now finally look into our newly created Amazon S3 bucket, we can also see that there is some traffic.



Amazon S3 > sw-vtl-replication

Objects Properties Permissions Management

Q Type a prefix and press Enter to search. Press ESC to clear.

Upload Create folder More

EU (Frankfurt)

	Name	Last modified	Size	Storage class
<input type="checkbox"/>	SWE28001.VTape	Jun 12, 2017 6:16:55 PM	2.0 KB	Standard
<input type="checkbox"/>	SWE28001.VTdat	Jun 12, 2017 6:16:56 PM	22.3 MB	Standard
<input type="checkbox"/>	SWE28001.VTidx	Jun 12, 2017 6:17:16 PM	2.9 KB	Standard

## Conclusion

With StarWind Cloud VTL for AWS and Veeam, you can leverage your backup data and airgap it to the cloud. Cloud storages prices start from approximately 0.002 US \$ per gigabyte. The more data you backup to the cloud, the cheaper it gets. With solutions like the one in this chapter, you can save money because you don't have to invest it in real hardware, energy, and cooling.

I hope that this guide will help you with the installation and configuration of StarWind Cloud VTL for AWS and Veeam.

## Chapter 11

# Configure Azure Direct Restore with Veeam 9.5

**By: Karl Widmer (VMWare vEXPERT / Veeam Vanguard)**

Leveraging Veeam Direct Restore to Azure can be a very beneficial feature for customers looking to have a cloud-based DR Solution at a reduced cost. This Chapter will walk through the setting required to configure this feature including the pre-requisite setup.

### Requirements

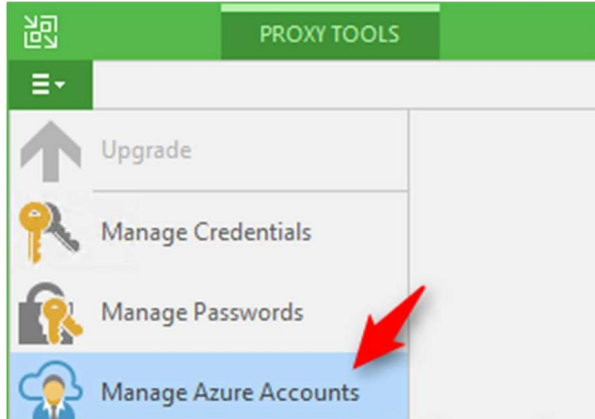
- Veeam Backup & Replication 9.5
  - Supported in Standard, Enterprise and Enterprise Plus
- Microsoft Azure Account
  - Pay-as-you-go or any other subscription based account
- Azure Storage (blob or general storage)
- Azure Virtual Network

**Note:** I recommend that you organize your lab configuration into a single Resource Group in Azure. This way you will have a very easy way to visualize what has been configured in the Azure Portal.

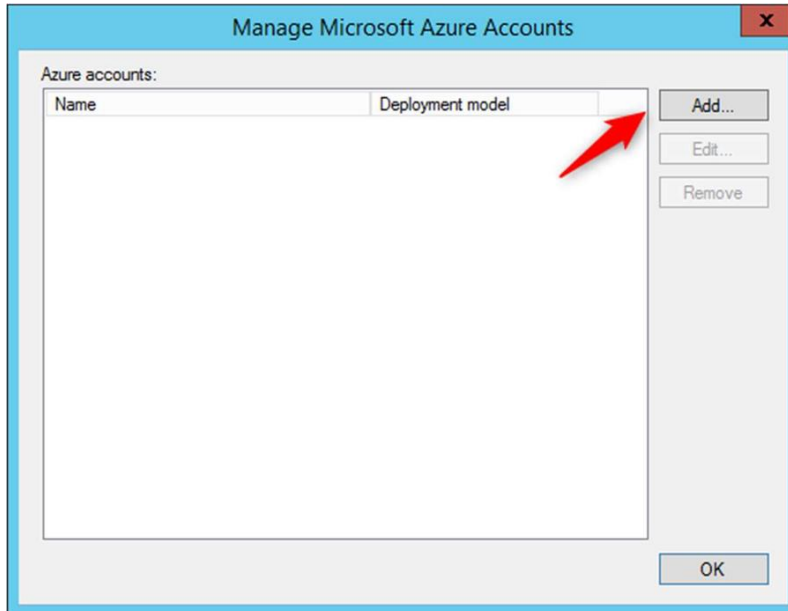
## Setup your Azure Proxy VM

The first step in configuring Direct Restore to Azure is setting up a Proxy Virtual Machine (VM).

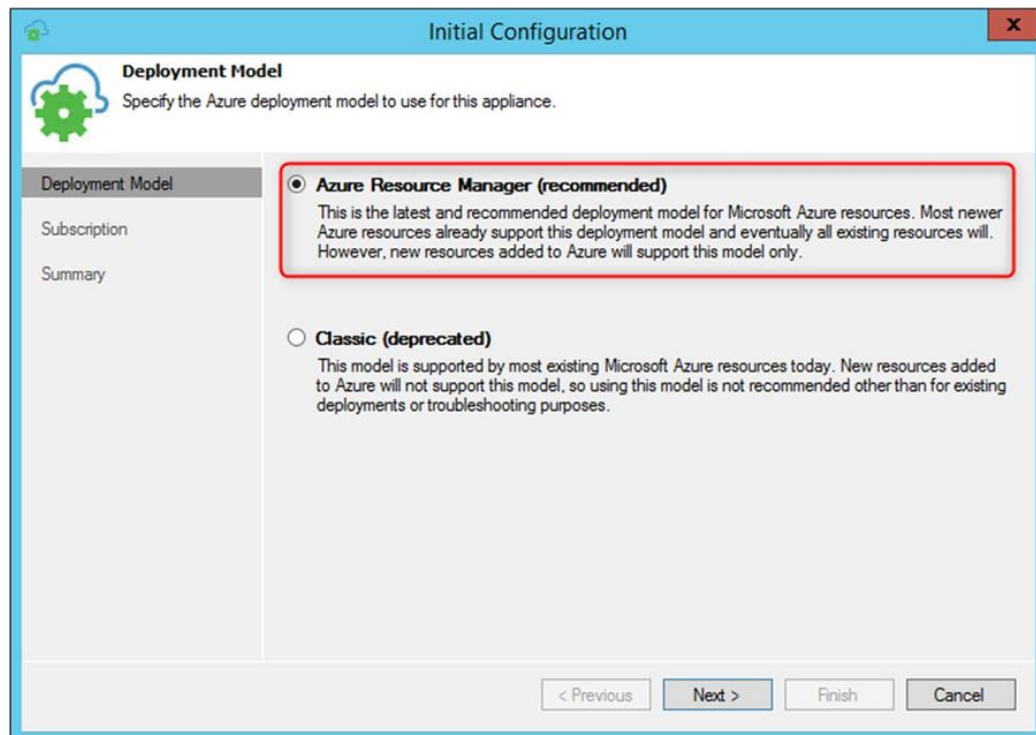
1. Open your Veeam Backup & Replication console.
2. Click on the Options button top left, then click “Manage Azure Accounts.”



3. On the Manage Microsoft Azure Accounts Page, click Add.



4. On the Deployment Model Page, Select Azure Resource Manager (Recommended), and click Next.



5. On the Subscription Page, specify your Azure account to access the Azure Resource Manager and click Next.

The screenshot shows the 'Initial Configuration' window of the Veeam backup software. The title bar is blue with the text 'Initial Configuration' and a close button. The main window has a light blue header with a gear icon and the text 'Subscription' and 'Import your Azure Resource Manager subscriptions.' Below the header is a sidebar with three items: 'Deployment Model', 'Subscription' (which is highlighted), and 'Summary'. The main content area is divided into two sections. The top section is titled 'Specify account to access your Azure Resource Manager subscriptions' and contains a text box labeled 'Account:' with a 'Configure account' link to its right. The bottom section is titled 'Linux support' and contains a checkbox labeled 'Enable restore of Linux-based computers' with the text 'You will need to provide the helper appliance settings on the next step of the wizard.' at the bottom. At the very bottom of the window are four buttons: '< Previous', 'Next >', 'Finish', and 'Cancel'.

Initial Configuration

**Subscription**  
Import your Azure Resource Manager subscriptions.

Deployment Model

**Subscription**

Summary

Specify account to access your Azure Resource Manager subscriptions

Account:

[Configure account](#)

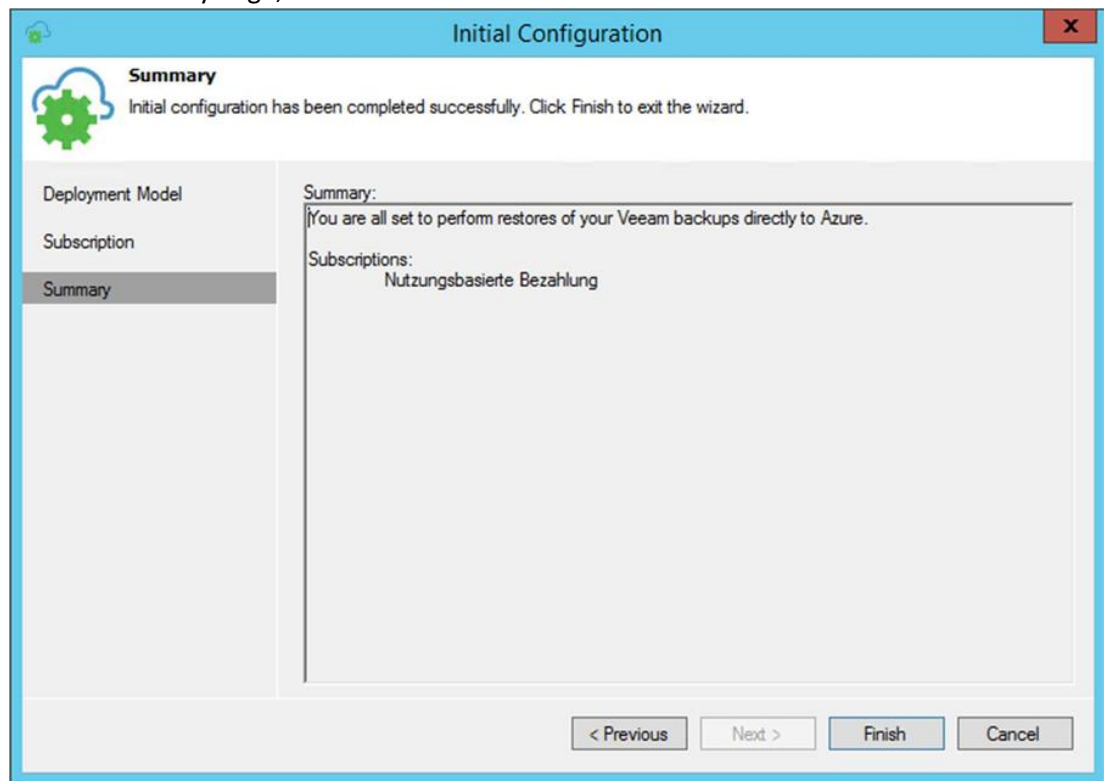
Linux support

☐ Enable restore of Linux-based computers

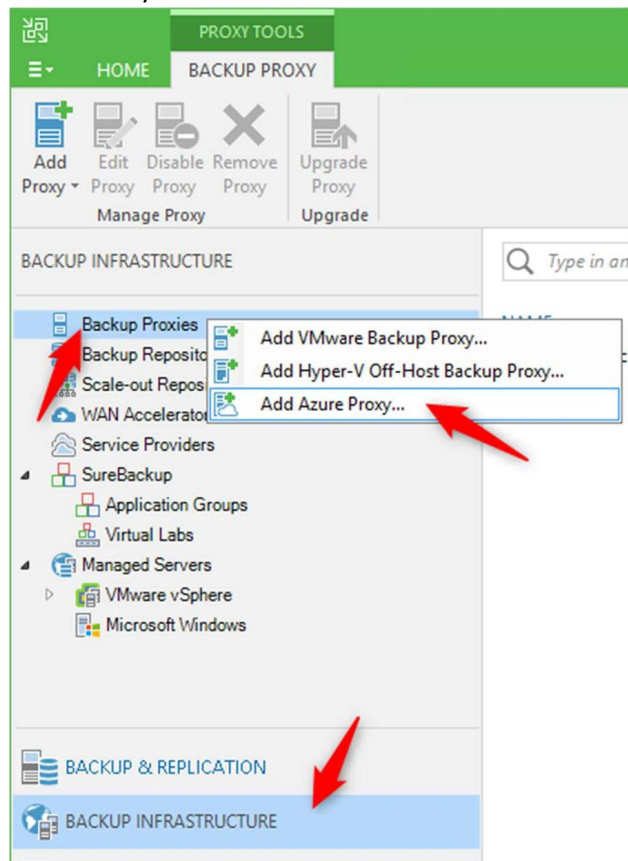
You will need to provide the helper appliance settings on the next step of the wizard.

< Previous   Next >   Finish   Cancel

6. On the Summary Page, click Finish.



7. In the Veeam Console, click Backup Infrastructure, click Backup Proxies, and click Add Azure Proxy.



8. On the Add Azure Proxy Name Page, Type in a name and click Next.

The screenshot shows the 'Add Azure Proxy' wizard window. The title bar is blue with a close button. The main area has a light blue header with a green plus icon and the text 'Name'. Below this, it says 'Type in a name and description. The specified name will be used as the Azure VM name and the computer name.' The left sidebar contains a list of steps: Name, Credentials, Deployment Model, Subscription, VM Size, Resource Group, Network, Apply, and Summary. The 'Name' step is selected. The main content area has two text boxes: 'Name:' with the value 'AzureWestEurope' and 'Description:' with the value 'Created by DRIFTAR\administrator at 08.12.2016 21:47.'. At the bottom, there is a 'Max concurrent tasks:' label and a spinner box set to '4'. The bottom right corner has four buttons: '< Previous', 'Next >', 'Finish', and 'Cancel'.

Note: No special characters allowed, and don't exceed 15 characters.

9. On the Add Azure Proxy Credentials Page, define the credentials for the Azure Proxy and click Next.

**Add Azure Proxy**

**Credentials**  
Specify local administrator account credentials to assign to the Azure VM.

**Name**

**Credentials**

Deployment Model

Subscription

VM Size

Resource Group

Network

Apply

Summary

Specify a local administrator account to create on the Azure proxy VM. The account name must be plain (without domain or machine name). The password must be at least 8 characters long and contain a mix of uppercase and lowercase characters, special symbols and digits.

Credentials: veeam (veeam, last edited: less than a day ago) Add...

[Manage accounts](#)

All management and restore traffic to the Azure proxy will be encapsulated into the single port specified below. The network traffic will also be encrypted, so no VPN to Azure is required.

Traffic port: 6181

< Previous Next > Finish Cancel

Note: You can create new and random credentials. These are used for the deployment of the proxy on Azure.

10. On the Add Azure Proxy Deployment Model Page, select Azure Resource Manager (Recommended), and click Next.

The screenshot shows the 'Add Azure Proxy' wizard window. The title bar is blue with the text 'Add Azure Proxy' and a close button. The main content area has a light blue header with a green plus icon and the text 'Deployment Model' and 'Specify the Microsoft Azure deployment model to use for this proxy VM.' Below this is a list of steps on the left: Name, Credentials, Deployment Model (selected), Subscription, VM Size, Resource Group, Network, Apply, and Summary. The main area shows two radio button options. The first option, 'Azure Resource Manager (recommended)', is selected and highlighted with a red border. It includes a description: 'This is the latest and recommended deployment model for Microsoft Azure resources. Most newer Azure resources already support this deployment model and eventually all existing resources will. However, new resources added to Azure will support this model only.' The second option, 'Classic (deprecated)', is unselected and includes a description: 'This model is supported by most existing Microsoft Azure resources today. New resources added to Azure will not support this model, so using this model is not recommended other than for existing deployments or troubleshooting purposes.' At the bottom right are four buttons: '< Previous', 'Next >', 'Finish', and 'Cancel'.

**Add Azure Proxy**

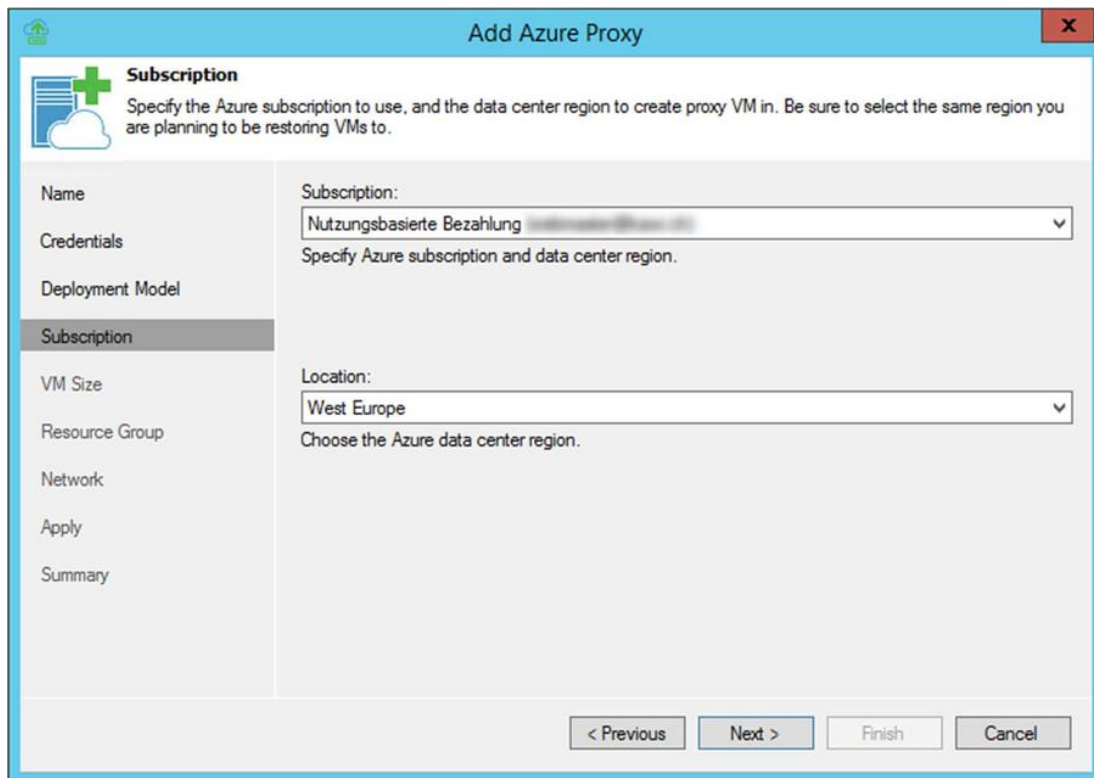
**Deployment Model**  
Specify the Microsoft Azure deployment model to use for this proxy VM.

☒ **Azure Resource Manager (recommended)**  
This is the latest and recommended deployment model for Microsoft Azure resources. Most newer Azure resources already support this deployment model and eventually all existing resources will. However, new resources added to Azure will support this model only.

☐ **Classic (deprecated)**  
This model is supported by most existing Microsoft Azure resources today. New resources added to Azure will not support this model, so using this model is not recommended other than for existing deployments or troubleshooting purposes.

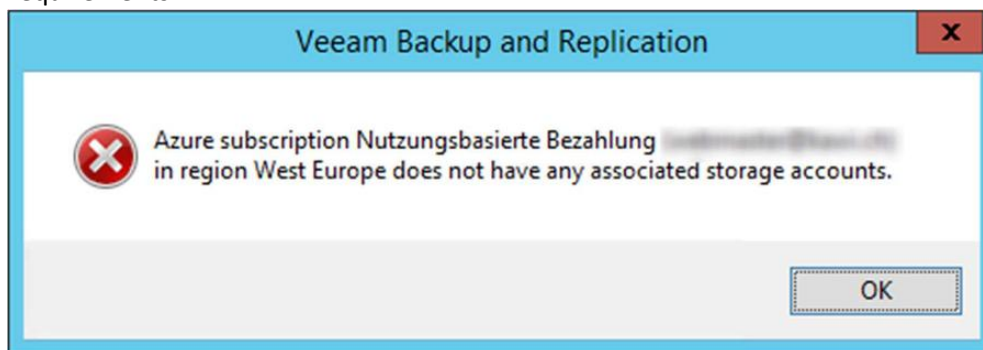
< Previous   Next >   Finish   Cancel

11. On the Add Azure Proxy Subscriptions Page, Specify the subscription model which you already use on Azure, choose the location and click Next.



The screenshot shows the 'Add Azure Proxy' wizard in Veeam Backup and Replication. The 'Subscription' step is active, indicated by a green plus icon and a blue bar. The wizard title is 'Add Azure Proxy'. The main area contains two dropdown menus: 'Subscription:' with the value 'Nutzungsbasierte Bezahlung' and 'Location:' with the value 'West Europe'. Below the 'Location:' dropdown is the text 'Choose the Azure data center region.' At the bottom, there are four buttons: '< Previous', 'Next >', 'Finish', and 'Cancel'.

12. If you receive this error message, then it's best to close the assistant and check the Azure requirements.



Note: You'll need storage and networking on Azure.

13. On the Add Azure Proxy VM Size Page, choose the size of your VM and click Next.

The screenshot shows the 'Add Azure Proxy' wizard window. The title bar says 'Add Azure Proxy'. The main window has a blue header with a green plus icon and the text 'VM Size' and 'Specify the storage account and disk type'. On the left is a sidebar with a list of steps: Name, Credentials, Deployment Model, Subscription, VM Size (highlighted), Resource Group, Network, Apply, and Summary. The main area is divided into two columns. The left column contains a 'Size:' dropdown menu with 'Basic\_A2 (2 cores, 3.5 GB memory)' selected. Below this are three rows of specifications: 'Cores: 2', 'Max disks: 4', and 'Memory: 3.5 GB'. The right column contains a 'Storage account:' dropdown menu with a blurred selection. Below this is the text 'Choose the storage account this Azure VM should be placed in.' At the bottom right are four buttons: '< Previous', 'Next >', 'Finish', and 'Cancel'.

Note: Depending on your specific need you can choose smaller or larger deployments.

14. On the Add Azure Proxy Resource Group Page, Create a new Resource Group and click Next.

The screenshot shows the 'Add Azure Proxy' dialog box with the 'Resource Group' tab selected. The dialog has a sidebar on the left with the following options: Name, Credentials, Deployment Model, Subscription, VM Size, Resource Group (selected), Network, Apply, and Summary. The main area is titled 'Resource Group' with a subtitle 'Specify the resource group to place the proxy VM into.' Below this, there are two radio button options: 'Place VM into the existing resource group:' (unselected) and 'Create a new resource group:' (selected). Under the 'Create a new resource group:' option, there is a 'Name:' field containing 'azurewesteurope'. Below that is a 'DNS name label:' field containing 'azurewesteurope'. A note states: 'DNS name label for RDP connections. To connect to the machine using Remote Desktop, use the following address in the connection settings:'. Below the note is a text box containing 'azurewesteurope.westeurope.cloudapp.azure.com'. At the bottom right, there are four buttons: '< Previous', 'Next >', 'Finish', and 'Cancel'.

**Add Azure Proxy**

**Resource Group**  
Specify the resource group to place the proxy VM into.

Name  
Credentials  
Deployment Model  
Subscription  
VM Size  
**Resource Group**  
Network  
Apply  
Summary

Specify the resource group that serves as a container to help you manage a collection of Azure resources.

☐ Place VM into the existing resource group:

☒ Create a new resource group:  
Name:

DNS name label:

DNS name label for RDP connections. To connect to the machine using Remote Desktop, use the following address in the connection settings:

< Previous   Next >   Finish   Cancel

15. On the Add Azure Proxy Network Page, select the Virtual Network and Subnet and click Next.

The screenshot shows the 'Add Azure Proxy' wizard window. The title bar is blue with the text 'Add Azure Proxy' and a close button. The window has a light blue header with a green plus icon and the text 'Network Specify the virtual network'. On the left is a sidebar with a list of steps: Name, Credentials, Deployment Model, Subscription, VM Size, Resource Group, Network (highlighted), Apply, and Summary. The main area contains two dropdown menus. The first is labeled 'Virtual network:' and has 'veeamnetwork' selected; below it is the text 'Specify the virtual network to connect proxy VM to.'. The second is labeled 'Subnet:' and has 'veeamsubnet' selected; below it is the text 'Choose an IP address range for the virtual network.'. At the bottom right are four buttons: '< Previous', 'Next >', 'Finish', and 'Cancel'.

**Add Azure Proxy**

**Network**  
Specify the virtual network

Virtual network:  
veeamnetwork  
Specify the virtual network to connect proxy VM to.

Subnet:  
veeamsubnet  
Choose an IP address range for the virtual network.

< Previous   Next >   Finish   Cancel

16. On the Add Azure Proxy Apply Page, review the settings and click Finish.

**Add Azure Proxy**

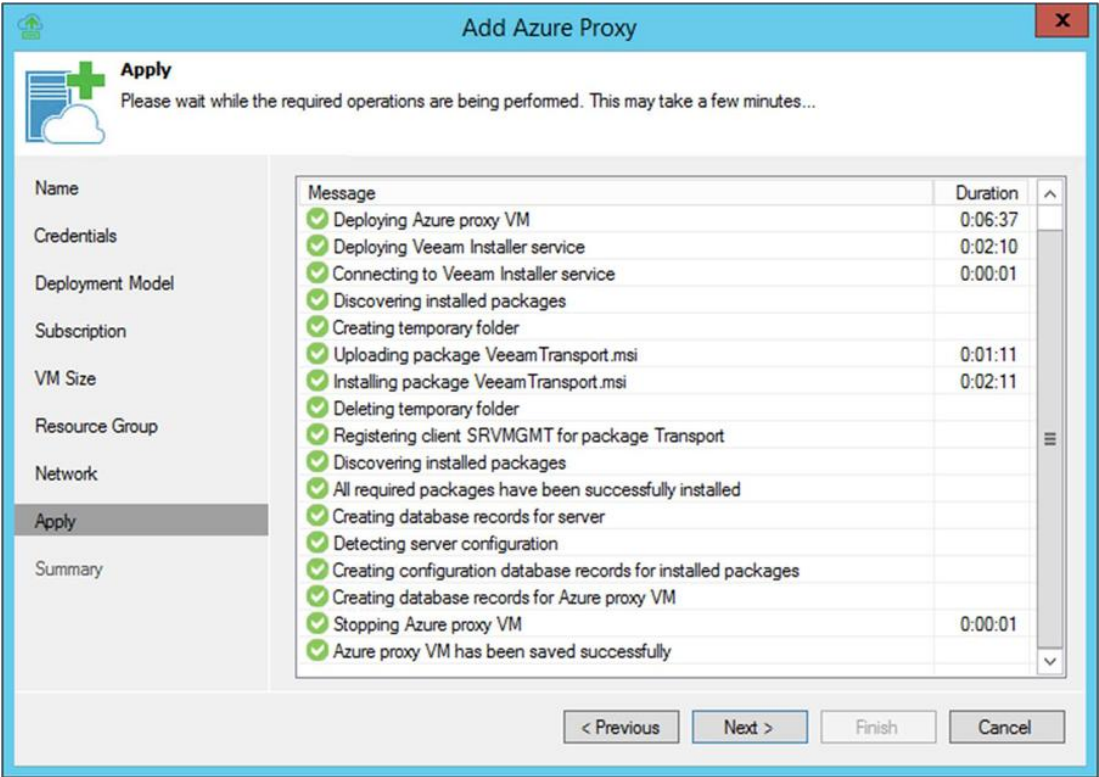
**Apply**

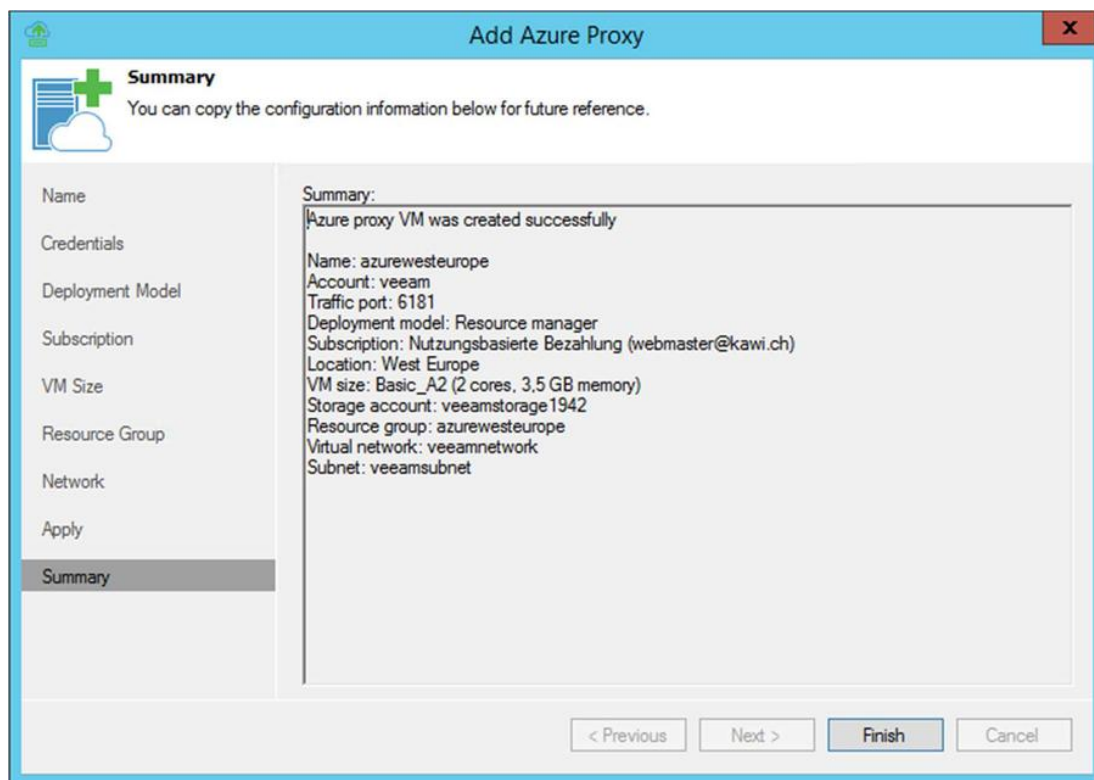
Please wait while the required operations are being performed. This may take a few minutes...

Name	Message	Duration
Credentials	✓ Starting infrastructure item creation job	
Deployment Model	✓ Deploying Azure proxy VM	0:06:37
Subscription	▶ Deploying Veeam Installer service	0:00:47
VM Size		
Resource Group		
Network		
<b>Apply</b>		
Summary		

< Previous      Next >      **Finish**      Cancel

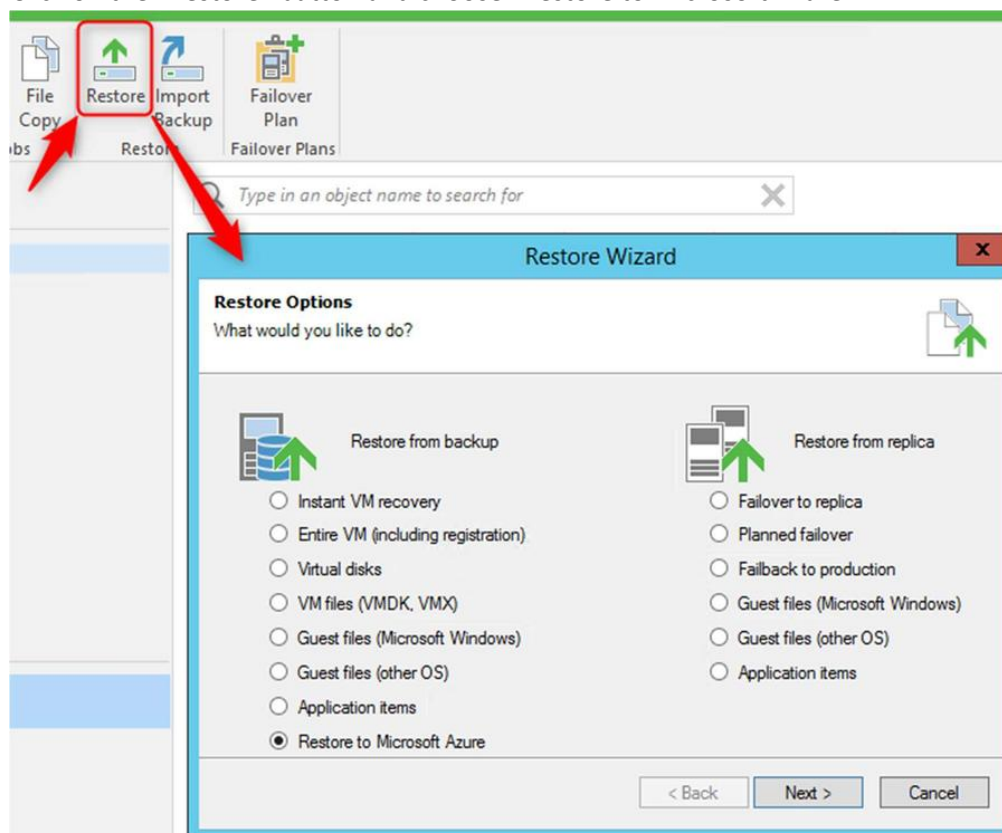
17. On the Add Azure Proxy Apply Page, review the changes and make sure everything deployed successfully and then click Finish.





## Restore a VM to Azure

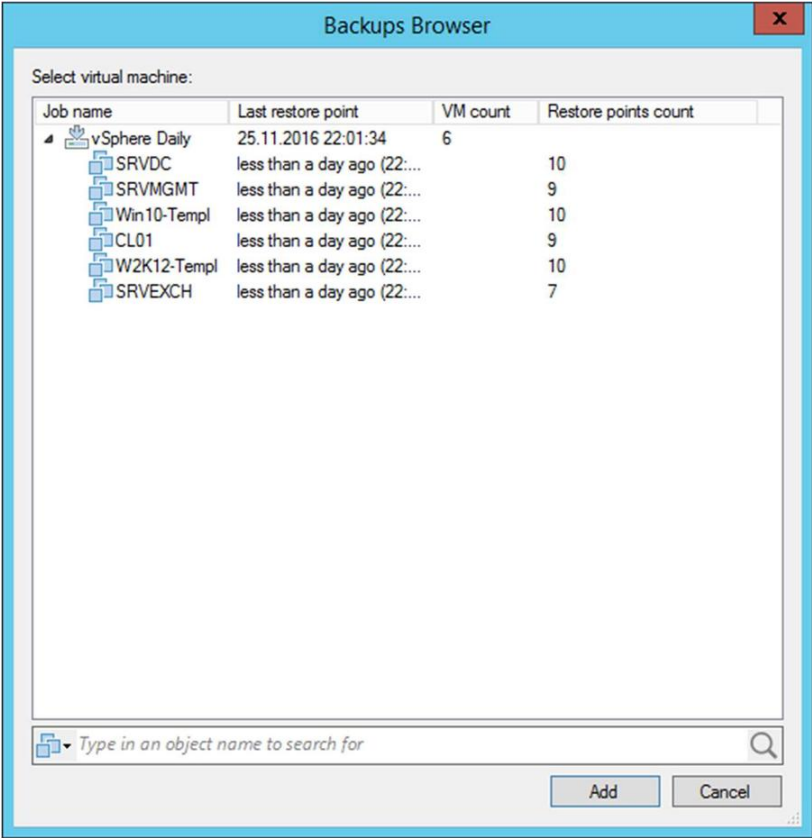
1. Open your Veeam Backup & Replication console.
2. Click on the “Restore” button and choose “Restore to Microsoft Azure.”



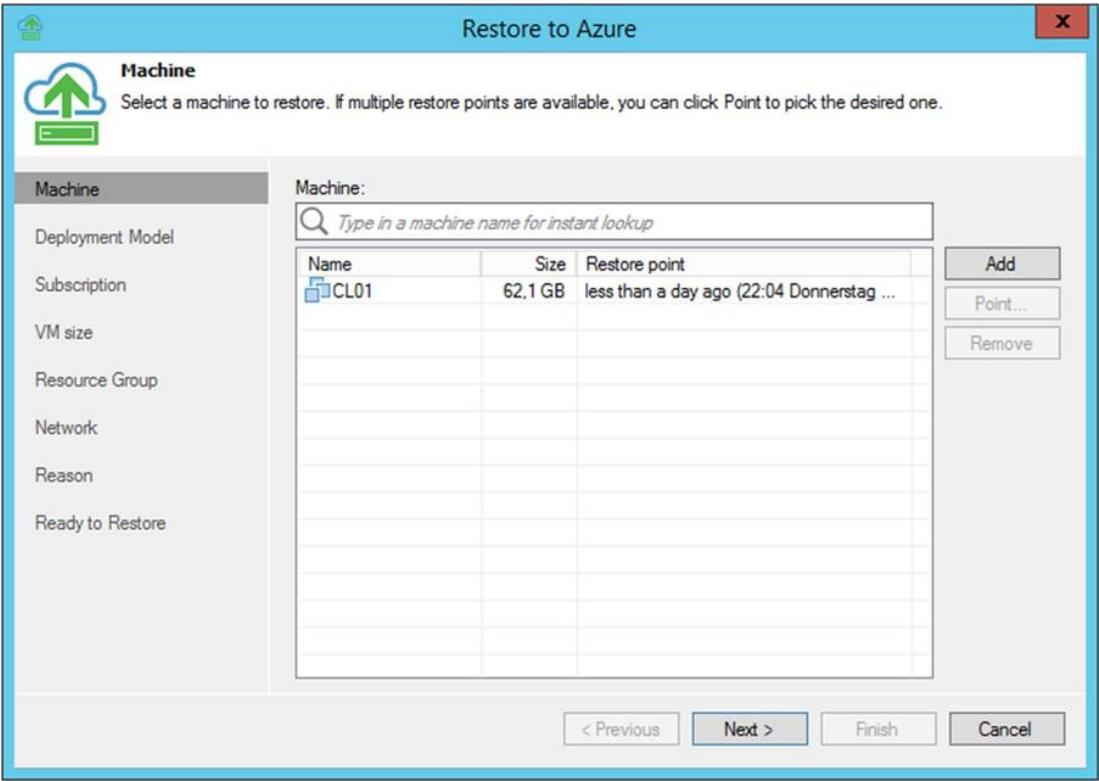
3. On the Restore to Azure Machine Page, click Add.

[illegible]

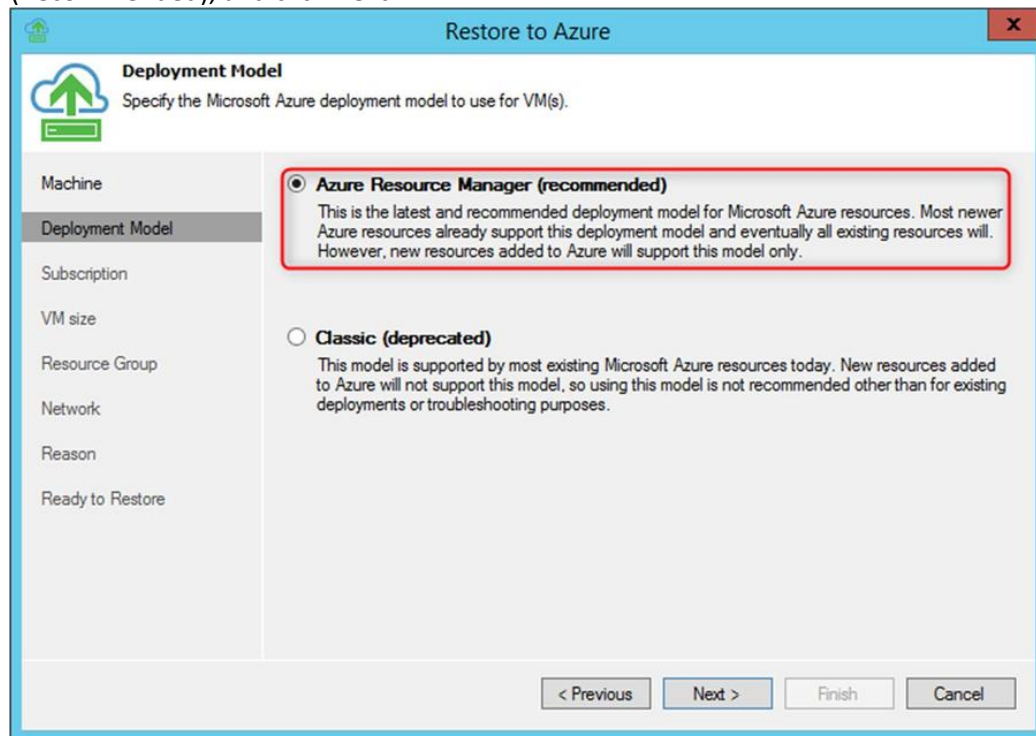
4. On the Backups Browser Page, Search your backups for the VM, then click Add.



5. Back on the Restore to Azure Machine Page, make sure your VM is in the list and click Next.



6. On the Restore to Azure Deployment Model Page, select Azure Resource Manager (Recommended), and click Next.



7. On the Restore to Azure Subscription Page, select your subscription, location, and Azure Proxy VM, and click Next.

**Restore to Azure**

**Subscription**  
Specify Azure subscription and data center region to restore to.

Machine  
Deployment Model  
**Subscription**  
VM size  
Resource Group  
Network  
Reason  
Ready to Restore

Subscription:  
Nutzungsbasierte Bezahlung

Specify Azure subscription and data center region to restore to.

Location:  
West Europe

Choose the Azure data center region where the VMs should be restored to.

☒ Use Azure proxy VM (recommended)

Proxy VM:  
azurewesteurope

Using a proxy VM running in target Azure data center significantly improves restore performance.

< Previous   Next >   Finish   Cancel

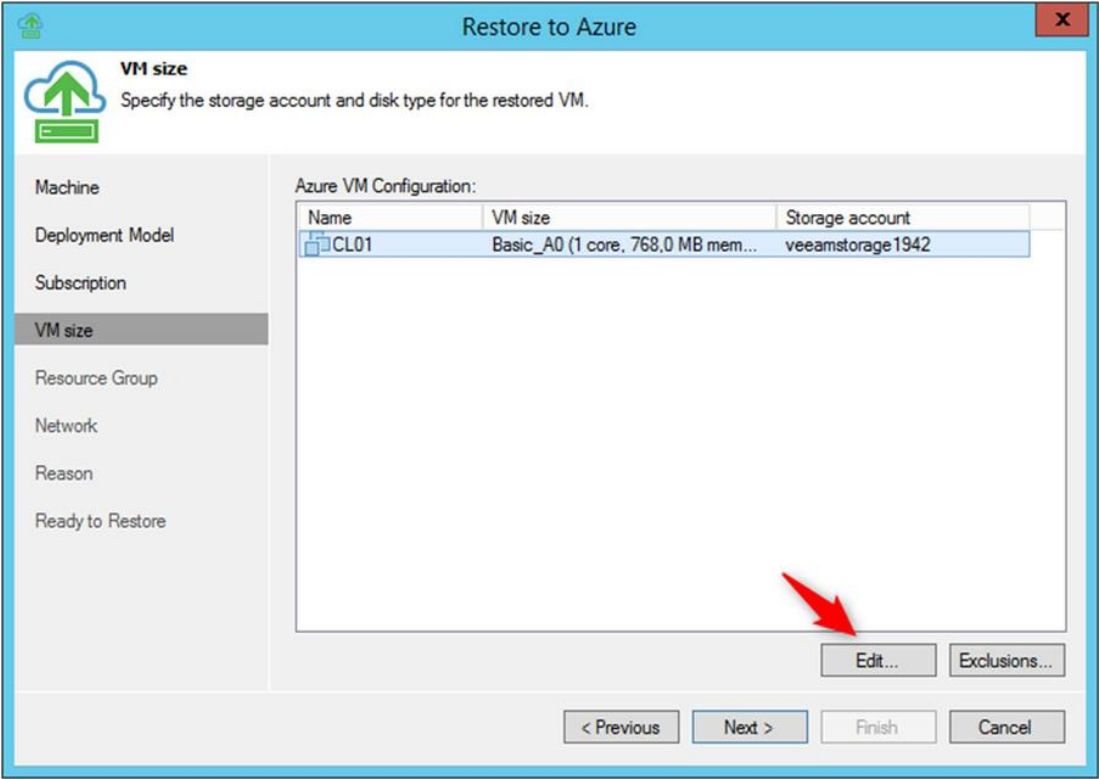
8. On the Veeam Backup and Replication popup, click No. We don't need the Linux Appliance at this time.

**Veeam Backup and Replication**

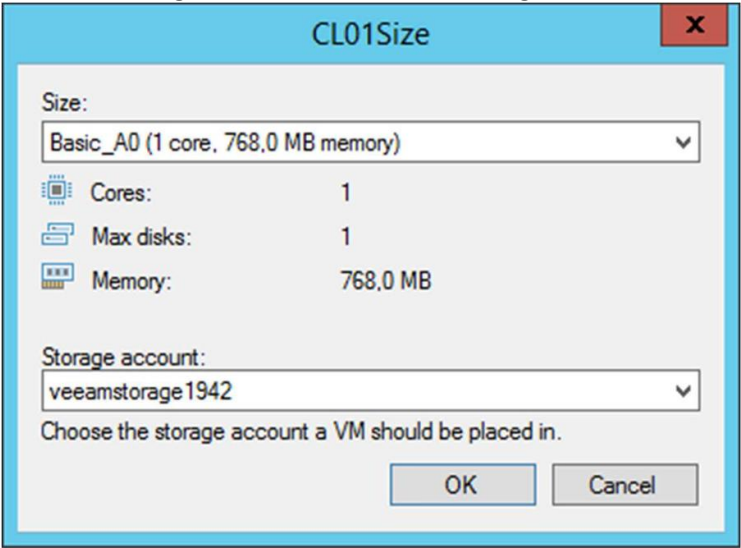
Restoring Linux computers require helper appliance.  
Deploy the appliance now?

Yes   No

9. On the Restore to Azure VM Size Page, click Edit.

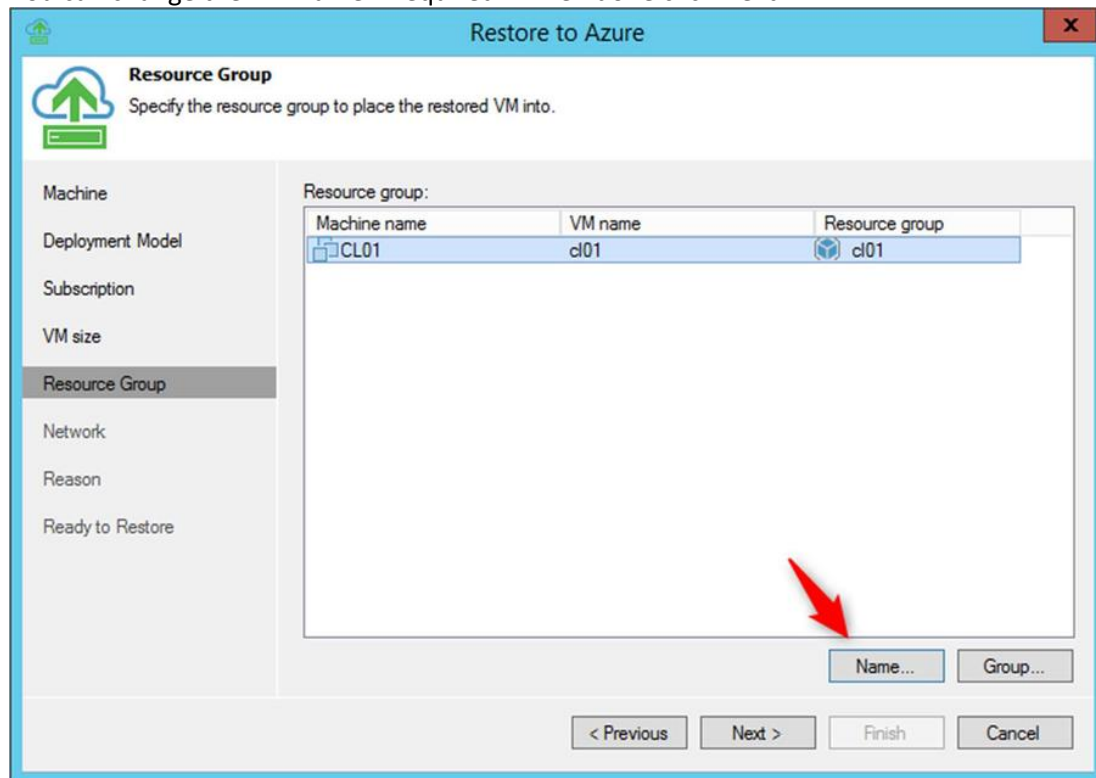


10. On the Size Page, choose VM size and storage account and click OK.

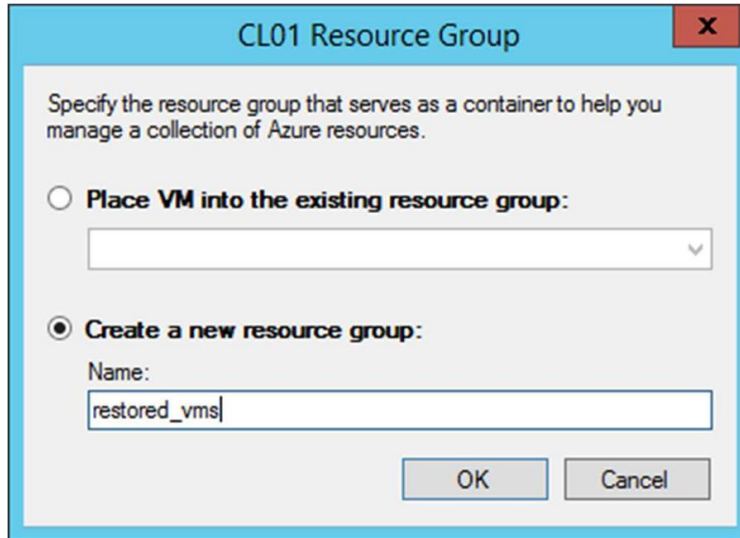


Note: Depending on your need you'll need probably some adjustment here.

11. On the Restore to Azure Resource Group Page, click on your VM and then click Name. You can change the VM name if required. When done click Next.



12. On the Resource Group Page, Create a new Resource Group called `restored_vms` and click OK.



CL01 Resource Group

Specify the resource group that serves as a container to help you manage a collection of Azure resources.

☐ Place VM into the existing resource group:

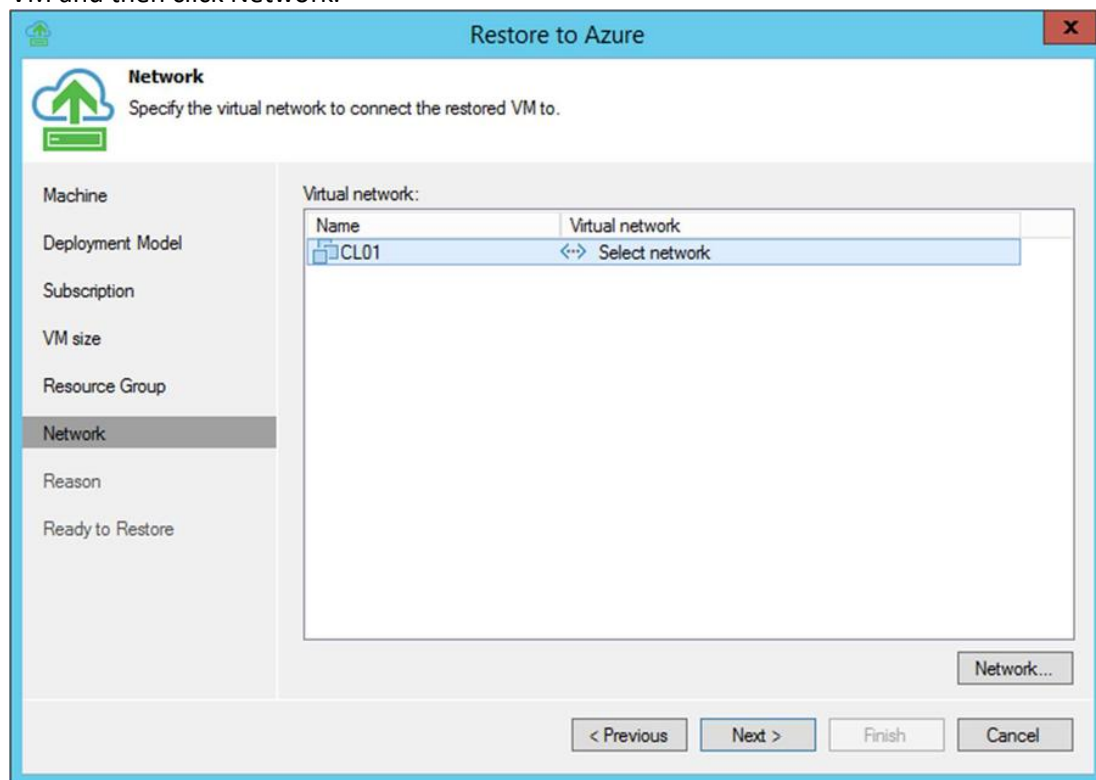
☐ **Create a new resource group:**

Name:

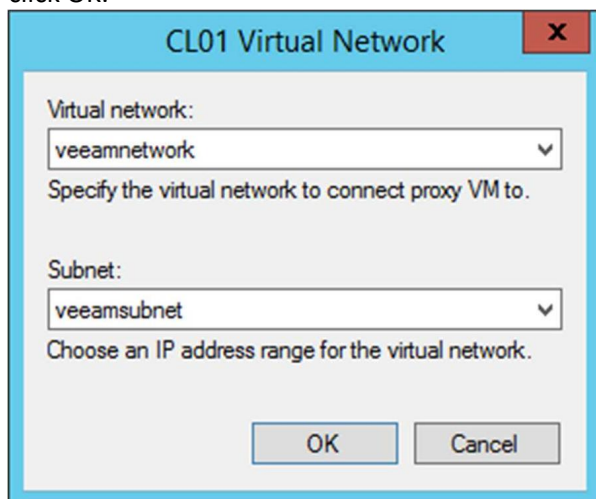
OK Cancel

Note: Depending on your (already existing) Azure deployment it's probably a good idea to have things organized.

13. On the Restore to Azure Network Page, specify the networking for your VM. Click on the VM and then click Network.

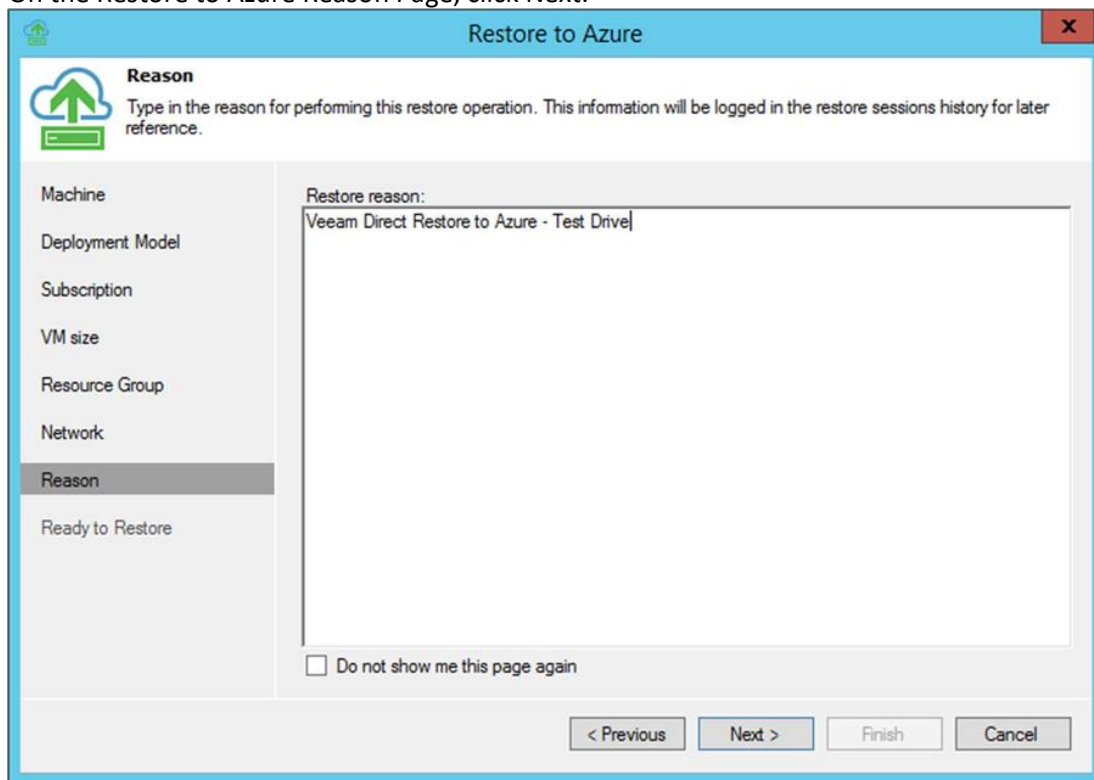


14. On the Virtual Network page, choose your existing network and subnet on Azure and click OK.



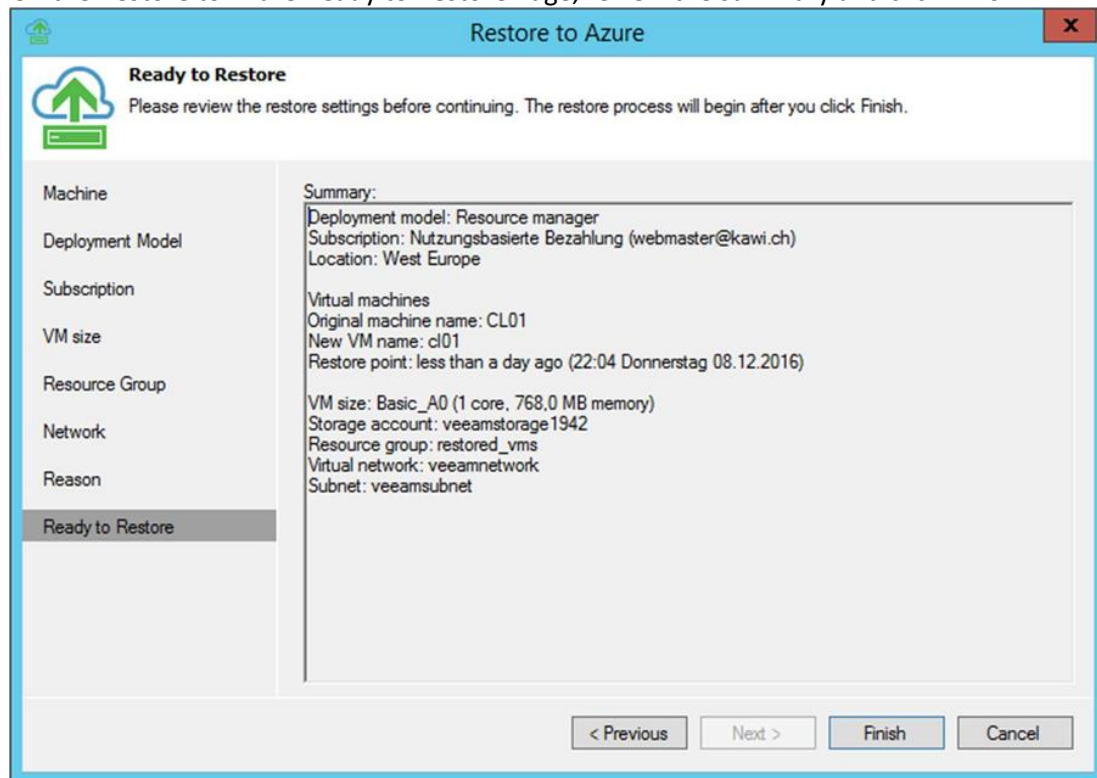
The dialog box is titled "CL01 Virtual Network" and has a close button (X) in the top right corner. It contains two dropdown menus. The first is labeled "Virtual network:" and has "veeamnetwork" selected. Below it is the text "Specify the virtual network to connect proxy VM to." The second dropdown menu is labeled "Subnet:" and has "veeamsubnet" selected. Below it is the text "Choose an IP address range for the virtual network." At the bottom are "OK" and "Cancel" buttons.

15. On the Restore to Azure Reason Page, click Next.

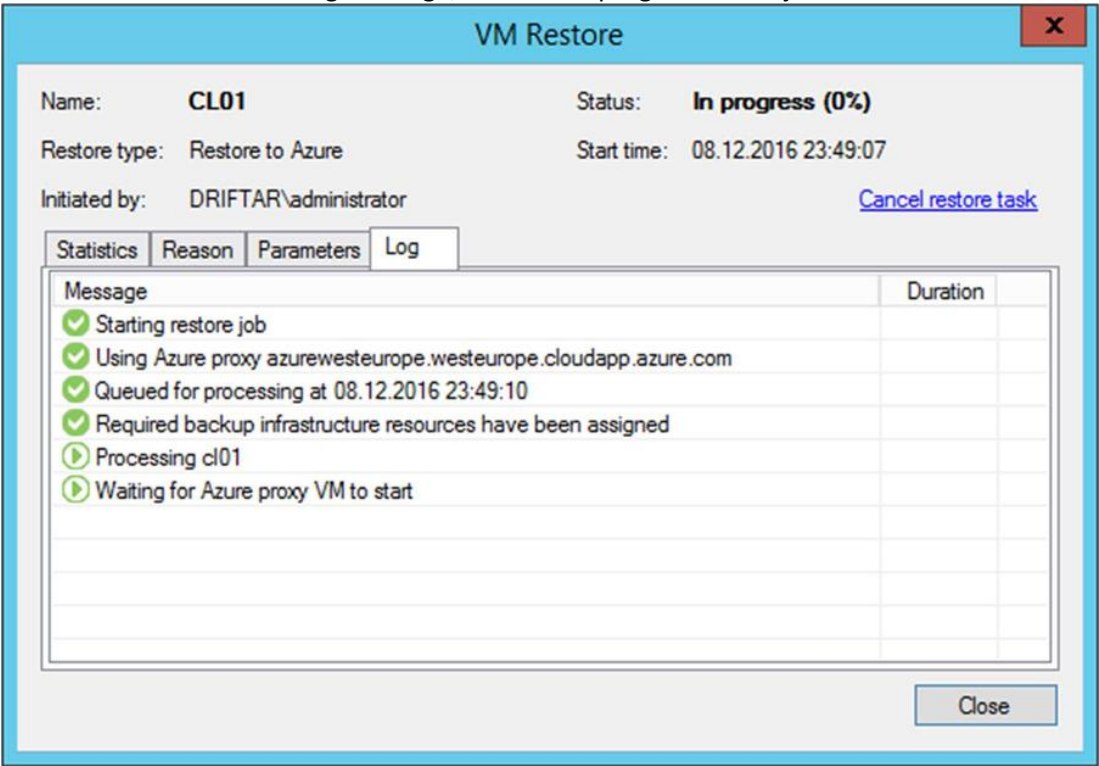


The dialog box is titled "Restore to Azure" and has a close button (X) in the top right corner. It features a sidebar on the left with a list of options: Machine, Deployment Model, Subscription, VM size, Resource Group, Network, Reason (which is highlighted), and Ready to Restore. The main area is titled "Reason" and contains a text box with the text "Veeam Direct Restore to Azure - Test Drive". Above the text box is a small icon of a cloud with an upward arrow and a server icon. Below the text box is a checkbox labeled "Do not show me this page again". At the bottom are four buttons: "< Previous", "Next >", "Finish", and "Cancel".

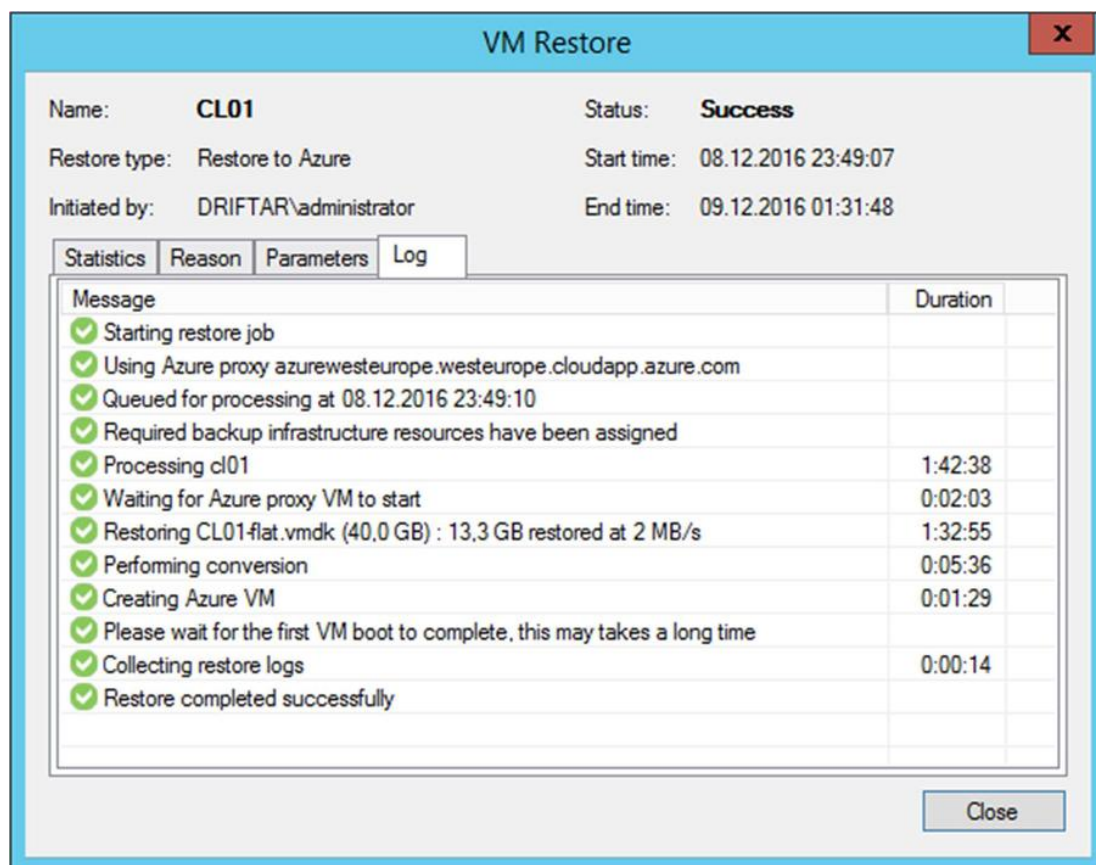
16. On the Restore to Azure Ready to Restore Page, review the summary and click Finish.



17. On the VM Restore Job Progress Page, review the progress of the job and click close.



It will take a while to restore to any cloud service. The amount of time depends on the network connectivity, size of VM, and amount of data. This VM used in the lab took approximately 1.75 hours to complete.



## Chapter 12

# Configure 2 Factor Authentication (2FA) for Veeam with DUO

**By: Rhys Hammond (VMWare vEXPERT / Veeam Vanguard)**

Something that all Veeam administrators should consider is how secure the underlying servers running your Veeam software really are. To help improve security, I always try and run through a few recommendations with each Veeam administrator I work with,

- Inbound connectivity to backup servers from the Internet must not be allowed (3389 anyone?)
- Any accounts used for RDP access must not have Local Administrator privileges on jump servers, and you must never use the saved credentials functionality for RDP access or any other remote console connections.
- Ensure timely guest OS updates on backup infrastructure servers

An excellent resource for keeping up to date on Veeam security recommendations is [here](#). I like to check it out every 3-6 months to ensure I'm still making the right recommendations to my customers.

One other thing I like to recommend in addition to the best practices above is enabling 2FA (Two-Factor Authentication) for all logon sessions to underlying servers running Veeam components such as the VBR server, proxies and especially repositories. With 2FA, even if an attacker illegally acquires the correct username and password, the attacker is also required to gain access to the device used to receive the 2FA verification code. Often this device is a mobile phone or a security token which can easily be disabled if lost or stolen.

While 2FA for Veeam consoles is presently not possible (it is a heavily requested feature though) and even with 2FA for login sessions, there is still a risk that an attacker can access Veeam infrastructure via a Veeam Console running from another machine. This is why off-site/offline backups are so so critical in today's world of ransomware. Leveraging Veeam Cloud Connect Backup with its Insider Protection feature is a great way to efficiently protect against this kind of risk.

This chapter will go into detail on how to quickly and easily enable 2FA for RDP and local logon sessions connecting to your Veeam server.

In this article, we'll be leveraging Duo, a robust two-factor authentication provider that can be added to not only our Veeam servers but also VPN, email, web portal, cloud services, and more. After Duo has been successfully configured, any user approves a secondary authentication request pushed to a Duo Mobile smartphone app. Users can also authenticate by answering a phone call or by entering a one-time passcode generated by the Duo Mobile app, a compatible hardware token, or received via SMS.

The great thing about Duo is that is free for the first 10 users which is sufficient for most small to medium Veeam environments.



The general overview of what need we need to complete are as follows;

Download the mobile app (via Apple App Store or Google Play) – Use the “push” feature to authenticate into our admin account.

Protect an application – Choose your first application to protect to protect using the Duo Admin Panel.

Enroll our first end user – Add a user and device manually, or consider other options, like self-enrollment.

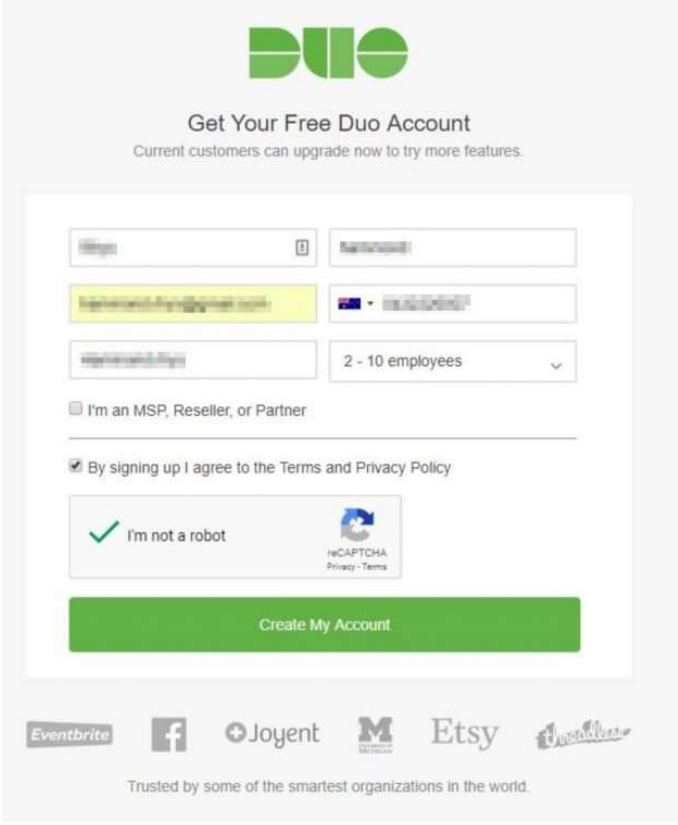
## Important Notes

- Installing Duo Authentication for Windows Logon adds two-factor authentication to all Windows login attempts, whether via a local console or over RDP unless you select the “Only prompt for Duo authentication when logging in via RDP” option in the installer. If two-factor is enabled for both RDP and console logons, it may be bypassed by restarting Windows into Safe Mode (e.g., in case of a configuration error). If you wish to protect local console logons with Duo, please see the FAQ for some guidance on securing your Windows installation appropriately.
- Duo Authentication for Windows Logon doesn’t support inline self-service enrollment. We recommend using bulk enrollment to send your users unique self-enrollment links via email. Read the enrollment documentation to learn more.
- Additional configuration may be required to log in using a Microsoft attached account. See Can I Use Duo with a Microsoft Account? For more information.
- Windows users must have passwords to log in to the computer. Users with blank passwords may not login after Duo Authentication installation.

## Creating a DUO Account

1. To get started we first must create a Duo account at <https://signup.duo.com/>, during this first step we are prompted to provide our name, email address, a phone number, company/account name and the size of the organization. Since I was using this in my

home lab I've selected 2-10 employees.



The image shows the Duo account creation page. At the top is the Duo logo, followed by the heading "Get Your Free Duo Account" and a subtext "Current customers can upgrade now to try more features." Below this is a registration form with several input fields: a name field, an email field, a password field, a company name field, and a country dropdown menu. A dropdown menu for employee count is set to "2 - 10 employees". There is a checkbox for "I'm an MSP, Reseller, or Partner" and a checked checkbox for "By signing up I agree to the Terms and Privacy Policy". A reCAPTCHA widget is present with the text "I'm not a robot" and a green checkmark. At the bottom of the form is a large green button labeled "Create My Account". Below the form, there are logos for Eventbrite, Facebook, Joyent, Medium, Etsy, and Kickstarter, followed by the text "Trusted by some of the smartest organizations in the world."


**DUO**

**Get Your Free Duo Account**  
Current customers can upgrade now to try more features.




2 - 10 employees

☐ I'm an MSP, Reseller, or Partner

☒ By signing up I agree to the Terms and Privacy Policy

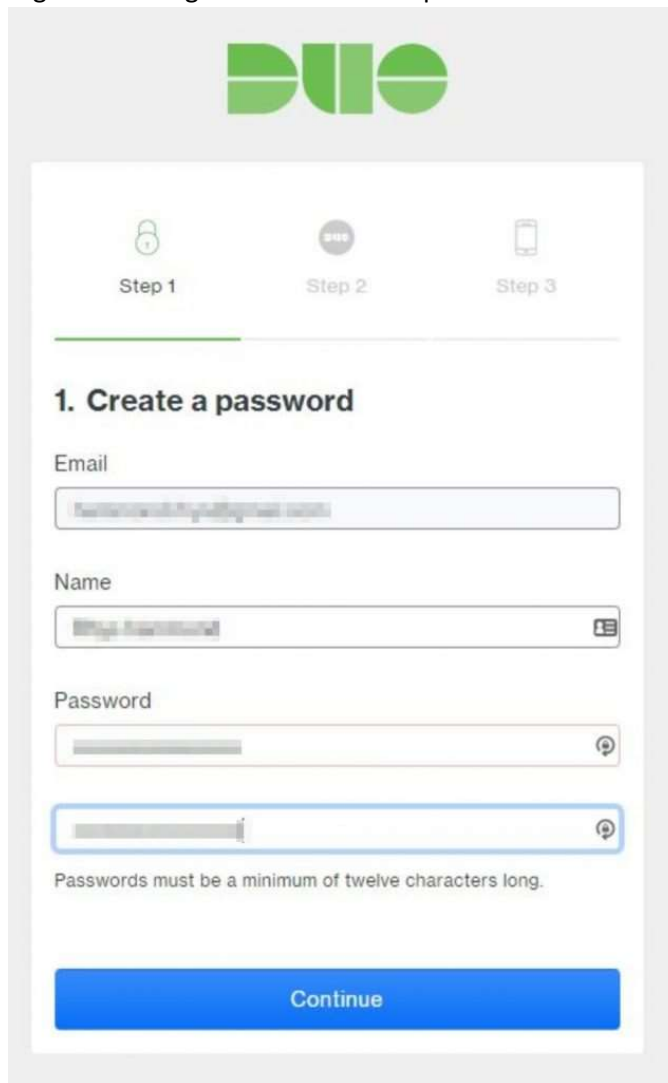
☒ I'm not a robot 

**Create My Account**

Eventbrite  Joyent  Etsy 

Trusted by some of the smartest organizations in the world.

2. At the next step, we are prompted to create a password for the account that will be registered using the email address provided in the earlier step.

The screenshot shows the Duo Mobile account creation interface. At the top is the Duo logo. Below it are three steps: Step 1 (active), Step 2, and Step 3. The main heading is "1. Create a password". There are three input fields: "Email" (containing "test@test.com"), "Name" (containing "John Doe"), and "Password" (containing "123456789012"). Below the password field is a second password field (containing "123456789012") and a note: "Passwords must be a minimum of twelve characters long." At the bottom is a blue "Continue" button.

3. Next, the account registration process is going to display a QR code and prompt us to activate Duo Mobile; this is done by installing an app on a mobile phone.

To install the app, we need to launch the app store and search for "Duo Mobile."

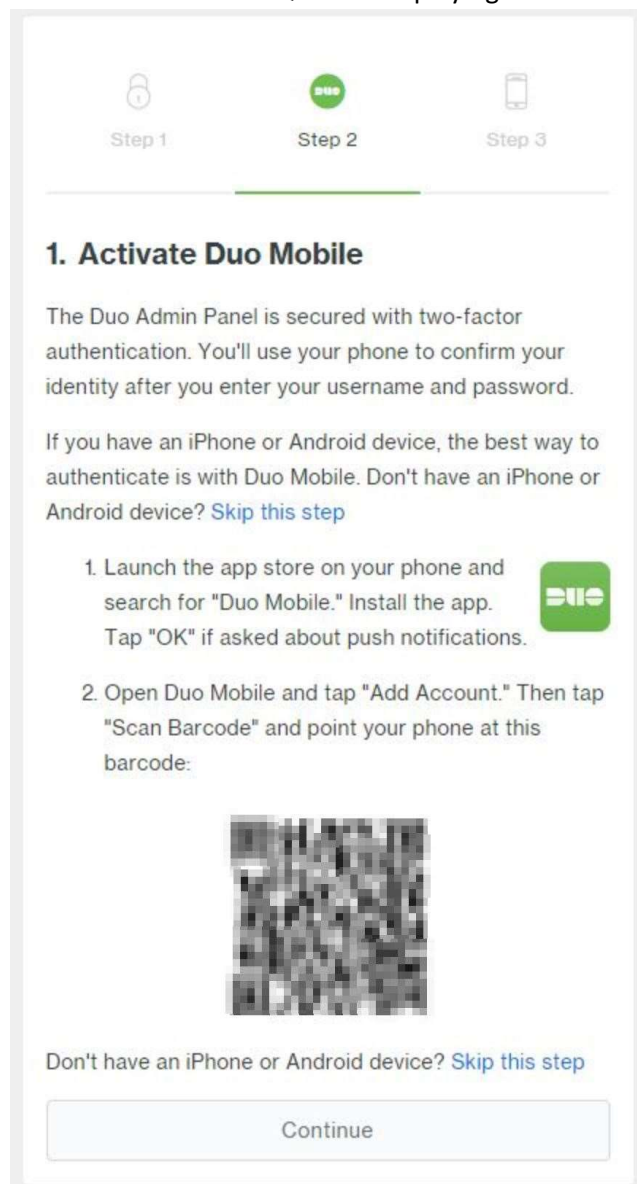
Alternatively, click on your respective marketplace below.

<https://play.google.com/store/apps/details?id=com.duosecurity.duomobile&hl=en>

<https://itunes.apple.com/au/app/duo-mobile/id422663827?mt=8>

After the app has been installed and opened we are prompted to either “Get started” or “Get my account back,” in this case, we want to select the first option which will then

allows us to scan that QR code displaying on the Duo account creation page.



4. Next, we need to provide a backup verification method. Duo recommend providing a mobile phone, not an office landline.

Step 1 Step 2 Step 3

## 1. Set a Backup

We'll use this phone number as a backup verification method.

Phone number

Tip: use your mobile phone, not an office landline.

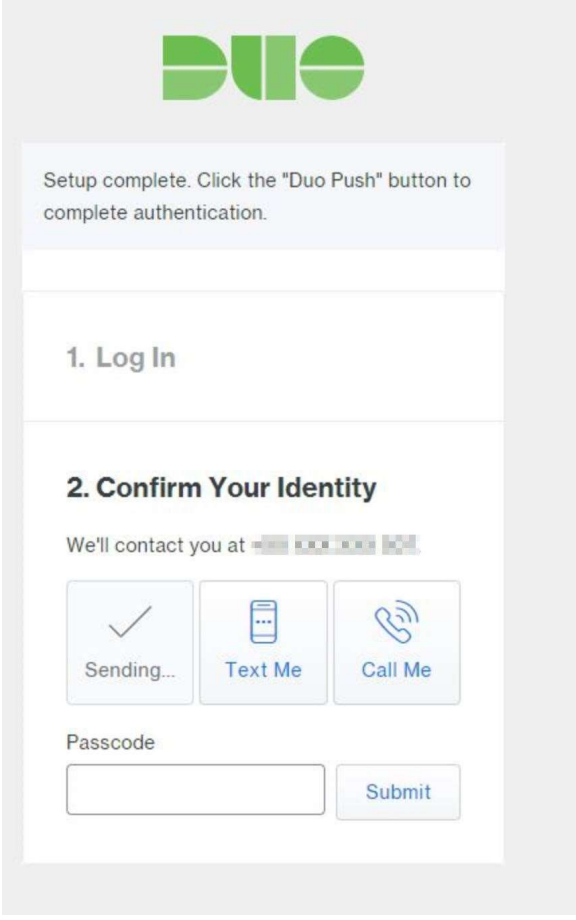
International number? Add your country code.

For example: "+44 7981-897555"

Finish

5. At this stage, our Duo account requires authentication to verify our identity. We can choose 'Duo Push' which will use the Duo Mobile app recently installed on our phone to

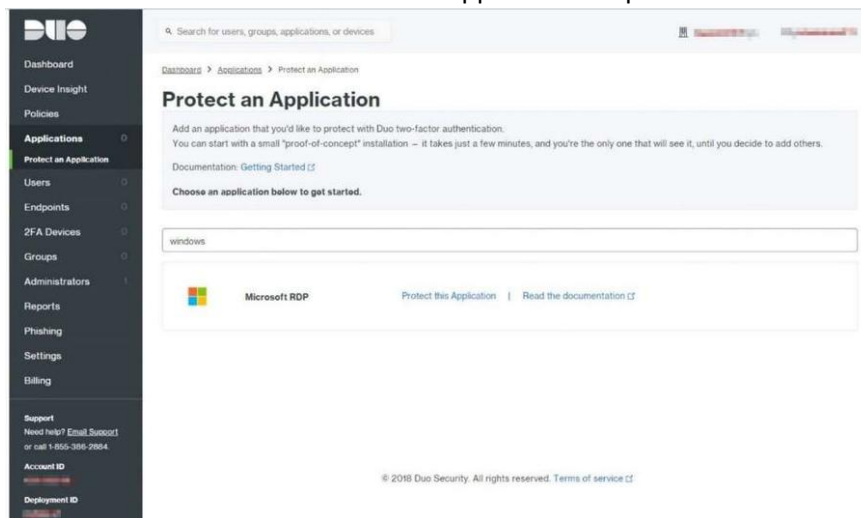
complete authentication.



The image shows a screenshot of the Duo mobile application interface. At the top is the Duo logo. Below it, a light blue box contains the text: "Setup complete. Click the 'Duo Push' button to complete authentication." The interface is divided into two main sections. The first section, titled "1. Log In", is currently empty. The second section, titled "2. Confirm Your Identity", displays a message: "We'll contact you at [redacted phone number]". Below this message are three buttons: "Sending..." (with a checkmark icon), "Text Me" (with a smartphone icon), and "Call Me" (with a telephone handset icon). At the bottom of the section, there is a "Passcode" label, a text input field, and a "Submit" button.

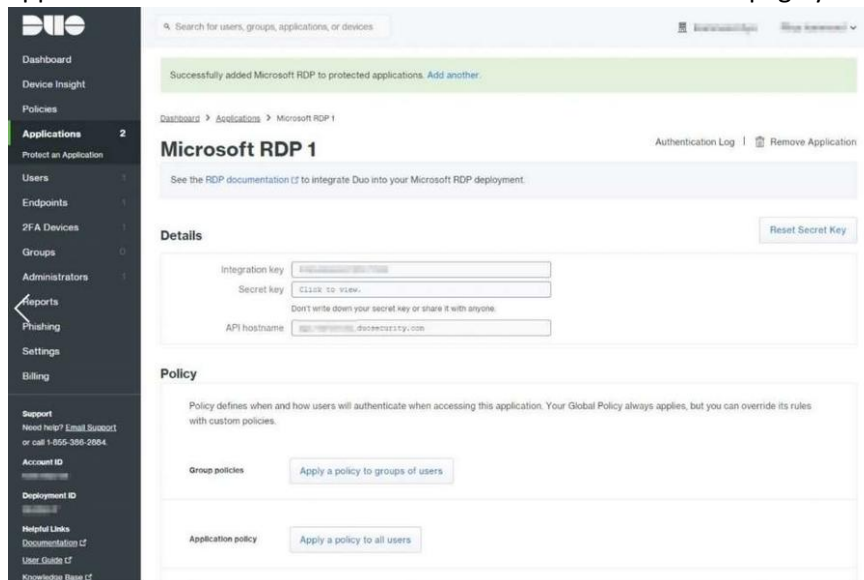
## Configuring DUO

1. After completing authentication we are taken into the Duo Admin console, here we want to expand 'Applications' on the left-hand navigation pane and click 'Protect an Application.' In the search field, we type 'Microsoft RDP'; we should see a single result for Microsoft RDP. Click 'Protect this Application' to proceed.



2. The next screen that appears states that Duo has 'Successfully added Microsoft RDP to protected applications.' Here we want to take note of three important values; they are as follows;
  - Integration key
  - Secret key
  - API Hostname

We need these values when installing the 'Duo Authentication for Windows Logon' application onto our Veeam server later on so don't close this page yet.



**Important Note:** Before proceeding, if possible, take a snapshot of the Veeam server and ensure the veeam configuration backup completed successfully recently. We're about to enable 2FA on the server, and by the very nature of how 2FA works, if it's misconfigured, you will be denied when attempting to log into the server. 2FA may be bypassed by restarting Windows into Safe Mode.

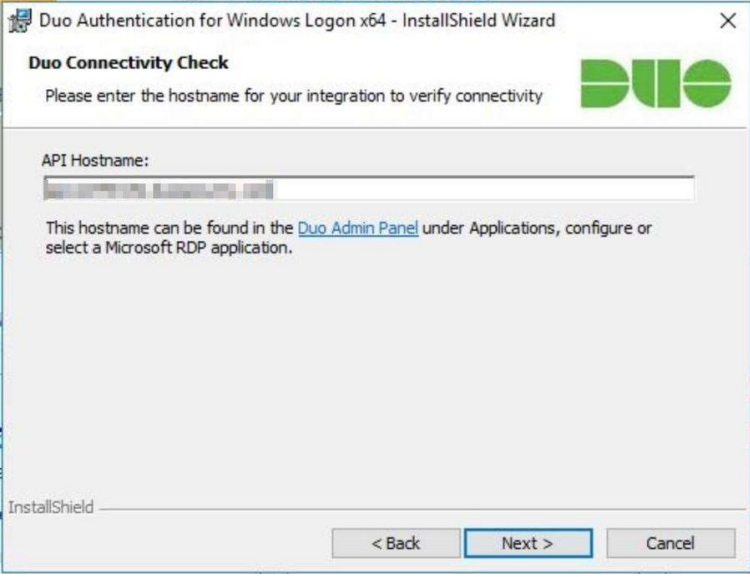
## Installing the DUO Client

1. To continue with deployment, we need to download the 'Duo Authentication for Windows Logon' Installer Package which will be installed onto our Veeam server. To download the installer click here: <https://dl.duosecurity.com/duo-win-login-latest.exe> Once the first server is configured we can add more servers in the Duo console admin page, I recommend adding 2FA for at least any windows backup repositories. Once we have downloaded the installer package, we want to start the installation on the Veeam server to enable 2FA authentication during logon.



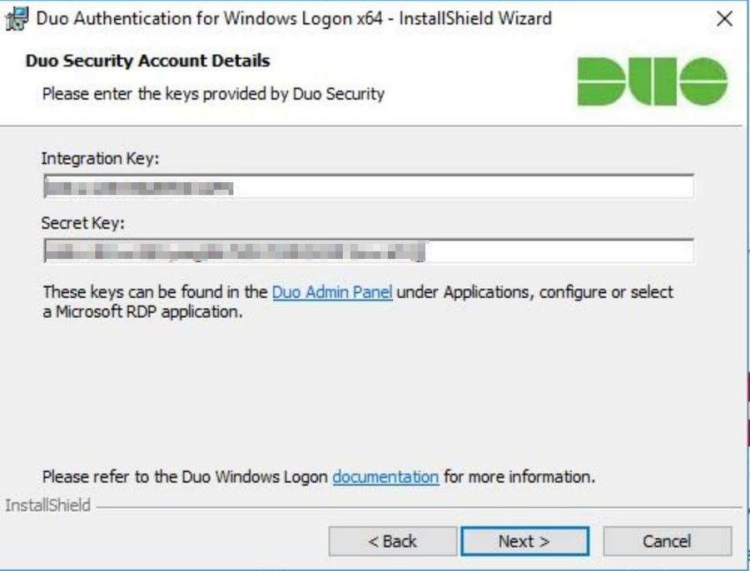
2. The first thing the installer requires us to enter is the API hostname which can be found in the Duo Admin Panel under Applications, configure or Select a Microsoft RDP

application.



The image shows a Windows dialog box titled "Duo Authentication for Windows Logon x64 - InstallShield Wizard". The main heading is "Duo Connectivity Check". Below it, the text says "Please enter the hostname for your integration to verify connectivity". There is a text input field for "API Hostname:". Below the field, a note states: "This hostname can be found in the [Duo Admin Panel](#) under Applications, configure or select a Microsoft RDP application." At the bottom, there are three buttons: "< Back", "Next >", and "Cancel". The "Next >" button is highlighted with a blue border. The InstallShield logo is in the bottom left corner.

3. Next, we provide the integration key and secret key.

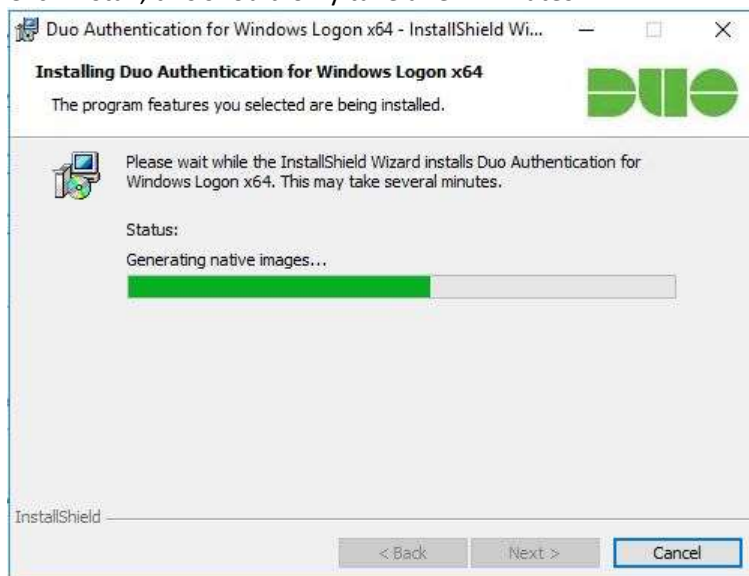


The image shows a Windows dialog box titled "Duo Authentication for Windows Logon x64 - InstallShield Wizard". The main heading is "Duo Security Account Details". Below it, the text says "Please enter the keys provided by Duo Security". There are two text input fields: "Integration Key:" and "Secret Key:". Below the fields, a note states: "These keys can be found in the [Duo Admin Panel](#) under Applications, configure or select a Microsoft RDP application." At the bottom, there is a line of text: "Please refer to the Duo Windows Logon [documentation](#) for more information." Below this, there are three buttons: "< Back", "Next >", and "Cancel". The "Next >" button is highlighted with a blue border. The InstallShield logo is in the bottom left corner.

4. After providing the integration and secret key, we are prompted to confirm a few options. I just left these on defaults.

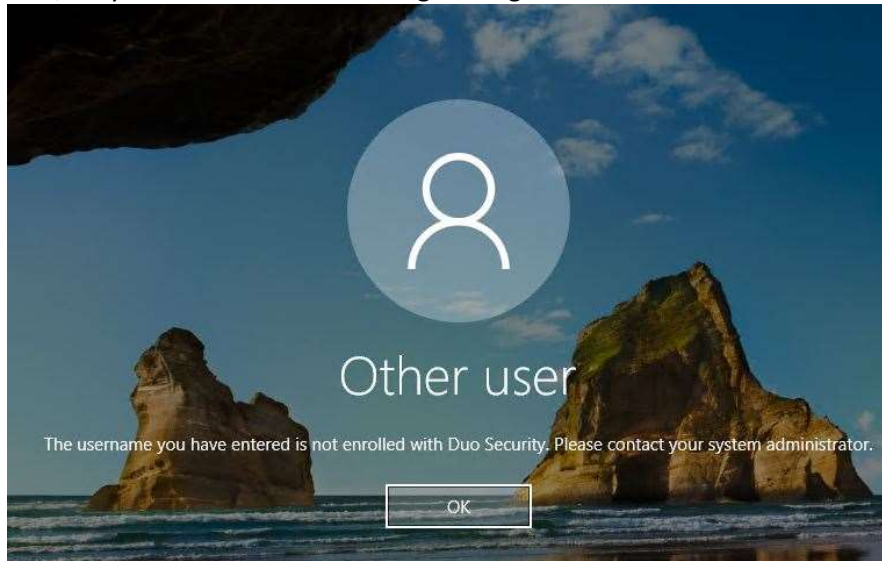


5. Click install; this should only take a few minutes.



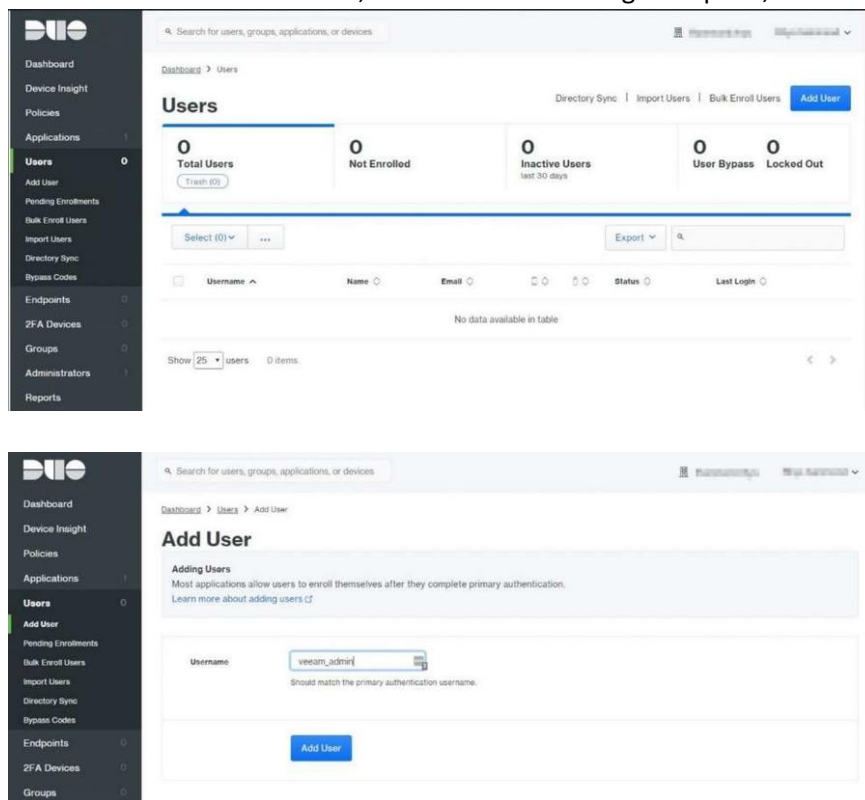
6. Once installed, the next time someone RDPs into the VBR server they will be prompted to authenticate via Duo, be mindful that until the user's account has been configured in

Duo, they will see this error message though.



## Adding DUO Users

1. To add our first user to Duo, we first need to either manually enroll or use bulk enrollment. Since I don't have an Active Directory in my home lab, I'll be using the manual method.
2. From the Duo Admin Console, on the left-hand navigation pane, we click 'Add User.'



**Note:** The username must match our Windows logon name otherwise the user will fail to login. In my case, my local administrator account has been renamed to 'veeam\_admin,' so this is the username I'll need to enter.


3. In the next step, we need to specify an email address for our user. We can also specify a username alias and whether 2FA is active, bypassed, or access is disabled altogether.

The screenshot shows the Duo user management interface for a user named **veeam\_admin**. At the top right, there are links for [Logs](#), [Send Enrollment Email](#), and [Send to Trash](#). Below the user name, a message states: "This user has not enrolled yet. See our [enrollment documentation](#) to learn more about enrolling users." The main form contains several sections: 

- Username:** A text field containing "veeam\_admin" with a dropdown arrow on the right.
- Username Alias 1:** An empty text field.
- Add alias:** A link with the text "Users can have up to 4 aliases".
- Real Name:** An empty text field.
- Email:** A text field containing "veeam\_admin@veeam.com" with a yellow highlight.
- Status:** A section with three radio button options:
  - Active:** Selected. Description: "Require two-factor authentication (default)".
  - Bypass:** Description: "Skip two-factor authentication".
  - Disabled:** Description: "Automatically deny access". Below these options is a note: "This controls the user's two-factor authentication process."
- Groups:** A section stating "You don't have any editable groups. [Add one.](#)" and a footer note: "Groups can be used for management, reporting, and policy. [Learn more about groups](#)."

4. The email address specified during user creation will be sent an enrollment email from Duo, a unique link will be in the email which is used to start enrollment for the device that the user wishes to utilize for 2FA, this could be a phone, tablet, or another supported device. Clicking on the link in the email starts the device registration process.

In my case, I was registering my mobile phone as the device.



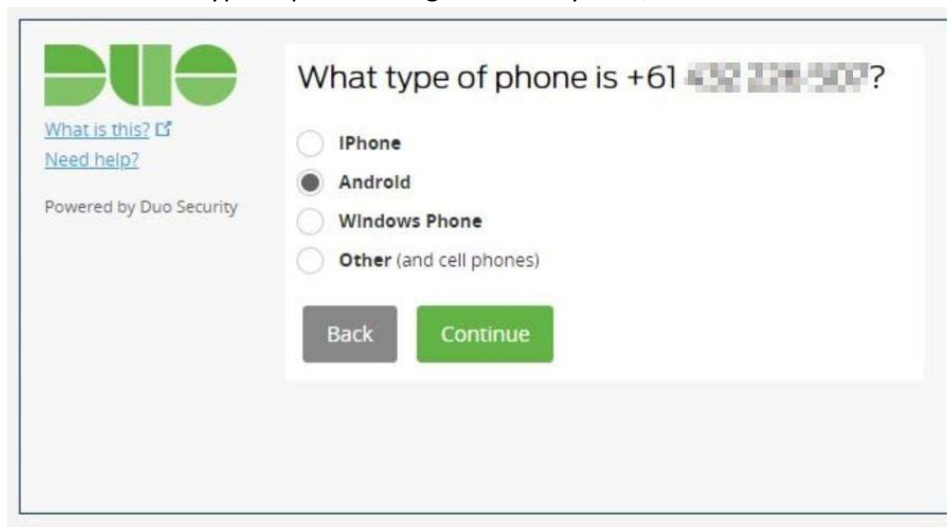
The image shows the Duo Security registration interface. On the left, there is a Duo logo, links for "What is this?" and "Need help?", and the text "Powered by Duo Security". The main heading is "What type of device are you adding?". Below this heading are four radio button options: "Mobile phone" (which is selected and has a green "RECOMMENDED" label), "Tablet (iPad, Nexus 7, etc.)", "Landline", and "U2F". At the bottom of the options is a green "Continue" button.

5. We enter our user's mobile phone number.



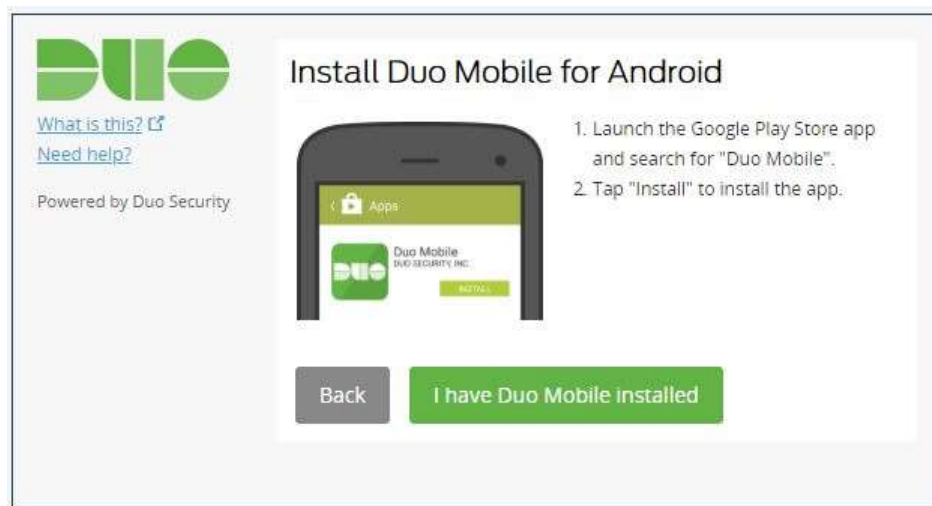
The image shows the Duo Security registration interface for entering a phone number. On the left, there is a Duo logo, links for "What is this?" and "Need help?", and the text "Powered by Duo Security". The main heading is "Enter your phone number". Below this heading is a dropdown menu showing "Australia". To the right of the dropdown is a text input field containing "+61" followed by a masked number "0000000000" and a green checkmark. Below the input field is the text "ex: (4) 1234 5678". Below this is a checkbox that is checked, followed by the text "You entered 0000000000. Is this the correct number?". At the bottom are two buttons: a grey "Back" button and a green "Continue" button.

6. We confirm the type of phone being used. In my case, an Android.



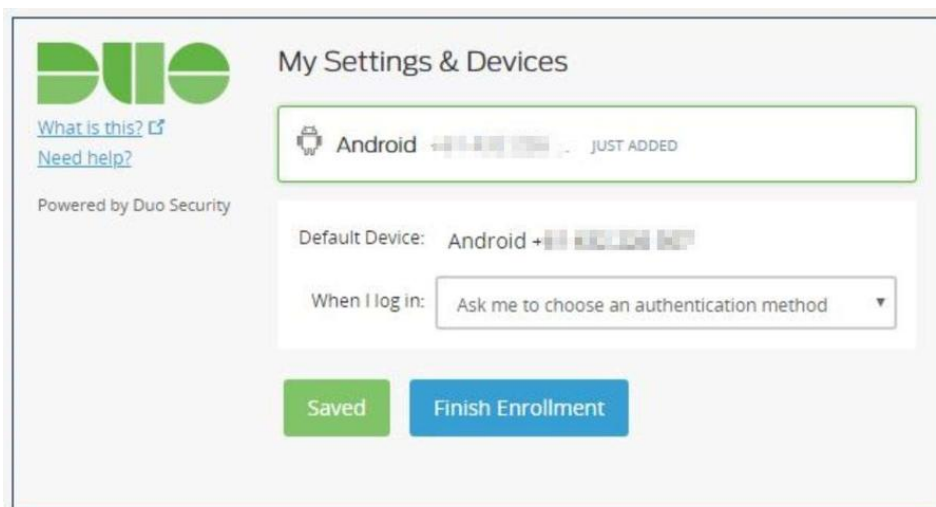
The screenshot shows the Duo Security mobile app interface. On the left, there is a sidebar with the Duo logo, links for "What is this?" and "Need help?", and the text "Powered by Duo Security". The main content area has the heading "What type of phone is +61 433 238 900?" and four radio button options: "iPhone", "Android" (which is selected), "Windows Phone", and "Other (and cell phones)". At the bottom of the main area are two buttons: "Back" and "Continue".

7. Duo wants us to install the Duo Mobile for Android app, in my case I already had it Installed.

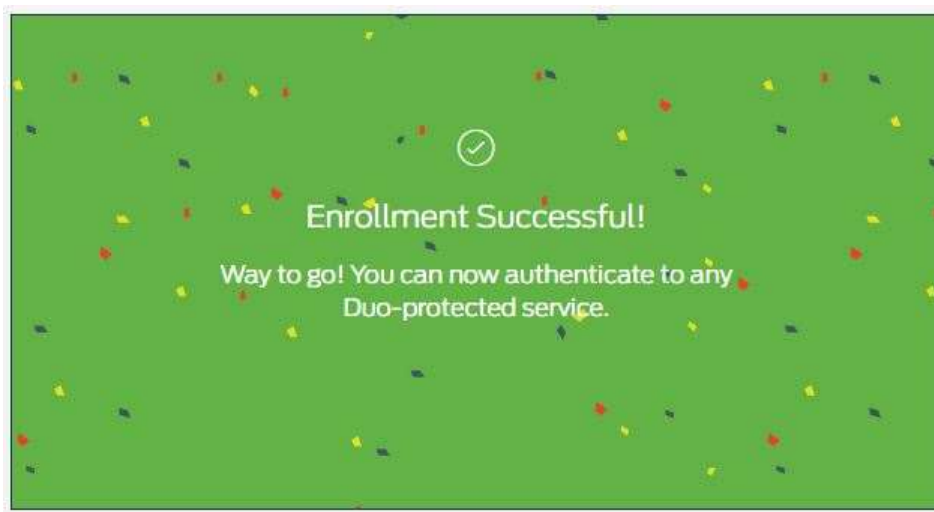


The screenshot shows the Duo Security mobile app interface. On the left, there is a sidebar with the Duo logo, links for "What is this?" and "Need help?", and the text "Powered by Duo Security". The main content area has the heading "Install Duo Mobile for Android". Below the heading is an illustration of a smartphone screen showing the Google Play Store app page for "Duo Mobile" by "DUO SECURITY INC.", with an "INSTALL" button. To the right of the illustration are two numbered steps: "1. Launch the Google Play Store app and search for 'Duo Mobile'." and "2. Tap 'Install' to install the app." At the bottom of the main area are two buttons: "Back" and "I have Duo Mobile installed".

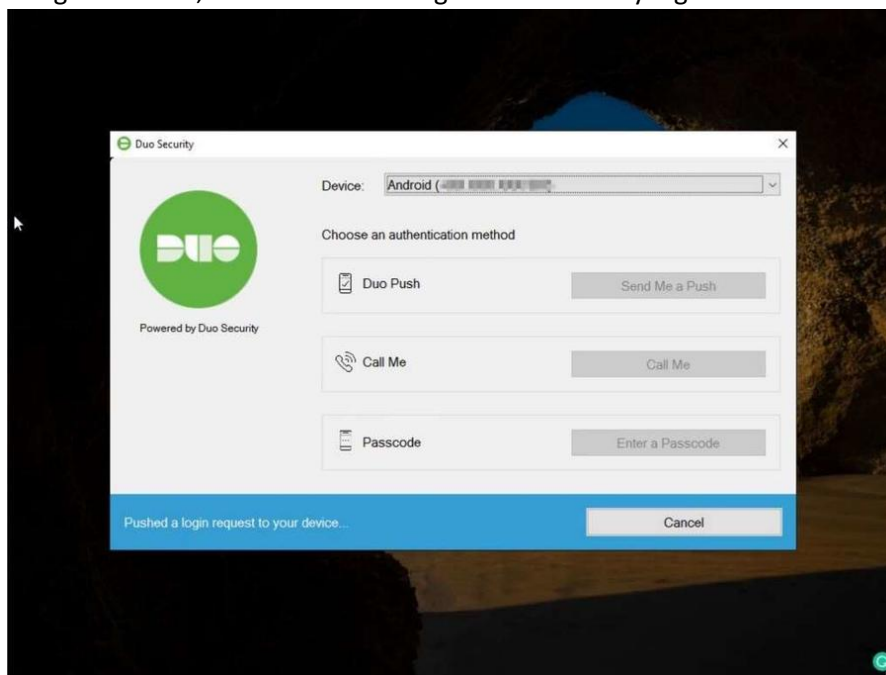
8. Another QR code is displayed, at this stage, we only open the Duo Mobile app from our user's phone, click the + symbol in the top right and register via the QR code.



9. Congratulations, we have just enrolled our first user and their 2FA device.



10. Congratulations, we have now configured 2FA for any logons to our Veeam server.





## Chapter 13

# The Secret Life of VBM Files

By: Rhys Hammond (VMWare vEXPERT / Veeam Vanguard)

### What do VMB files do?

VBM files are an XML dump of relevant backup metadata about the relevant backup job from the Veeam DB. A VBM file will be created for each backup and backup copy job.

<input type="checkbox"/> Name	Date modified	Type	Size
 Backup Job 018045-1.vbm	9/08/2017 12:18 PM	Veeam Backup & ...	15 KB
 Backup Job 018045-12017-08-09T121659.vbk	9/08/2017 12:18 PM	Veeam Backup & ...	31,100 KB

### Why do we create VMBs?

Before VBM files, Veeam would read metadata information from the backup files themselves. As Veeam was deployed into larger and larger environments it soon became apparent that this method of reading metadata information was too I/O intensive and needed to be improved, so the VBM file was born.

### When are VMB files Created?

VBM files are generated multiple times during a backup job run. The exact number depends on a few things such as the number of VMs in a job, how many concurrent tasks are running, whether a transform operation is performed during a given run, etc. Regarding the very first VBM file that is generated, Veeam collects all pertinent backup-related metadata for the running job from the Veeam DB which is then processed and stored in the Veeam.Backup.Manager process memory, this is otherwise known as a 'Full' read of the relevant backup metadata from the DB. This can be quite I/O expensive so all subsequent VBM file generations will use incremental 'Reads' of backup metadata from the Veeam DB to improve performance. This metadata information stored in the Veeam.Backup.Manager process memory is then written to a resulting VBM file on disk. One important thing to note is that while subsequent reads from the Veeam DB will be incremental, the VBM file creation operation is always full, meaning Veeam will always overwrite VBM file in its entirety.

Veeam has quite a bit of intelligence built around optimizing this VBM file generation process, for example, Veeam schedules multiple VBM generations to be merged into one VBM if multiple generation requests coincide in time, this results in less overhead and improved VBM file creation performance.

### **Where are VBM files stored?**

For simple repositories, there will be a single VBM per job located alongside the backup files. If you are using scale-out repositories, you will see a VBM for each job on every extent that the backup job has files residing on.

### **What is inside of a VBM file?**

It's a text-based file which you can open with a text editor and have a look inside. It contains a big description of everything that is in the backup, so all the VMs that are in the backup, all the restore points, application-specific data if you are doing application-aware processing, specific disk drives of the VMs, etc.

```
<BackupMeta Version="4" MetaSchemaVersion="2"
CreationTimeUtc="08/09/2017 02:18:01">
  <Backup Id="0853d842-1999-42e2-8d5b-1b750589b588"
JobId="c9a262ff-8d5d-4149-9499-fb3bc1c052c7" JobName="Backup Job
018045-1" JobType="4000" SourceType="3" TargetType="0"
JobTargetHostId="6745a759-2205-4cd2-b172-8ec8f7e60ef8"
JobTargetHostProtocol="0" RepositoryId="4fdalla9-a3b6-498b-aedd-
d74b57d05513" DirPath="D:\VeeamBackup\Backup Job 018045-1"
MetaFileName="Backup Job 018045-1.vbm" MetaVersion="3"
MetaUpdateTime="08/09/2017 12:18:00" BackupPlatform="6"
CreationTime="08/09/2017 12:16:58" EncryptionState="0" />
  <BackupMetaInfo>
    <Hosts>
      <Host Id="6745a759-2205-4cd2-b172-8ec8f7e60ef8" Moref=""
Name="This server" Type="3" Options="" HostInstanceId="a4ce5334-
b6bc-afb3-1342-f66cd6ba12ab" HostUniqueId="" />
    </Hosts>
    <Storages>
```

## When do we use VBM files?

The VBM file is used when something is wrong with the Veeam DB, so, for example, you create a new VBR, and we want to import the backup file. The VBM file is also critical for scale-out repositories because it tracks which files living on the extent belong to which backup.

## What happens if I don't have a VBM file?

While VBM files are not necessary to import or restore a backup, they will notably improve the backup import process. If you are trying to map existing backups files you will need to recreate: the VBM file though, that process is described here; <https://rhyshammond.com/creating-vbm-from-scratch/>

## Chapter 14

# Efficient Veeam NetApp Backups from Storage Snapshots

**By: Markus Kraus (VMWare vEXPERT / Veeam Vanguard)**

As a part of a larger Veeam project, I was looking for the most efficient Veeam setup with minimal impact to the whole virtualization and storage environment during the backup window. The primary requirement was to deliver constant performance at any time. My tests quickly turned out, that the Veeam NetApp Backup from Storage Snapshots is the best transport mode to reach this goal.

What means efficient in this context:

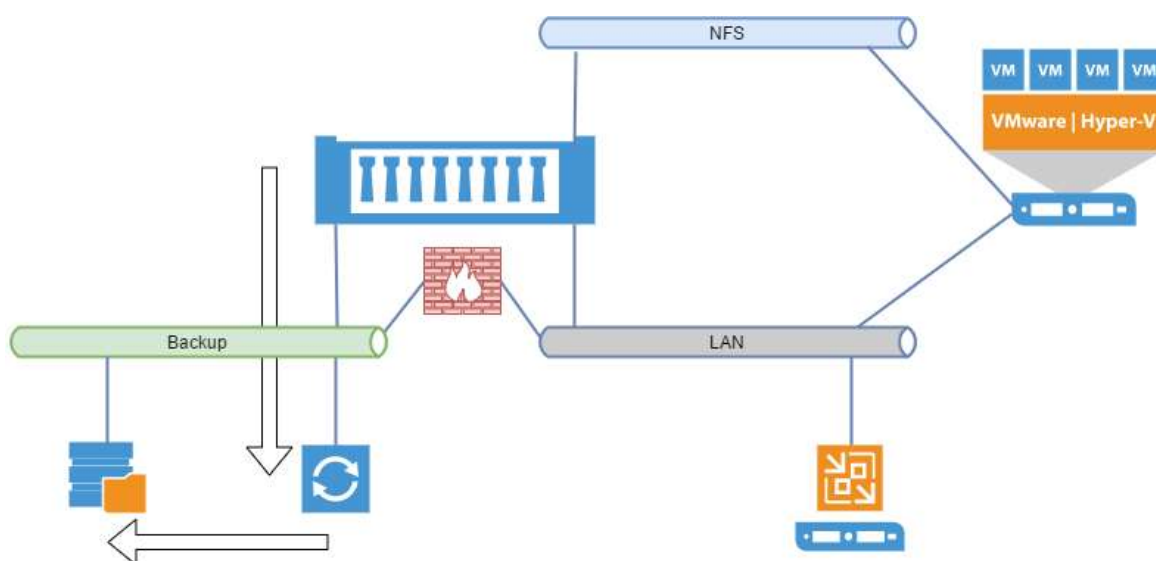
- Minimal impact to the production environment
- Additional costs should be avoided

Main parameters of the infrastructure

- VMware vSphere 6.5
  - NFS Datastores
- NetApp ONTAP 9
- Veeam Backup & Replication 9.5
- Isolated NFS network (no routing, etc.)

## Storage Design for NetApp Backup from Storage Snapshots

The awesome whitepaper NetApp and Veeam Backup & Replication 9.5: Configuration Guide and Best Practices from Stefan Renner shows a lot of concepts of efficient NetApp backups with Veeam. The premium solution is the NetApp integration of SnapMirror / SnapVault with backup from the secondary site. With this method only minimal impact on the components on the production site takes place. However, on the other side, this concept generates additional costs because of the additional array and the data duplication. Unfortunately, that is the reason why my design has to do without NetApp SnapMirror or NetApp SnapVault.



In addition to the base networks, LAN and NFS is a new network for the backup traffic needed within this design. The new backup network connects the Veeam backup proxy with an additional NetApp SVM network interface to transfer the data via Direct Storage Access.

The design only makes sense when dedicated interfaces are used for the Backup, NFS and LAN networks. The SVM which exports the VMware Datastores needs minimum two networks, Backup and NFS. The management network (LAN) is optional.

## Storage Snapshot procedure

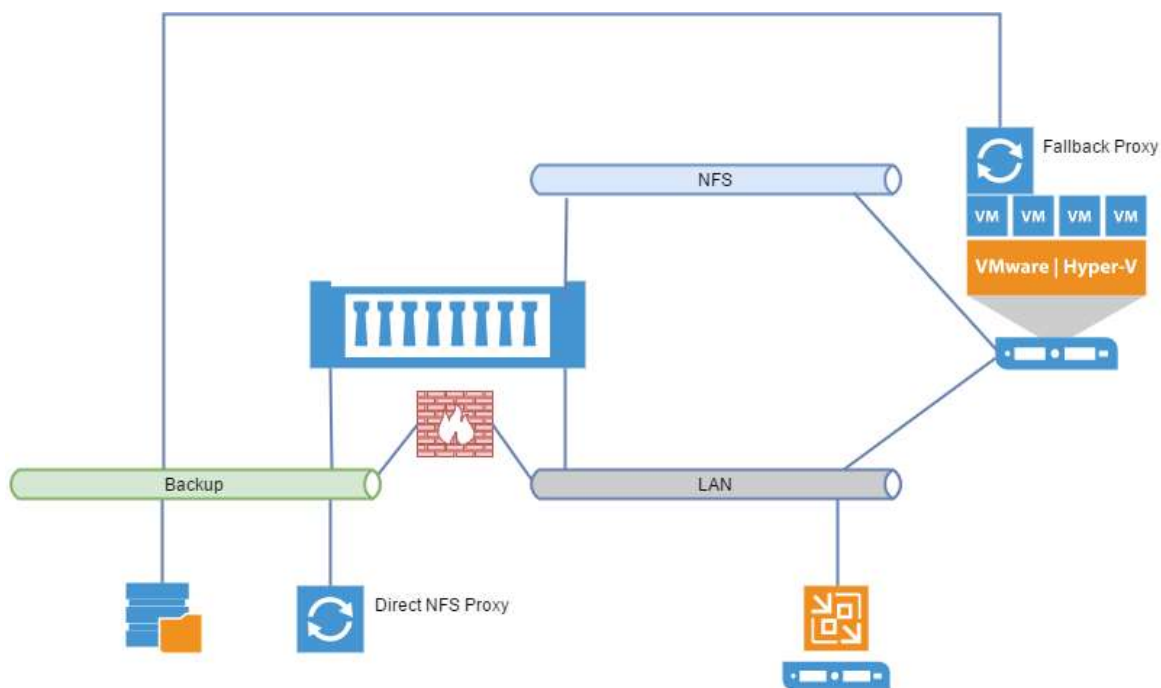
1. Veeam creates a VMware VM Snapshot, if necessary
2. Veeam creates NetApp Volume Snapshot
3. Veeam deletes VMware VM Snapshot, if necessary
4. NetApp creates a Read-Only export of the NetApp Volume Snapshot
5. Veeam Proxy reads Snapshot data with the integrated NFS Client
6. Veeam deletes NetApp Volume Snapshot

As you can see in the procedure is one of the main benefits of this concept that the VM snapshots are only a short period open compared to the Hot Add or network transfer mode. The result is a way quicker consolidation of the VM snapshot and a reduced impact to the VM itself. In combination with the backup traffic transfer over the dedicated NetApp interface is the impact to the production environment minimized.

Disadvantages of the concept:

- Restore will be processed via network transport mode
- If a VM Snapshot exists before the Veeam Backup starts the network transport mode will also be used

It is possible to enhance the concept with an additional Proxy for virtual appliance transport mode; this proxy will be used if VM Snapshot exists. If the additional backup proxy will be placed in the backup network the traffic through the firewall will be dramatically reduced.



Even with the additional Hot Add Proxy, the restore process will use network transport mode. The only way (at the moment) to prevent Veeam from falling back from Direct NFS Restore to Network Transport Mode is access to the NFS Network with the Direct NFS Proxy.

## Storage Snapshot Configuration

For the setup of this backup method, we need to take a look at three components, the NetApp SVM configuration, the Veeam setup and the vSphere infrastructure.

NetApp SVM configuration

To serve NFS Datastores to ESXi hosts from a NetApp SVM only one interface in the NFS network is necessary. However, for this concept one additional interface for the backup traffic must be added.

Network						
Subnets Network Interfaces Ethernet Ports Broadcast Domains FC/FCoE Adapters IPspaces						
Create Edit Delete Status Migrate Send to Home Refresh						
Interface Name	Storage Virtual Machine (SV...	IP Address/WWPN	Current Port	Is Home Port	Data Protocol Access	
svm-nfs_backup	svm-esxi	192.168.4.104	netapp-01:e0b	Yes	nfs	
svm-nfs_data	svm-esxi	10.100.10.100	netapp-01:e0d	Yes	nfs	

In my test environment, the interface SVM-nfs\_data represents the NFS Network and SVM-nfs\_backup the backup network. As mentioned earlier, both interfaces use different Ethernet ports as backing.

To serve NFS traffic to the ESXi hosts and the backup proxy both networks need to be added to the export policy.

Add Edit Delete Move Up Move Down Refresh					
Rule Index	Client	Access Protocols	Read-Only Rule	Read/Write Rule	
1	192.168.4.102	NFS	Any	Never	
3	10.100.10.101	NFS	Any	Any	

Veeam Setup

- 1. The first step to enable the Storage Integration is adding the NetApp cluster to the Veeam Backup & Replication server (in my test lab the version 9.5 Update 3 is used). At this point, only the management traffic between Veeam Backup & Replication server and NetApp cluster IP happens. The Direct NFS Proxy will be used in further steps like the scan of the NetApp Volumes.

The image displays two sequential screenshots of the 'Edit NetApp Storage' wizard in the Veeam Backup & Replication console.

**Top Screenshot: Name Step**

- Title:** Edit NetApp Storage
- Section:** Name
- Instruction:** Register NetApp storage by specifying DNS name or IP address.
- Fields:**
  - Management server DNS name or IP address:** 192.168.3.103
  - Description:** Lab NetApp Cluster
- Buttons:** < Previous, Next >, Finish, Cancel

**Bottom Screenshot: Credentials Step**

- Title:** Edit NetApp Storage
- Section:** Credentials
- Instruction:** Type in storage administrator credentials.
- Fields:**
  - Credentials:** admin (Created by Powershell at 29.06.2017 20:57:28 for h...)
  - Protocol:** HTTPS
  - Port:** 443
- Buttons:** < Previous, Next >, Finish, Cancel

2. It is also possible to use the default option "Create required NFS export rules automatically." I removed this options to do some further tests with modified export

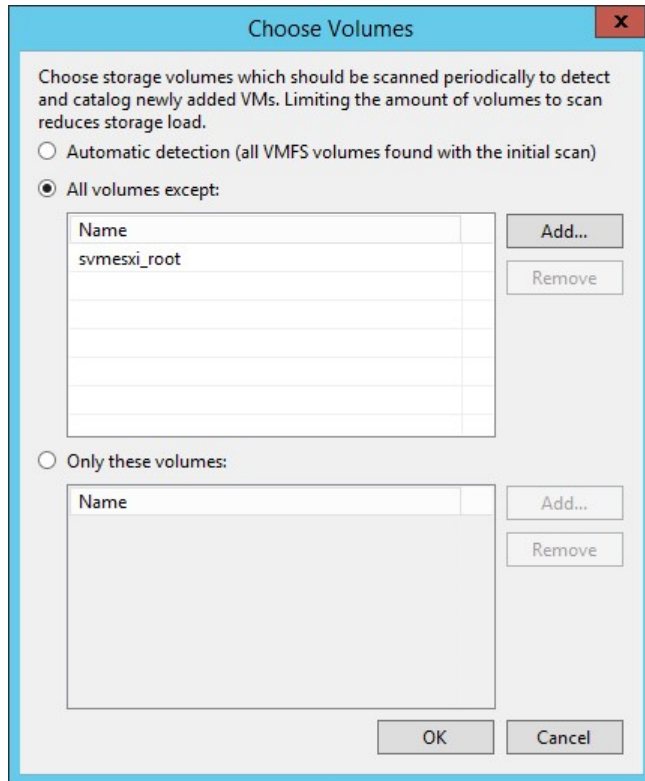
rules.

The screenshot shows the 'Edit NetApp Storage' dialog box with the 'Access Options' tab selected. The dialog has a sidebar on the left with tabs: 'Name', 'Credentials', 'Access Options' (selected), and 'Summary'. The main area contains the following options:

- Protocol to use:**
  - ☐ Fibre Channel (FC)
  - ☐ iSCSI
  - ☒ NFS
  - ☐ Create required NFS export rules automatically
- Volumes to scan:**
  - Text field: Rescan all volumes except selected
  - Button: Choose...
- Backup proxies to use:**
  - Text field: Veeam-02.lab.local
  - Button: Choose...

At the bottom, there are four buttons: '< Previous', 'Next >' (highlighted with a mouse cursor), 'Finish', and 'Cancel'.

3. Excluding all NetApp SVM root volumes will speed up the scan process in larger environments a little bit.



4. Especially the existing NetApp licenses are very interesting; the available licenses have influenced the possible restore options on the NetApp side. You can find the procedure on page 11 (Restore: NFS Protocol, (ONTAP)) in the whitepaper NetApp and Veeam Backup & Replication 9.5: Configuration Guide and Best Practices from Stefan Renner.
5. With the Veeam PowerShell SnapIn we can gather some more details about the managed NetApp Cluster:

```
PS C:\> (Get-NetAppHost).NAoptions
```

```
ConnectionOptions          :  
Veeam.Backup.SanPlugin.NetApp.CDomNaHostConnectionOptions  
  
DomContainer               :  
Veeam.Backup.Common.CDomContainer  
  
HostType                   : NaCluster  
  
IsMetroClusterEnabled     : False  
  
MetroClusterPartner       :  
  
IsMetroClusterAlive       : False  
  
IsHAPairEnabled            : False  
  
VolumesRescanMode         : ExceptExcluded  
  
SelectedSanProtocols      : NFS  
  
CreateNfsExportRulesAutomatically : False  
  
IsRescanProxyAutoSelect   : False  
  
HAPairPartner              :  
  
IsNeedToShowRetentionForSnapMirror : False  
  
License                    : FlexClone, SnapRestore, Fcp,  
Iscsi, Nfs, SnapVaultPrimary, SnapVaultSecondary,  
                             SnapMirror  
  
IsVfilerLicensed          : False  
  
IsFlexCloneLicensed       : True  
  
IsSnapRestoreLicensed     : True  
  
IsFcpLicensed             : True  
  
IsIscsiLicensed           : True
```

IsNfsLicensed	: True
IsSnapVaultPrimary	: True
IsSnapVaultSecondary	: True
IsSnapMirror	: True
IsHAPairLicensed	: False

- The next component to configure is the Direct NFS Proxy. This proxy type has unlike the Hot Add proxy no dependency to the vSphere VM that needs to be backed up. This proxy type can even be a physical server or can run in a different vCenter.

**Edit VMware Proxy**

**Server**  
Choose server for new backup proxy. You can only select between Microsoft Windows servers added to the managed servers which are not proxies already.

**Server**  
Choose server:  
Veeam-02.lab.local Add New...

**Proxy description:**  
Created by Veeam-01\Administrator at 04.01.2018 08:04.

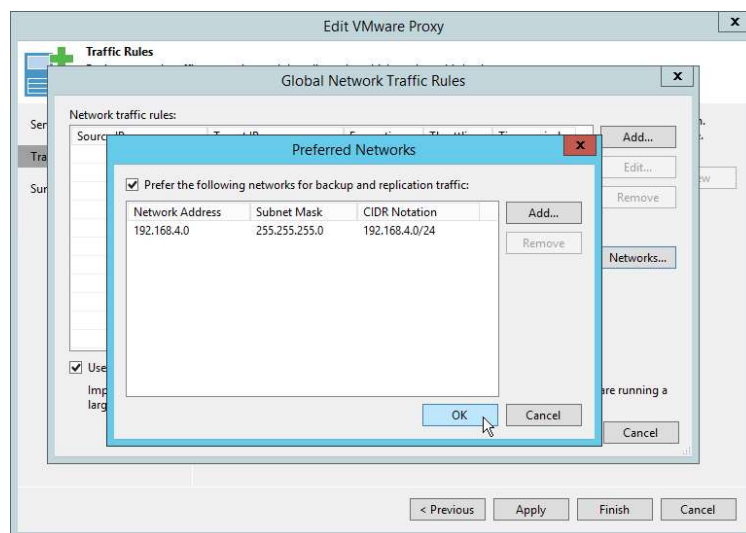
**Transport mode:**  
Direct storage access Choose...

**Connected datastores:**  
Manual selection Choose...

**Max concurrent tasks:**  
2

< Previous Next > Finish Cancel

- Within this design, the configuration of a Preferred Backup Network is an optional step. It is only necessary when the NetApp SVM can be accessed via different networks or proxies.



8. The final setup uses these backup proxies, one for Direct Storage Access and one für Hot Add:

```
PS C:\> (Get-VBRViProxy -Name Veeam-02.lab.local).Options
```

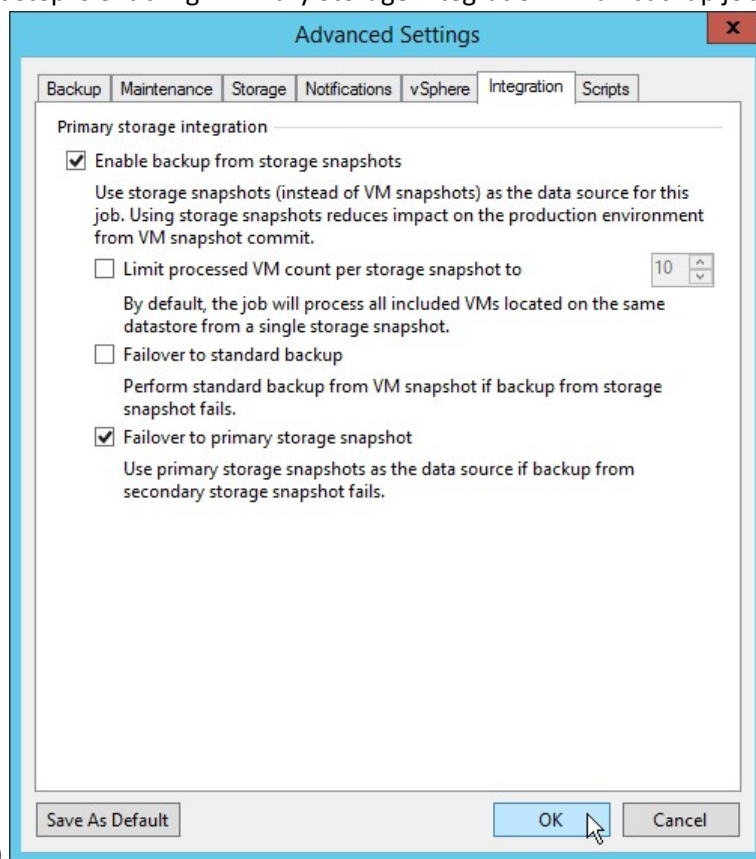
```
TransportMode           : San
FailoverToNetwork       : True
UseSsl                   : False
IsAutoVddkMode          : True
IsAutoDetectDisks       : False
MaxTasksCount           : 2
```

```
IsAutoDetectAffinityRepositories : True

PS C:\> (Get-VBRViProxy -Name Veeam-03.lab.local).Options

TransportMode           : HotAdd
FailoverToNetwork       : True
UseSsl                   : False
IsAutoVddkMode           : True
IsAutoDetectDisks       : True
MaxTasksCount           : 1
IsAutoDetectAffinityRepositories : True
```

9. The last step is enabling “Primary Storage integration” in all backup jobs (it’s the default



option).

Is this option in one or more jobs not enabled but the Veeam and NetApp setup allow Veeam NetApp Backup from Storage Snapshots a Direct NFS Backup without Storage Snapshot will be done (thanks for the clarification Niels Engelen!).

Task log without Storage Integration:

```
Using backup proxy Veeam-02.lab.local for disk Festplatte 1 [nfs]
```

Task log with Storage Integration:

Using backup proxy Veeam-02.lab.local for retrieving Festplatte 1 data from a storage snapshot

### VSphere Infrastructure

My test lab only has minimalistic vSphere Setup. However, even in a production setup, no additional configuration is necessary to leverage Veeam NetApp Backup from Storage Snapshots.

esxi-01.lab.local

AKTIONEN

Übersicht

Überwachen

Konfigurieren

Berechtigungen

VMs

Ressourcenpools

Datenspeicher

Netzwerke

Name	Status	Typ	Gerät	Datenspe...	Kapazität	Frei
 esx-01_local	 Normal	VMFS 6	mpx.vmhba0:C0:T1:L0:1		39,75 GB	35,17 GB
 vol_vmware_01	 Normal	NFS 3	10.100.10.100:/vol_vmware_01		675 GB	65,88 GB

Note: For more details about a proper VMware NFS setup with NetApp, please refer to the NetApp TR- 4597 VMware vSphere with ONTAP: <https://www.netapp.com/us/media/tr-4597.pdf>

## Chapter 15

# Sizing your Backup Storage and Saving Money

**By: Eugene Kashperovetskyi (VMWare vEXPERT / Veeam Vanguard)**

My parents made sure the phrase “planning is the key to success” became the mantra of my childhood. Moreover, for a good reason too. Whether it was a hiking weekend, fishing trip or an exotic vacation, planning was essential to the pleasant memories I took from these adventures directly because a lack thereof very well could’ve led to disastrous results.

As it turns out, it’s now a mantra I apply every day toward sizing cloud environments.

When working with clients, a topic we frequently discuss is the current state of their environment and its anticipated growth over a period. More often than not, this is a “guesstimate” science due to the complexity of most setups and the lack of statistical data for greenfield deployments. These facts lead some IT professionals to manage their environments via the “learn and adjust on the fly” approach. While that may work in some scenarios, it’s hardly a reliable method for managing critical business assets. Having a ballpark idea of future state requirements goes a long way.

The same principle applies to backup and disaster recovery strategy. For just a small time investment, proper planning can save you thousands of dollars.

Covered in six steps below, the goal of this exercise is to help you understand two things: 1) the process for making sure your backup jobs will work given specific retention requirements, recovery targets, and backup windows; and 2) the method for getting just enough space for your current needs while simultaneously creating a roadmap with expansion milestones.

(If you want to skip all the details and access a very helpful sizing calculator, jump down to step four!)

## Step One – Assess the Job Setup and Key Requirements

Let's assume we have five critical workloads that need to be backed up. They each require high availability and a long-term retention policy:

- VM1 – Active Directory
- VM2 – SQL server
- VM3 – App server
- VM4-5 – Web frontend

First off, we need to make a note of our RPO/RTO and data retention policy. We can't arrive at an accurate sizing estimate without them.

In this case, all VMs require the same RPO/RTO, as they belong to the same application group:

- RPO – 6 hours
- RTO – 24 hours

The retention policy is as follows:

Last 7 days of backups

Monthly backup job copies

Next, we need to gather some key info regarding the VMs, data and backup windows.

–Operating System — All Guest OS's are Microsoft Windows 2012 R2

–Change Rate — All data on the VMs, except Active Directory Domain controller, have a change rate of 5%

–Source Data Size — The total size of our VMs is 800GB:

- VM1 – 100GB
- VM2 – 250GB

- VM3-5 – 150GB

–Backup Windows — The idle/non-busy time periods deemed suitable for backups are established as such: Central Time Zone: 7-9 a.m., 1-2 p.m., 5-6 p.m., and 1-2 a.m.

–Read/Write — The storage can handle up to 300MB/s reads and 300MB/s writes at a given time

–Uplinks — The environment is configured to use 2 x 1Gbps uplinks to the Backup server, no LACP — 1Gbps maximum at a time

## Step Two – Make sure the backup Windows Align with RPO/RTO

Now we have to review the above information to ensure there aren't any potential conflicts with our requirements.

In this scenario, the backup windows conflict with the 6-hour RPO. The time between 5 p.m. to 1 a.m. is spaced by eight hours, which causes the violation. On the other hand, 1-5 p.m. is spaced by four hours and is not considered a violation, since the RPO is met (the oldest restore point is no older than 6 hours).

To mitigate the risk of an RPO policy violation, we decide to re-arrange the non-busy periods and adjust application-level tasks to allow for a backup schedule within the following windows: 7-9 a.m., 1-2 p.m., 7-9 p.m., and 1-2 a.m. This may require a meeting or two to sort out, but that's why we're planning!

To ensure the 24-hour RTO is achievable, simply:

1. Take the total size of the initial, Day 1 data set: 800GB1
2. Determine how long the backup will take based on the storage read/write speed:  
300MB/s read and 300MB/s write

Because our storage also handles production VMs, let's estimate that only 50% of the top capabilities will be available for backup/restore purposes (150MB/s).

3. Initial Backup time:  $800\text{GB} / 150\text{MBps} = 1.5$  hours

Plenty of time to spare!

## Step Three – Make Sure the network can handle the Backup Jobs

We also want to make sure our network equipment can tolerate the extra load.

Out of two available 1Gbps uplinks, only half the overall capacity (1Gbps) can be allotted for backup purposes. Meaning, the initial backup will take an estimated two hours. We'll use this value for all further calculations on timing over 1.5 hours at the storage level.

One last thing before getting to the sizing: Before we can calculate the total storage requirements, we have to make sure our 1-hour windows, occurring four times daily, are sufficient for copying incremental backups, based on the rate of change.

—Every incremental backup is expected to contain up to 5% changed data, which means we need to copy 40GB every six hours. Based on our 150MBps benchmark, this will take approximately 15-20 min.

All good!

## Step Four – Calculate Total Backup Storage Requirements

How much space do we need to store the backup data? As with everything in IT, that depends on the type of data and the configuration chosen for the backups.

In our scenario, two types of jobs will run the Backup Job (our last seven days of backups) and the Backup Job Copy (our monthly backup copies).

To make our calculations, we'll use this handy calculator developed by one of Veeam's talented team members:

<http://rps.dewin.me/>

Calculating Storage for the Backup Job

The default backup job type in Veeam v9 is Forever Forward Incremental, in which no Synthetic Full backups are created. This type of job allows for substantial space savings, allowing us to avoid the need for longer backup windows and restore times.

Using the info from our assessment, fill out the calculator fields as follows:

## The Restore Point Simulator






Current version : 0.3.2  
Feedback via @tdewin or on [GitHub](#)  
RPS heavily relies on some opensource [javascript frameworks](#)

### Quick Presets

Forever Incremental ▼





























☐ Manual Run

### Configuration

Style	Incremental ▼	
Used Size GB	800	
Retention Points	28	
Change Rate	5% Optimistic ▼	
Data left after reduction	50% (100GB > 50GB) Conservative ▼	
Interval	Every 6 Hours ▼	
Time Growth Simulation <input type="checkbox"/>	1 Year ▼	10% ▼ 

This input will deliver the following retention interval schedule, with a total storage size of 1,360GB. This includes 420GB workspace.

### Result

Retention		File	Size	Modify Date	Point Date
28		full.vbk	<u>400 GB</u>	2016-04-20 We 22	2016-04-20 We 22
27		incremental.vib	<u>20 GB</u>	2016-04-21 Th 04	2016-04-21 Th 04
26		incremental.vib	<u>20 GB</u>	2016-04-21 Th 10	2016-04-21 Th 10
25		incremental.vib	<u>20 GB</u>	2016-04-21 Th 16	2016-04-21 Th 16
24		incremental.vib	<u>20 GB</u>	2016-04-21 Th 22	2016-04-21 Th 22
23		incremental.vib	<u>20 GB</u>	2016-04-22 Fr 04	2016-04-22 Fr 04
22		incremental.vib	<u>20 GB</u>	2016-04-22 Fr 10	2016-04-22 Fr 10
21		incremental.vib	<u>20 GB</u>	2016-04-22 Fr 16	2016-04-22 Fr 16
20		incremental.vib	<u>20 GB</u>	2016-04-22 Fr 22	2016-04-22 Fr 22
19		incremental.vib	<u>20 GB</u>	2016-04-23 Sa 04	2016-04-23 Sa 04
18		incremental.vib	<u>20 GB</u>	2016-04-23 Sa 10	2016-04-23 Sa 10
17		incremental.vib	<u>20 GB</u>	2016-04-23 Sa 16	2016-04-23 Sa 16
16		incremental.vib	<u>20 GB</u>	2016-04-23 Sa 22	2016-04-23 Sa 22
15		incremental.vib	<u>20 GB</u>	2016-04-24 Su 04	2016-04-24 Su 04
14		incremental.vib	<u>20 GB</u>	2016-04-24 Su 10	2016-04-24 Su 10
13		incremental.vib	<u>20 GB</u>	2016-04-24 Su 16	2016-04-24 Su 16
12		incremental.vib	<u>20 GB</u>	2016-04-24 Su 22	2016-04-24 Su 22
11		incremental.vib	<u>20 GB</u>	2016-04-25 Mo 04	2016-04-25 Mo 04
10		incremental.vib	<u>20 GB</u>	2016-04-25 Mo 10	2016-04-25 Mo 10
9		incremental.vib	<u>20 GB</u>	2016-04-25 Mo 16	2016-04-25 Mo 16
8		incremental.vib	<u>20 GB</u>	2016-04-25 Mo 22	2016-04-25 Mo 22
7		incremental.vib	<u>20 GB</u>	2016-04-26 Tu 04	2016-04-26 Tu 04
6		incremental.vib	<u>20 GB</u>	2016-04-26 Tu 10	2016-04-26 Tu 10
5		incremental.vib	<u>20 GB</u>	2016-04-26 Tu 16	2016-04-26 Tu 16
4		incremental.vib	<u>20 GB</u>	2016-04-26 Tu 22	2016-04-26 Tu 22
3		incremental.vib	<u>20 GB</u>	2016-04-27 We 04	2016-04-27 We 04
2		incremental.vib	<u>20 GB</u>	2016-04-27 We 10	2016-04-27 We 10
1		incremental.vib	<u>20 GB</u>	2016-04-27 We 16	2016-04-27 We 16
			940 GB		
Work Space			+420 GB		
			<u>1360 GB</u>		

The dates and times on the right represent the 6-hour intervals over the 7-day retention span. While knowing those dates and times aren't as critical for short-term backup retention policy, the impact it has over the long term is enormous, as we'll see in our provisioning plan.

### Calculating Storage for the Backup Copy Job

For the monthly Backup Copy Jobs, use these inputs.

Quick Presets

Backup Copy Job

☐ Manual Run

Simulate

Configuration

Style  
Used Size GB  
Retention Points  
Change Rate  
Data left after reduction  
Interval  
Time Growth Simulation ☐

Backup Copy Job

800

0

5% Optimistic

50% (100GB > 50GB) Conservative

Daily

1 Year

10%

Backup Copy Job Specific

Weekly

0

Monthly

12

Quarterly

0

Yearly

0

That will produce the following:

Result

Retention		File	Size	Modify Date	Point Date
-1 0W 12M 0Q 0Y		full.vbk	<u>400 GB</u>	2016-05-01 Su 22	2016-05-01 Su 22
-1 0W 11M 0Q 0Y		full.vbk	<u>400 GB</u>	2016-06-05 Su 22	2016-06-05 Su 22
-1 0W 10M 0Q 0Y		full.vbk	<u>400 GB</u>	2016-07-03 Su 22	2016-07-03 Su 22
-1 0W 9M 0Q 0Y		full.vbk	<u>400 GB</u>	2016-08-07 Su 22	2016-08-07 Su 22
-1 0W 8M 0Q 0Y		full.vbk	<u>400 GB</u>	2016-09-04 Su 22	2016-09-04 Su 22
-1 0W 7M 0Q 0Y		full.vbk	<u>400 GB</u>	2016-10-02 Su 22	2016-10-02 Su 22
-1 0W 6M 0Q 0Y		full.vbk	<u>400 GB</u>	2016-11-06 Su 22	2016-11-06 Su 22
-1 0W 5M 0Q 0Y		full.vbk	<u>400 GB</u>	2016-12-04 Su 22	2016-12-04 Su 22
-1 0W 4M 0Q 0Y		full.vbk	<u>400 GB</u>	2017-01-01 Su 22	2017-01-01 Su 22
-1 0W 3M 0Q 0Y		full.vbk	<u>400 GB</u>	2017-02-05 Su 22	2017-02-05 Su 22
-1 0W 2M 0Q 0Y		full.vbk	<u>400 GB</u>	2017-03-05 Su 22	2017-03-05 Su 22
-1 0W 1M 0Q 0Y		full.vbk	<u>400 GB</u>	2017-04-02 Su 22	2017-04-02 Su 22
1 (1)		full.vbk	<u>400 GB</u>	2017-04-03 Mo 22	2017-04-02 Su 22
			5200 GB		
			Work Space		
			+420 GB		
			5620 GB		

We now know the total space needed to store all of the data given the requirements: 5,620GB + 1,360GB = ~7TB.

## Step Five – Map your Provisioning Plan

Since the total space can be estimated before we even begin the first job, why not provision the 7TB and call it a day? Simple: Because that would cost us some serious money.

We know that on Day 1 our Backup Copy Job will not require all the estimated space. Therefore, we can provision a smaller amount of storage, sufficient for the needs of the first four months and then add the space gradually.

For example, if we choose not to provision based on the milestones accounted for above, we'd have to request all 7TB of space starting Day 1, costing us about \$7,200/year.

Compare that to provisioning the space gradually:

Day 1 – provision space required for the first 3 months –  $4 \times 400\text{GB} = 1600\text{ GB}$ , \$160/mo, \$480/period

Day 90 – provision space required for 6 months –  $7 \times 400\text{GB} = 2800\text{ GB}$ , \$280/mo, \$840/period

Day 180 – provision space required for 9 months –  $10 \times 400\text{GB} = 4\text{TB}$ , \$400/mo, \$1200/period

Day 240 – provision all space – \$1800/period

The grand total for the year in this case is \$4,320 – nearly \$3,000 less than “no planning ahead” plan.

## Step Six – Plan How to Spend your Savings

Storage planning is not new conceptually, but you'd be surprised how often it's skipped during development and deployment stages. Don't make that mistake.

Planning not only allows you to validate your decisions at an early stage of the project and also gives company management confidence that your management is efficient and sound. The cherry on top – all extra savings that can be applied once you have a definitive set of milestones and know precisely when to provision and how much to add/remove. In our example, it may have saved us \$3000/year, but for many organizations, savings can balloon to 10X that size.

## Chapter 16

# Join us at MVPDays and meet great MVP's 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](http://www.checkyourlogs.net)

[www.mvpdays.com](http://www.mvpdays.com)

## Video Training

For video-based training, see the following site:

[www.mvpdays.com](http://www.mvpdays.com)

## Live Instructor-led Classes

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

- [www.truesec.com](http://www.truesec.com)
- [www.checkyourlogs.net](http://www.checkyourlogs.net)

## Consulting Services

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

