

# Extend your Availability strategy to the cloud with Veeam and Microsoft Azure

A configuration guide for  
deploying Veeam Cloud Connect  
*for the Enterprise* in Microsoft  
Azure Marketplace

1 + 1 = 3  
Microsoft + Veeam  
Better Together

# Contents

<b>Introduction.....</b>	<b>4</b>
<b>Architecture and scenario.....</b>	<b>4</b>
Scenario .....	4
Prerequisites.....	4
Single VM Veeam Cloud Connect architecture .....	5
Roles and concepts .....	5
<b>Azure preparations.....</b>	<b>8</b>
Sign up for an Azure subscription.....	8
Billing.....	8
Log in to the Management Portal .....	9
Create the Veeam Cloud Connect <i>for the Enterprise VM</i> .....	10
Configure a DNS name .....	19
Create and assign a repository disk .....	20
<b>Configure Veeam Backup &amp; Replication .....</b>	<b>23</b>
Initial configuration .....	23
Configure Veeam Backup & Replication repository .....	24
Configure Veeam Cloud Connect .....	28
Manage a certificate.....	28
Cloud gateway.....	29
Create tenant.....	29
<b>Distributed Veeam Cloud Connect infrastructure .....</b>	<b>32</b>
Scenario .....	33
Deploy and configure additional repositories .....	33
Configure virtual networking in Azure .....	55
Deploy and configure additional cloud gateways.....	65

<b>The on-premises side!</b> .....	<b>82</b>
Connect to a service provider .....	83
Create a Backup Copy job .....	85
<b>Restore</b> .....	<b>89</b>
Test 1: Recovering an Active Directory item .....	90
Test 2: Recovering a file .....	93
Test 3: Recovering a virtual disk .....	95
Test 4: Recovering the entire virtual machine .....	99
<b>Conclusion</b> .....	<b>104</b>
<b>Appendix A: Using WAN acceleration</b> .....	<b>105</b>
<b>About the Authors</b> .....	<b>107</b>

## Introduction

This white paper explains how IT administrators can use Veeam® Availability Suite™ deployed via Microsoft Azure Marketplace to build cloud backup services to serve internal customers, such as different subsidiaries or departments, referred to as tenants later in this document.

Available within Microsoft Azure Marketplace, [Veeam Cloud Connect for the Enterprise](#) allows IT administrators to configure multi-tenant backup as a service within just a few minutes. Veeam Cloud Connect provides all the capabilities you need to manage cloud backup repositories, including setting up tenants, assigning quotas and tracking usage. Cloud backup repositories for each tenant are completely isolated from one another and allow internal customers to store their off-premises backup files safely. Veeam Availability Suite fully supports the encryption of all backup files. Just as important, Veeam customers can connect to their cloud repositories directly from the Veeam Backup & Replication™ console — seamlessly, securely and with a standard internet connection. All traffic is sent over a single consolidated port utilizing a TLS connection, and no VPN is required.

Ability for subsidiaries and departments to send backup copies to a Veeam Cloud Connect repository is included in the solution they already use: [Veeam Availability Suite](#), [Veeam Backup & Replication](#)™, [Veeam Backup Essentials](#)™ and [Veeam Agent for Microsoft Windows](#). There's no need to change the solution on premises. Internal IT operating the repositories will need to purchase licenses for Veeam Cloud Connect for the Enterprise running within Microsoft Azure. Eligible customers can purchase them from any Veeam reseller as a yearly subscription per VM. More information on purchasing [here](#).

## Architecture and scenario

It is important to understand the key components and usage scenarios of Veeam Cloud Connect *for the Enterprise* solution in Azure.

### Scenario

This white paper will walk you through the basic, first-time setup and deployment of Veeam Cloud Connect *for the Enterprise* within Azure Marketplace, and a more advanced, distributed and scale-out one. This paper will also simulate the onboarding of a subsidiary (or tenant), and demonstrate how to test your service.

By the end of this white paper, the tenant will have on-premises, local backups as well as backups within the Azure cloud infrastructure that you, IT administrators, will offer. The tenant will also have performed four types of restores to ensure success as well as correct functionality.

### Prerequisites

The scenario covered in this white paper is built on the following prerequisites and configurations:

1. You should have a basic understanding of Veeam Cloud Connect, its capabilities and features. Before you proceed, please familiarize yourself with the information and resources available: <http://www.veeam.com/cloud-connect.html>.
2. For the on-premises infrastructure:
  - The Veeam Backup & Replication server is deployed and functioning within the on-premises infrastructure.
  - The infrastructure is running on Microsoft Hyper-V 2016, but the same concepts and processes apply to the VMware vSphere scenarios as well.
  - Backups of one or more virtual machines (VMs) are taken daily or more frequently. In this scenario, the tenant protects several VMs, with both Windows and Linux guest operating systems.
  - The Veeam Backup & Replication server has internet connectivity.

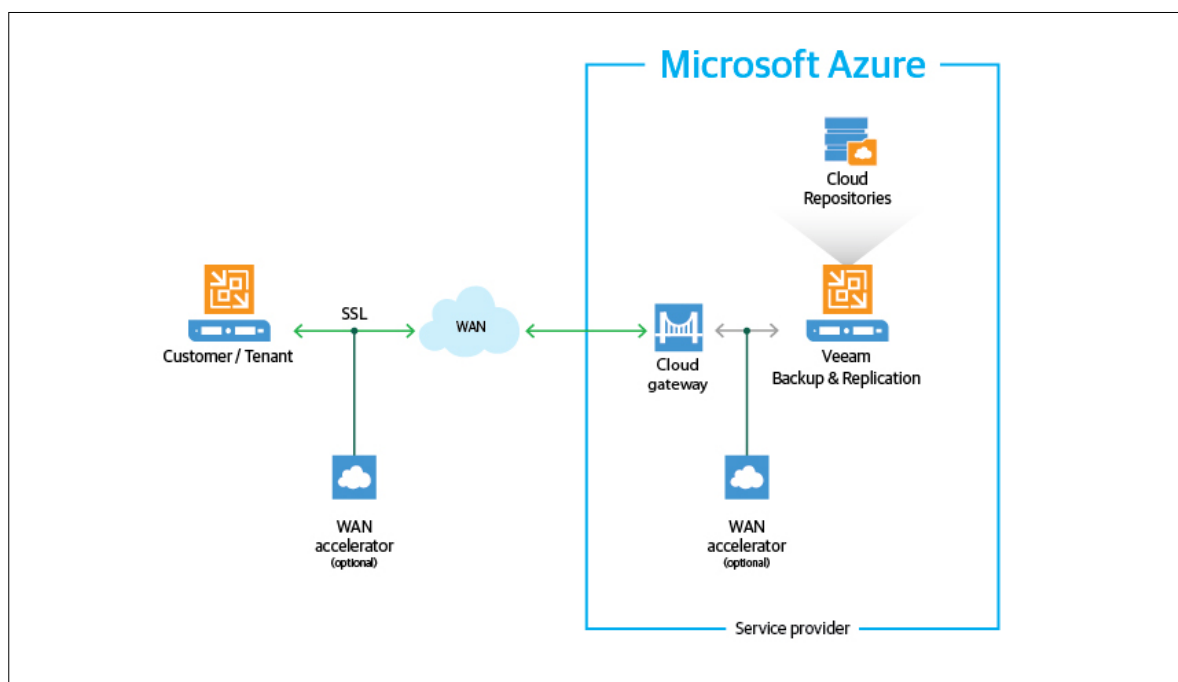


3. For the cloud infrastructure:

- An active and current Microsoft Azure subscription. NOTE: If you don't have a subscription, please see [Sign up for an Azure subscription](#) in this white paper.
- The customer must have a Microsoft Enterprise Agreement (EA) or a VMware Enterprise License Agreement (ELA).

## Single VM Veeam Cloud Connect architecture

For this scenario, the Veeam Cloud Connect setup will be a VM deployed from Microsoft Azure Marketplace, and all the different roles will be configured within this single machine. While this is not an ideal setup for a production environment, we will start with a simple configuration that will fit small deployments. Plus, when you start with a single VM to begin, it is very simple and easy to scale-out to multiple VMs later by deployment through Azure Marketplace. With this, you are now able to distribute the Veeam roles to those servers. This scenario is covered in the [Distributed Veeam Cloud Connect infrastructure](#) section.



## Roles and concepts

### Cloud infrastructure

Data communication in Azure is between two parties – the Veeam Cloud Connect *for the Enterprise* environment and the Veeam Backup & Replication infrastructure running on premises.

- The Veeam Cloud Connect *for the Enterprise* environment is an infrastructure running within Microsoft Azure that provides cloud repository services.
- The Veeam Backup & Replication infrastructure running on premises is sending the VM data off site. This data is stored as an image-based backup in the cloud repository within the Azure environment.

Within Azure, the IT administrators will configure the Veeam Cloud Connect infrastructure. This environment is needed to provide Backup as a Service (BaaS) to a subsidiary or department. As part of this process, the IT administrators will take the following steps:

- Determine their backup repository requirements. This storage will be used as the tenant's cloud repositories.
- Set up SSL certificates to enable secure communication in the Veeam Cloud Connect infrastructure
- Create cloud gateways
- Create tenant user accounts
- Manage tenant's accounts and data for proper functionality of the Veeam Cloud Connect infrastructure

### **Tenant**

The tenant will perform the following tasks:

- Connect to the Veeam Cloud Connect infrastructure in order to use cloud repositories
- Configure and run backup jobs and Backup Copy jobs, and perform restore operations targeted at Veeam Cloud Connect repositories

### **Veeam Cloud Connect infrastructure**

To expose cloud repository resources to tenants, IT administrators must configure the Veeam Cloud Connect infrastructure. The Veeam Cloud Connect infrastructure is comprised of the following components:

- Components on the cloud side:
  - Veeam Backup & Replication server providing Veeam Cloud Connect services
  - SSL certificate
  - Cloud repository
  - Cloud gateway
  - Target WAN accelerator [optional]
- Components on tenant side:
  - Tenant's Veeam backup server
  - Source WAN accelerator [optional]

### **Cloud Veeam Backup & Replication environment**

The Veeam Cloud Connect infrastructure runs as a component within their Veeam Backup & Replication server within Azure, and acts as the configuration and control center for the solution.

On this server, the Veeam backup cloud service runs a Microsoft Windows service that is responsible for:

- Validating tenant's credentials and access rights to assigned resources
- Providing access to the cloud repository for tenants
- Controlling transport services that work with the cloud repository
- Communicating with the Veeam Backup & Replication database

### **SSL certificate**

SSL certificates are not being used for encrypting data stored on the cloud repository. If the tenant wants to encrypt data, he or she needs to enable encryption within their on-premises Veeam Backup & Replication environment. This task is performed at the proxy level before data is copied to the cloud repository. This should be considered an absolute must for a tenant. Sending unencrypted backup files to a cloud platform would allow for increased chances of data compromise.

## **Cloud repository**

The cloud repository is a storage location within Azure where tenants store their VM data. Tenants can use the cloud repository as a target for their respective Veeam backup jobs as well as their Veeam Backup Copy jobs. The data stored within the cloud repository is also treated as a source from which they can restore their data.

The cloud repository is a multi-tenant repository configured within the cloud backup infrastructure. It is built on top of a standard Veeam Backup & Replication repository. Being a multi-tenant storage resource, the cloud repository still appears as a logically separate backup repository to each and every tenant. Data in the underlying cloud repositories is segregated and isolated. This gives every tenant their own underlying repository where their data is stored. Tenants do not have access, or visibility to the other tenants' data within the cloud repository. Each tenant is provided with access to their data resulting in a complete multi-tenant experience. These tenants have no access to data belonging to other tenants residing on the same cloud repository.

## **Cloud gateway**

The Veeam Cloud Connect infrastructure is configured within Microsoft Azure. This configuration is transparent from the tenant view. Tenants are only aware of cloud repositories that they are provisioned to utilize and use them as they would use a regular local backup repository. Veeam Backup & Replication servers on the tenant side do not communicate with the cloud repository directly. Data communication and transfer in the cloud is carried out via one or more cloud gateways.

The cloud gateway is a service running on a Windows server that resides on the cloud side and acts as a communication point in the cloud. The cloud gateway routes commands and traffic between the Microsoft Azure partner, tenants and the cloud repository. The cloud gateway is deployed within Azure and requires internet access as these are the devices that are used as the main point of entry for all inbound and outbound traffic. In larger deployments, IT administrators would deploy multiple cloud gateway servers and optionally configure a single DNS name with round robin for the pool of cloud gateway servers – the cloud gateway service will automatically load balance.

## **WAN accelerator [optional]**

WAN accelerators are an optional component in the Veeam Cloud Connect infrastructure. Tenants may use WAN accelerators for Backup Copy jobs targeted at the cloud repository.

WAN accelerators deployed in the cloud run the same services and perform the same role as WAN accelerators in an on-premises backup infrastructure. When configuring Veeam Backup Copy jobs, tenants can choose to exchange data over a direct channel or communicate with the cloud repository via a pair of WAN accelerators. To pass VM data via WAN accelerators, the IT administrators and tenants must configure WAN accelerators in the following way:

- The source WAN accelerator is configured on tenant side
- The target WAN accelerator is configured on the cloud side

## **Tenant Veeam Backup & Replication server**

To connect to the cloud and use the cloud repository service running with Microsoft Azure, tenants utilize Veeam Backup & Replication servers deployed within their environment.

Once connected, tenants configure necessary jobs and perform data protection and disaster recovery tasks targeted at the cloud repository. All tasks are performed by the tenants themselves. Cloud administrators only sets up the Veeam Cloud Connect infrastructure and exposes storage resources to tenants with the use of the cloud repository.

## Azure preparations

Before going into details about deploying the Veeam solution from Azure Marketplace, there are a few things you need to know about the scope of this white paper.

This white paper won't go too deep into the configuration of Microsoft Azure. We will start by deploying a single VM from a Veeam template in Azure Marketplace, and attach a repository disk for storing backup copies. The Veeam Backup & Replication server running in Microsoft Azure won't be joined to any domain and will be a single, multi-purpose server. The distributed, scale-out deployment with multiple VMs will be covered in the second part of this white paper.

It's important to note that Azure offers two different deployment models: the new Azure Resource Manager model (ARM) and the classic model (ASM). This white paper only focuses on deployment in ARM, as it represents the future of the platform.

To begin:

- Sign up for an Azure subscription
- Deploy the Veeam Cloud Connect *for the Enterprise* offering from Azure Marketplace
- Create and assign a data disk to the VM
- Add additional Veeam Backup & Replication proxy server(s), cloud gateway(s) and repository server(s) as necessary

## Sign up for an Azure subscription

Before you can start, you need a Microsoft Azure subscription. If you do not already have an Azure subscription, you can request a 30-day trial through <https://azure.microsoft.com/en-us/free/>.

The sign-up process requires a mobile phone number (to receive a verification code through SMS), a credit card (you won't get billed during the trial, but it is required for proof of identity) and a Microsoft account username (formerly Windows Live ID). Once you've signed up, you can start your deployment.

Please note that at the time of writing this white paper, trial subscriptions were given a \$200 (or 170 euros) credit in Azure, but that may not be the case going forward. If you hit that spending limit or pass the 30-day trial period, your account will be suspended. However, you have the option to upgrade the trial to a Pay-As-You-Go Azure subscription.

## Billing

One of the items you should keep an eye on is your usage. You can easily see what you are using by looking at the billing page of your subscription. Simply go to <http://azure.microsoft.com/en-us/account/>.

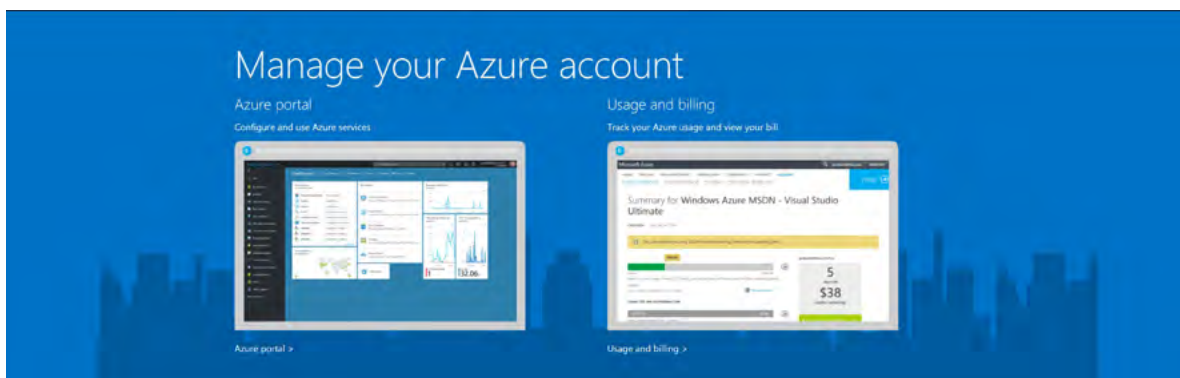


Figure 1: Manage your Azure account web page

From this page, you can go to the **Management Portal** (we will do that later) and to the **Usage and billing** portal. Select **Usage and billing** to see your usage and billing history.

After signing in with your Microsoft account, you will be able to see your subscription(s) and look at the usage.

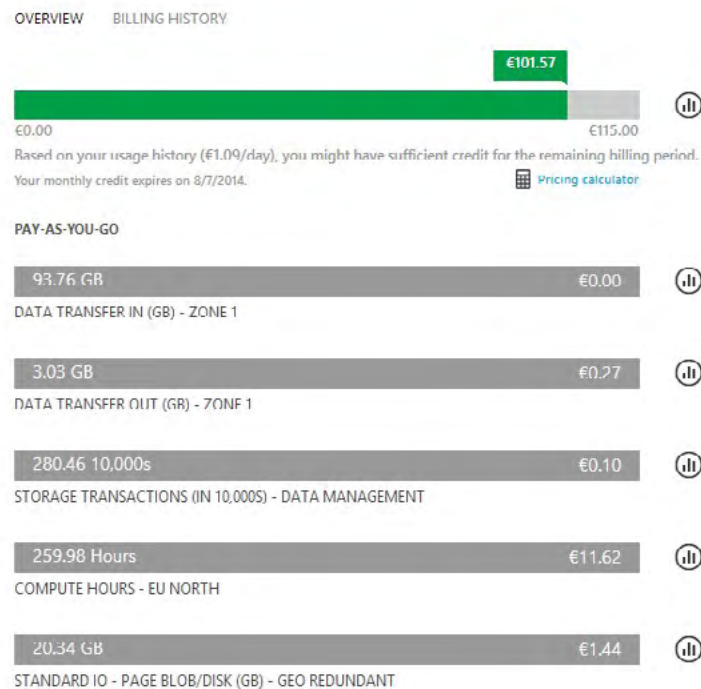


Figure 2: Usage history

The **Usage and billing** page shows a lot of data in addition to your current usage. You can view your historical billing and even download a CSV with lots of valuable data to view what you are exactly consuming.

## Log in to the Management Portal

Once you have a subscription, you can go to the **Management Portal** and start deploying some resources. You can go to <https://portal.azure.com>.

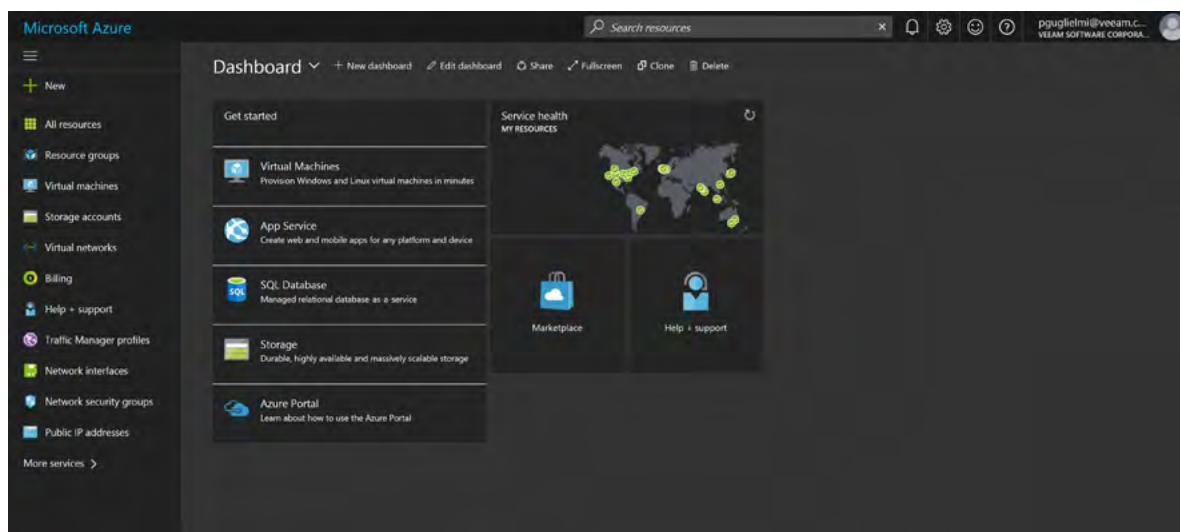


Figure 3: Azure Management Portal

Now that you have signed in, you are ready to create some resources. It's important to know that the person who created the subscription is the only one who is able to access it. If you want more administrators to have access, you will need to add them to your subscription.

More information on co-administrators, roles and more can be found here: <https://docs.microsoft.com/en-us/azure/billing-add-change-azure-subscription-administrator>

## Create the Veeam Cloud Connect for the Enterprise VM

The Veeam Cloud Connect for the Enterprise is in Azure Marketplace. Simply go to Azure Marketplace to start deploying a preconfigured VM.

Go to Azure Marketplace in the portal:

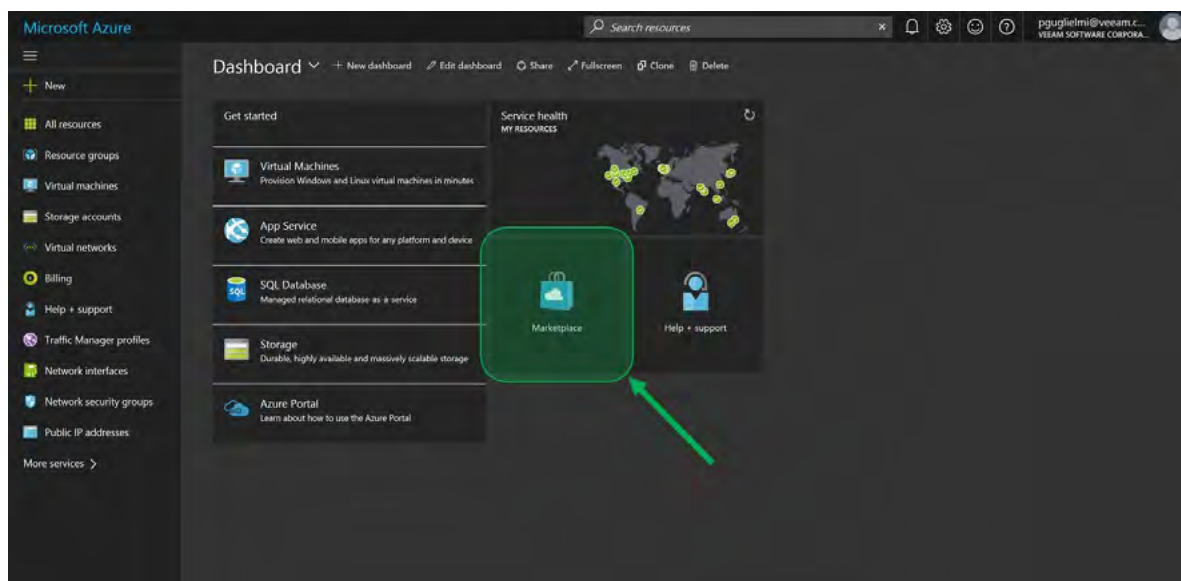


Figure 4: Azure Marketplace

In Azure Marketplace, search and find the Veeam Cloud Connect server.

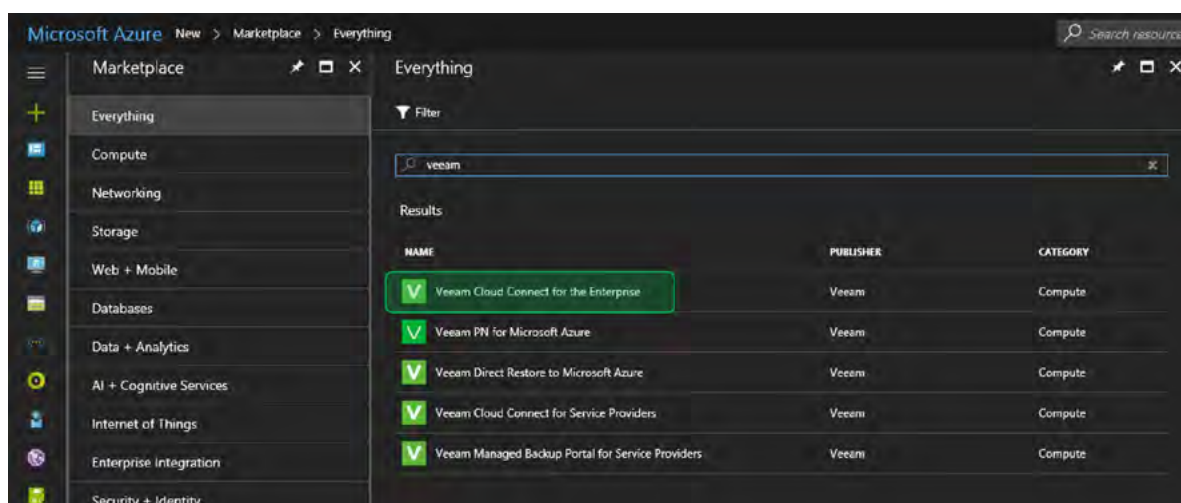


Figure 5: Search Azure Marketplace



Click on the Veeam Cloud Connect *for the Enterprise* offering and review the information. Then press **Create**.

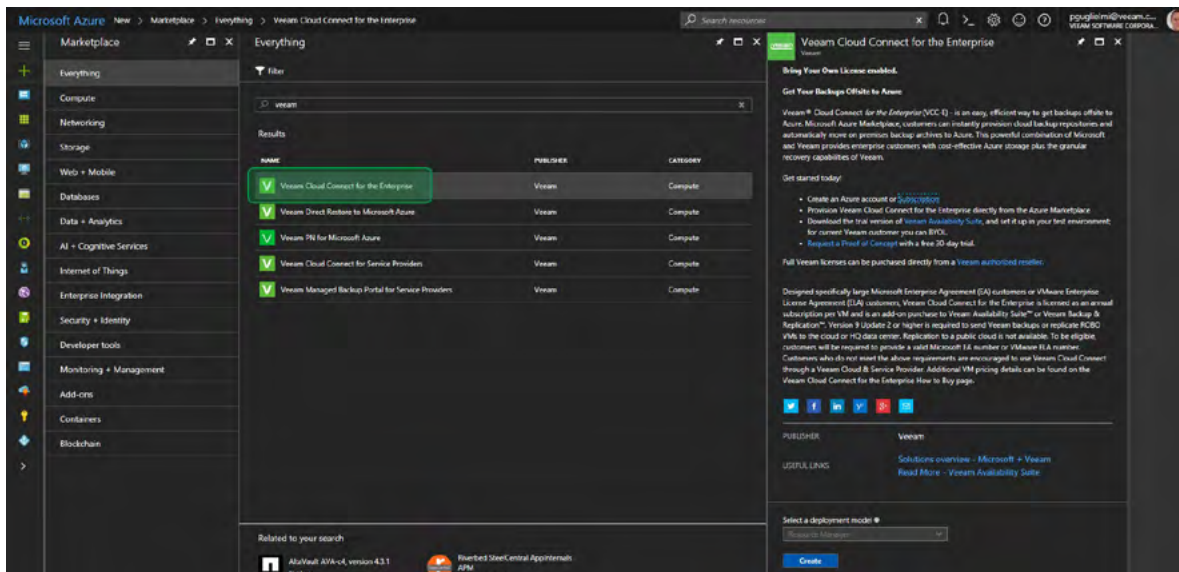


Figure 6: Select the Veeam Cloud Connect *for the Enterprise* offering

Next, fill in the specific settings for this machine. There are multiple options that you will need to decide on.

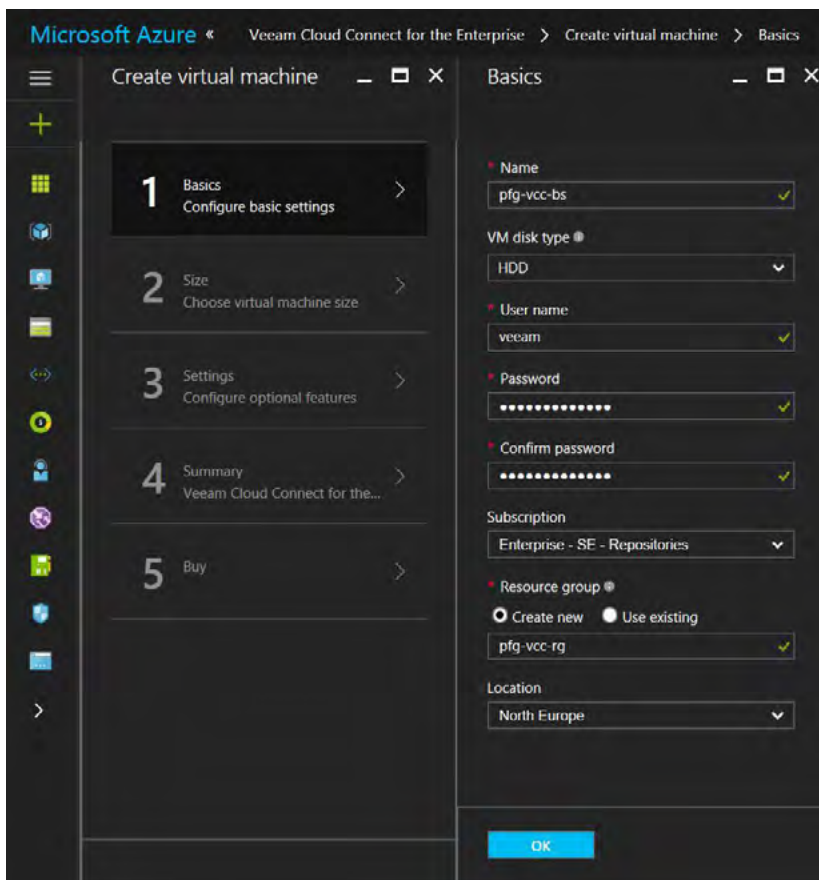


Figure 7: VM details



As you can see in Figure 7, we have already filled in the name of the server (**pfg-vcc-bs**), a username and a password for that VM. You will also need to decide on options such as the pricing tier, networking, storage, diagnostics, resource group and so on. Although Veeam's solution is very scalable and many of the options can be modified afterward, it is best to think this through before you start. One of the documents we recommend to review is the [reference architecture guide](#) for Veeam Cloud Connect. Although this is written for service providers who want to deploy their own infrastructure, many guidelines are also valid for Veeam Cloud Connect on Azure. Let's have a look at a few settings.

First of all, for the purpose of this white paper, we chose to create a new Azure resource group during the creation of the first VM, and North Europe for the location. All other VMs that will be deployed afterward will also be in the same resource group. An Azure resource group is typically used to logically group related resources that share the same lifecycle. For example, in the case of Veeam Cloud Connect, if you are deploying multiple servers that have different roles, you will be able to manage them together. We advise you to create a specific resource group for this service to make it easier when you scale out.

More information on Azure resource groups can be found here: <https://docs.microsoft.com/en-us/azure/virtual-machines/windows/infrastructure-resource-groups-guidelines>

A list of Azure regions can be found here: <https://azure.microsoft.com/en-us/regions/>

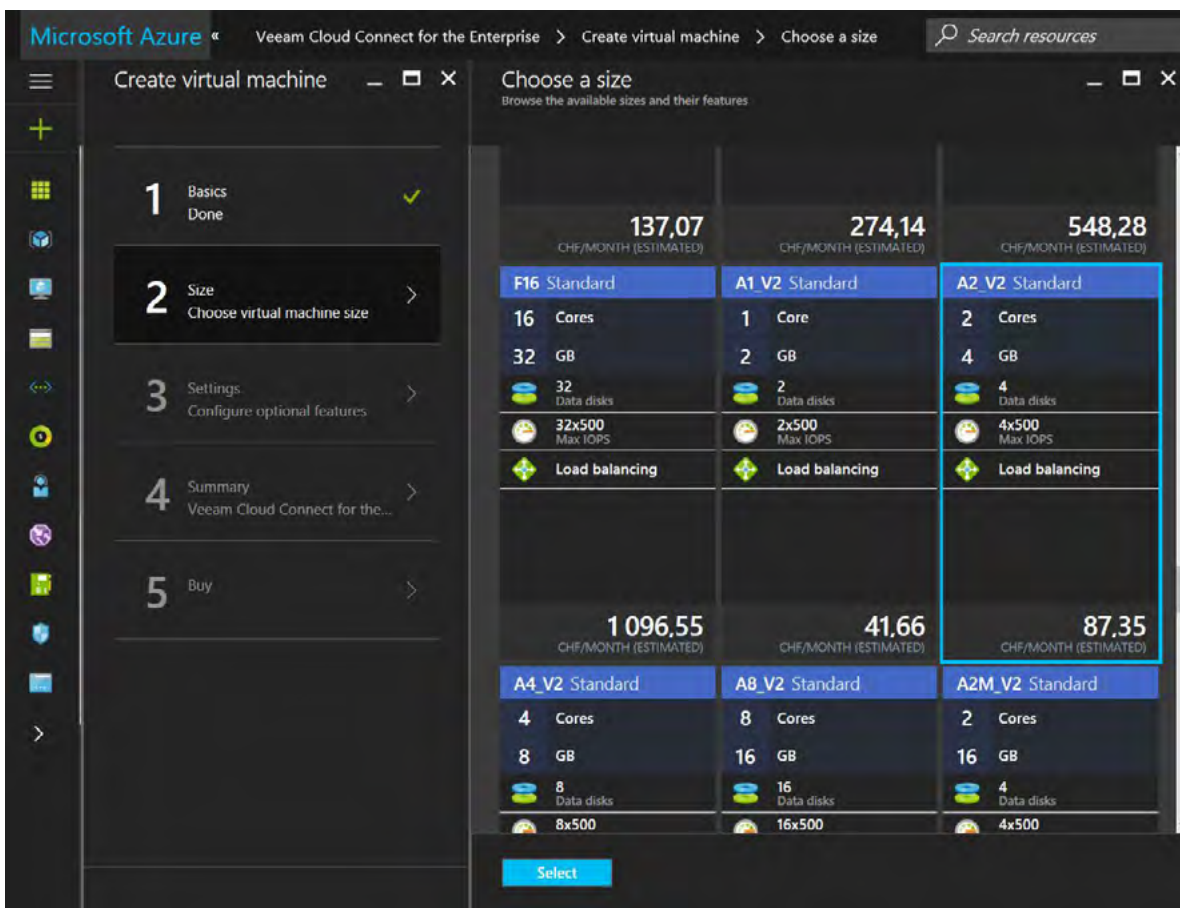


Figure 8: Pricing tiers

The pricing tier displays many different options – each for a different VM size. For example, an A0 VM can only hold one data disk and has only 25% of a CPU core and 0.75 GB of RAM. As you can imagine, this doesn't scale out very well and cannot contain a lot of data. An A2\_V2 has two cores and 4 GB and can hold up to four data disks. The A\_V2 series are particularly interesting as they have a little more RAM, faster disks, run on more modern CPUs and are cheaper than the Standard A series. More information about the different tiers and pricing can be found here: <https://azure.microsoft.com/en-us/pricing/details/virtual-machines/windows/>

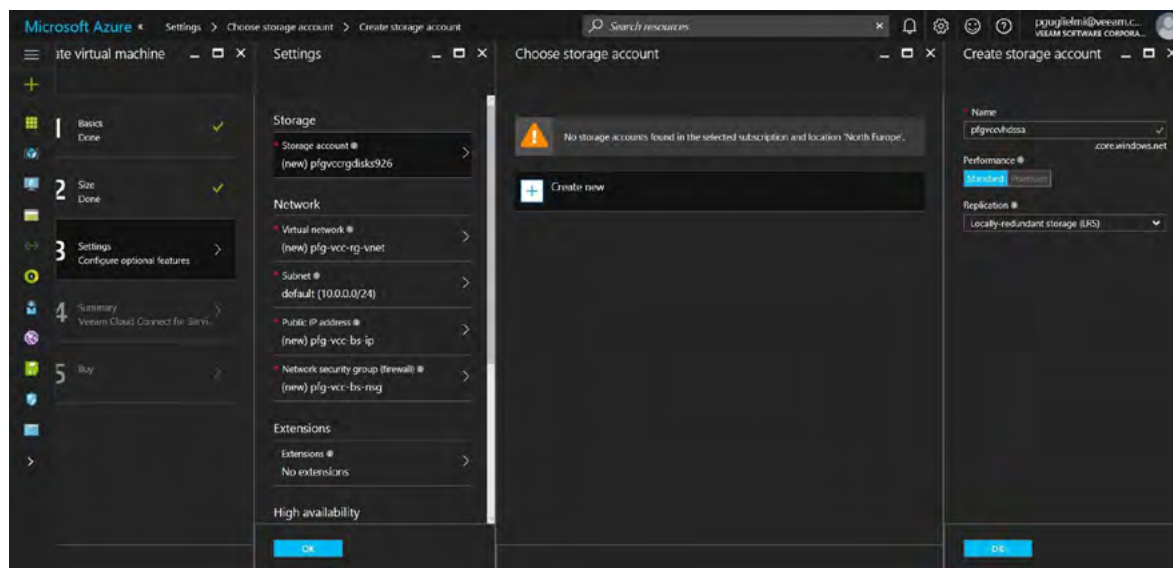


Figure 9: Optional configuration – Storage account

You can review and change quite a few settings on the optional configuration page. We'll first create a new standard storage account to store the Azure Virtual Machines' OS disks.

You can learn more about storage accounts performance and replication options here: <https://docs.microsoft.com/en-us/azure/storage/storage-introduction>




Figure 10: Optional configuration – Managed disks

Note: In February 2017, Microsoft added a new setting for disk management called **Managed Disks**. Although it brings some interesting new features, Managed Disks aren't required for the purpose of this white paper. To learn more about Azure Managed Disks, visit: <https://docs.microsoft.com/en-us/azure/storage/storage-managed-disks-overview>

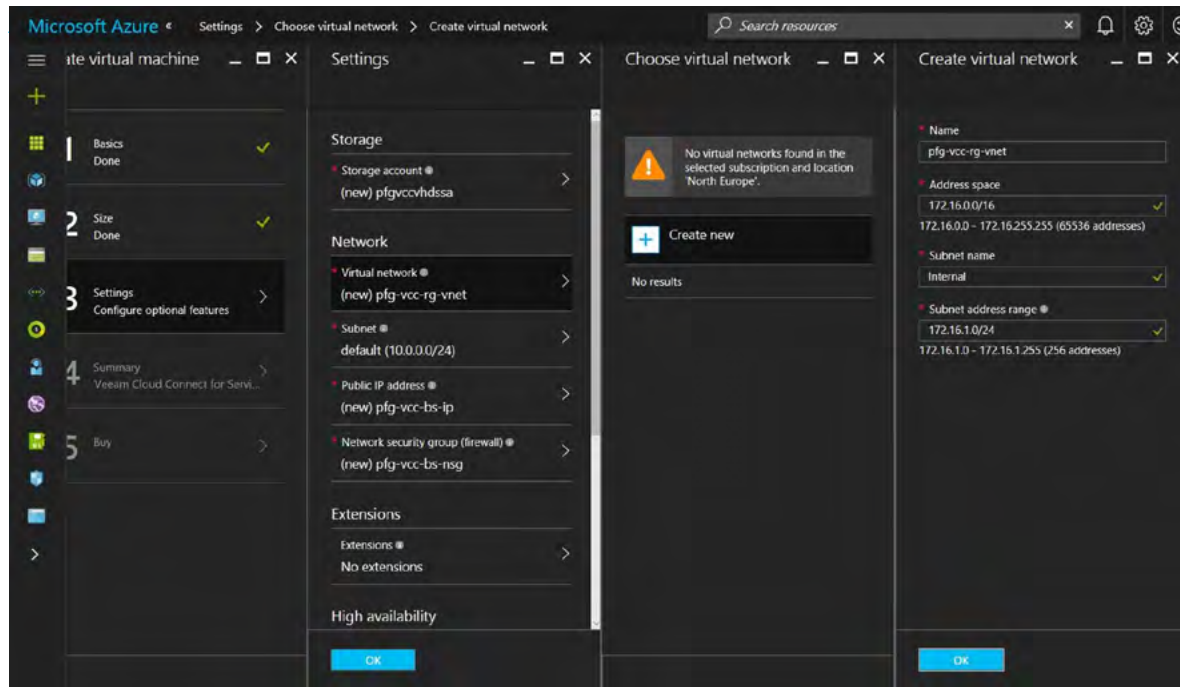


Figure 11: Optional configuration – Virtual Network

For the purposes of this white paper, we need to create a new virtual network with a new subnet inside. For now, we only create a new subnet named "Internal" inside the new virtual network. Another one will be created later for the needs of the distributed infrastructure.

**Note:** Doing so right from the start will allow us to place additional VMs in the right network perimeters to make sure that only the required Veeam components can be accessed from the internet when we'll describe the distributed architecture later in this white paper.

There are two types of IP addresses that can be allocated to a virtual machine in Azure: public and private. Public ones are used for communication with the internet, whereas private ones are used for communication within a virtual network. IP address allocation for both types can be either dynamic or static. For private IP addresses, allocation is always done through DHCP, meaning that if you set a static IP address to an Azure Virtual Machine, the selected IP address corresponds to a DHCP reservation. This will be necessary when you scale-out the Veeam Cloud Connect solution and the different components need to talk to each other.

**Note:** Because of the way it works in Azure, the configuration of network interfaces within the guest operating systems has to be dynamic. If you need static private IP addresses in your Azure Virtual Machines, you have to set them through the Azure portal.

However, the problem with dynamic public IP addresses is that they can change over time (e.g., when you stop the VMs in your service or de-allocate them). To solve this issue, you can work with a static public IP address that won't change over time unless you decide to disassociate it. Another way to work around that is to work with a proper DNS name configuration. This will be covered later in this white paper.

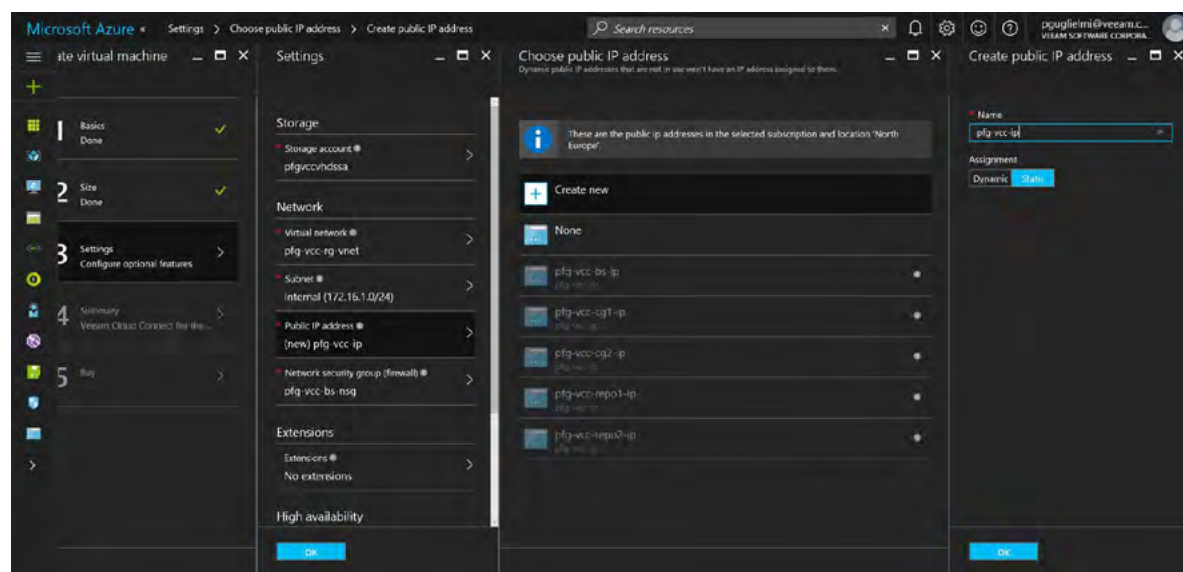


Figure 12: Public IP Address configuration

For more information about IP addresses in Azure: <https://docs.microsoft.com/en-us/azure/virtual-network/virtual-network-ip-addresses-overview-arm>

IP address pricing: <https://azure.microsoft.com/en-us/pricing/details/ip-addresses/>

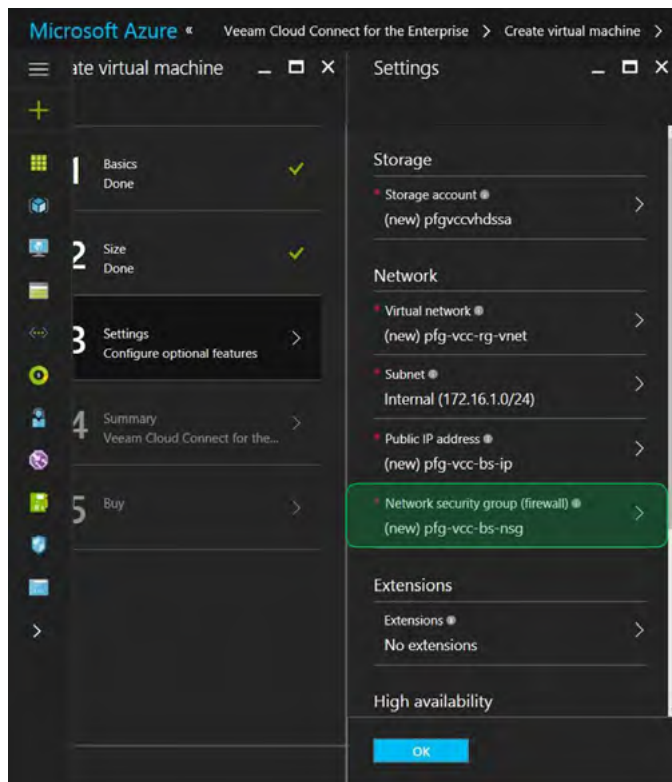


Figure 13: Network security group selection

As shown in Figure 13, we accept the default for the network security group selection and configuration. Because the virtual machine is deployed from a preconfigured template in Azure Marketplace, the default settings will be as follows:

- **Inbound rules:** Remote Desktop Protocol is allowed to enable remote connection to the virtual machine, therefore TCP port 3389 is open. Veeam Cloud Connect creates and manages a secure SSL tunnel in which all network communications are encapsulated which uses TCP port 6180. However, if you plan to use a different port for Veeam Cloud Connect, you'll need to update the inbound rules accordingly.
- **Outbound rules:** None.

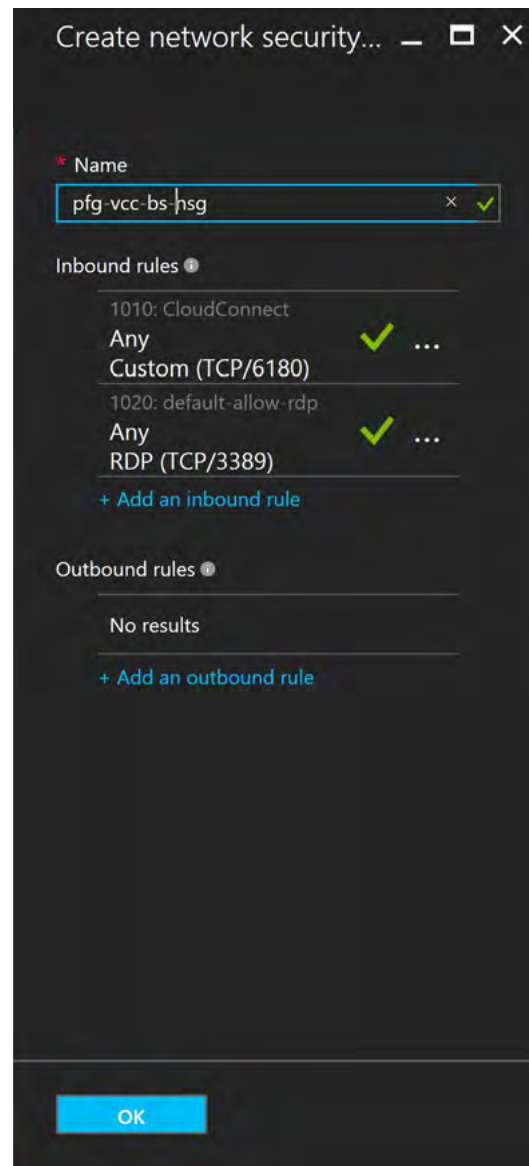


Figure 14: Network security group default rules

To finish the creation process, we choose to not enable any extension, Availability set or monitoring. Availability sets will be covered in the second part of this white paper. If you choose to enable monitoring, we recommend you create a new storage account dedicated to it in order to make sure that it won't disturb VHDs I/Os and have a negative impact on virtual machines performance, as monitoring generates logs and a lot of IOPS. Note that a single storage account has a maximum of 20,000 IOPS.



To learn more about Azure storage limits and scalability, visit: <https://docs.microsoft.com/en-us/azure/storage/storage-scalability-targets>

Finally, when everything is configured, you can review the summary (Figure 15) and click **OK**.

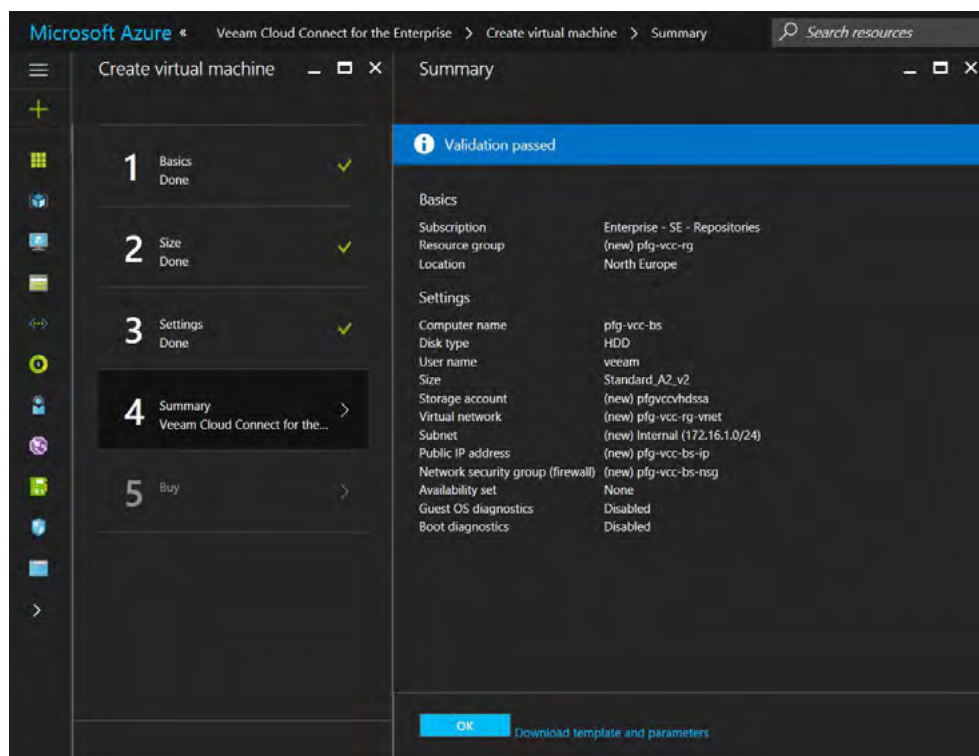


Figure 15: Create virtual machine summary

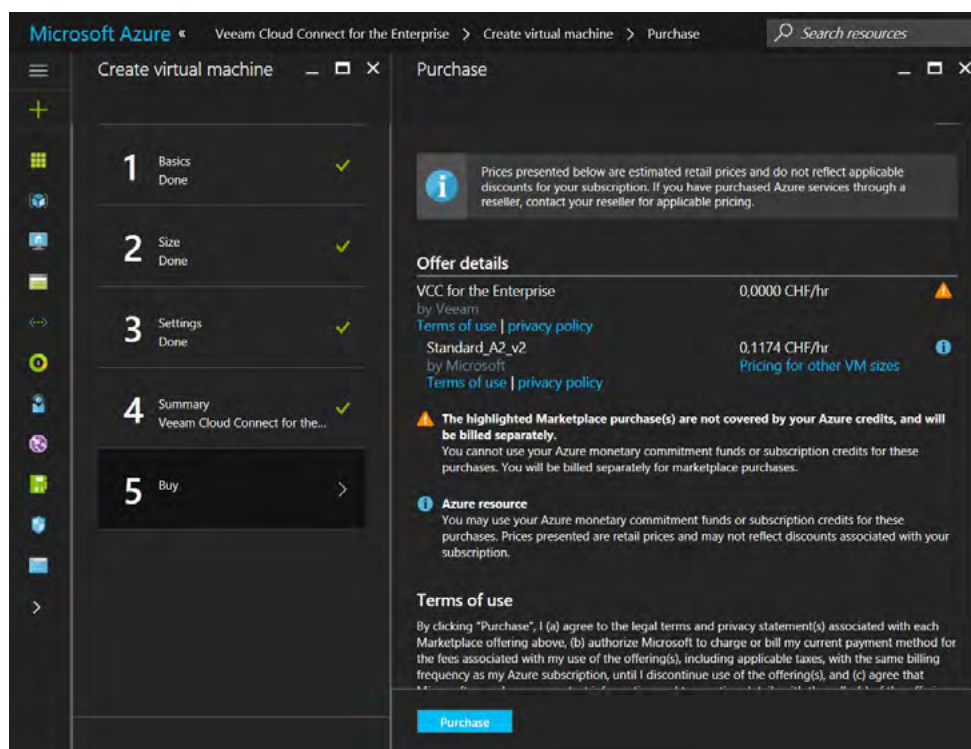


Figure 16: Offer details

When you receive a final confirmation with the details and legal terms, press **Purchase**.




Figure 17: Deployment in progress

When you are finished, Microsoft Azure will create and deploy the VM, and after a few minutes, the job will be finished and your VM will be ready to use. Now it is time to go to **Virtual Machines** and click on the newly deployed virtual machine to get an overview of it:

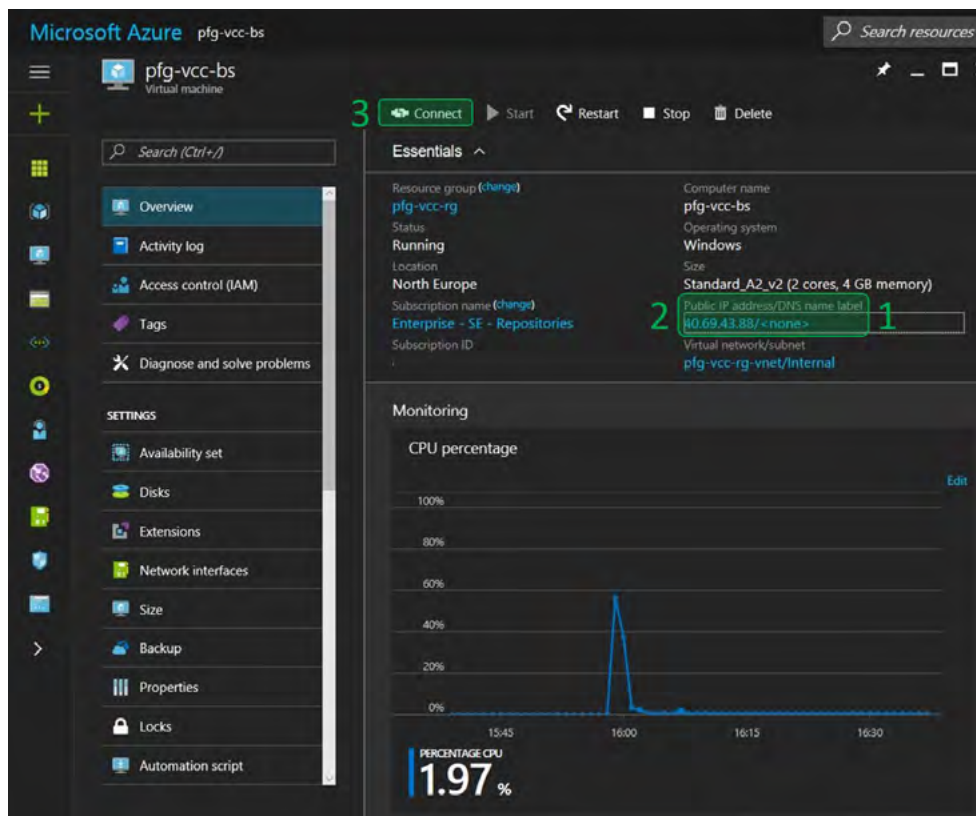


Figure 18: VM Ready to use

Please note some important items shown in Figure 18:

1. **DNS name label:** By default, no DNS name is configured. If you don't want or can't set a static public IP address, then it's better to set a DNS name label for your virtual machine. In this case, we are using a DNS name under the *.northeurope.cloudapp.azure.com* domain — the default MSFT domain for North Europe region. Yours may vary depending on the Azure region you've selected for deploying your VM, but you can use your own domain, which we recommend when you want to get this set up into production for customers. To use a custom DNS name for your Azure VMs, you can simply create a new CNAME (alias) record at your domain registrar and map it to the one in *your\_region.cloudapp.azure.com*.



2. **Public IP address:** This is your public IP address to which your subsidiaries and departments (and you, with RDP, PowerShell remoting) will connect. You receive this public IP when the VM is powered on. If the VM is stopped from the Azure portal, its status will turn into "Stopped (deallocated)" which means that you won't be charged for this VM anymore, but its public IP address will be released, and will most probably change at next power on. Which is why it is important to use either a static IP address or a DNS name.
3. **Connect:** This button allows you to RDP into your VM. This will automatically download an RDP connection with the correct settings. The only thing left to do is add your username and password to connect to the server.

## Configure a DNS name

Within the **Overview** blade of your virtual machine, click on the **Public IP Address/DNS name label** values to view their configuration.

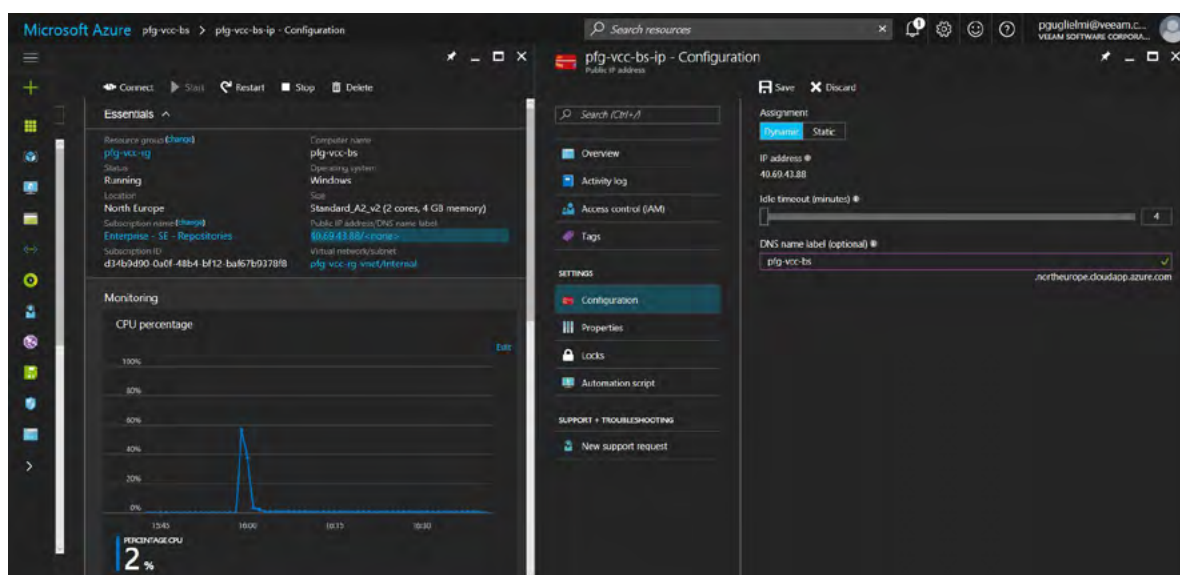


Figure 19: Public IP address and DNS name label configuration

The DNS name label is optional but must be unique within your chosen region. Once the configuration is OK, you can click on the **Save** button. Check the **Overview** again to confirm everything is configured as expected.

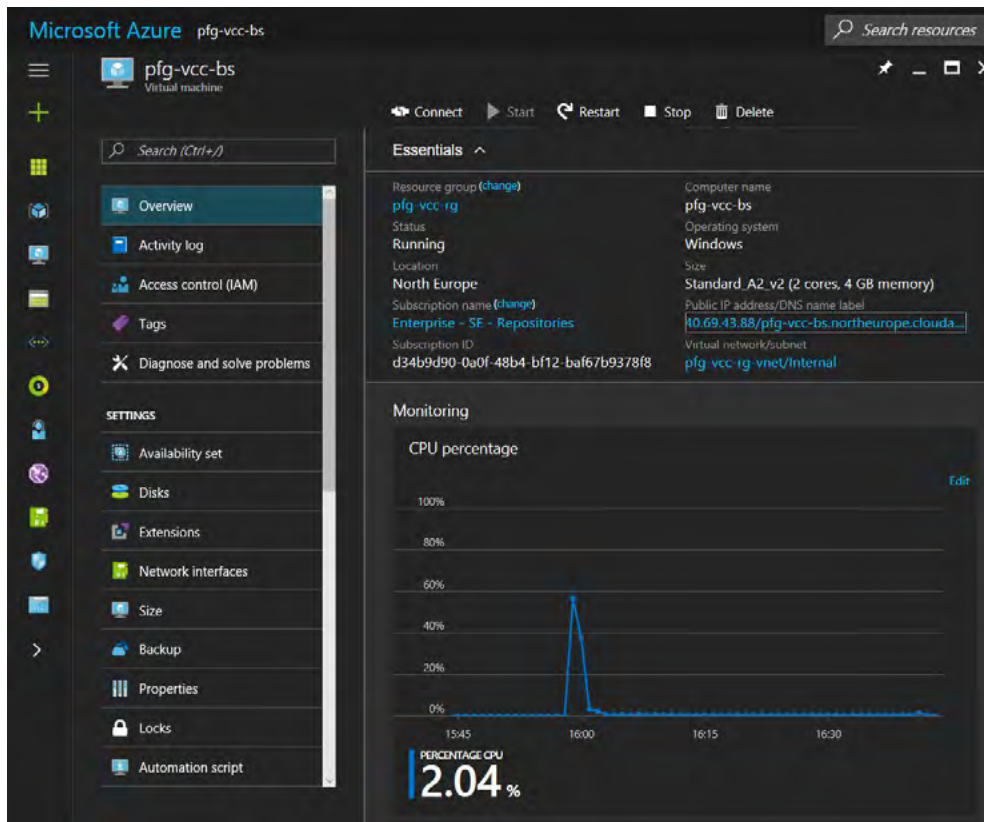


Figure 20: DNS name configured correctly

## Create and assign a repository disk

As explained in the [Architecture and scenario](#) section, you also need a repository to store your tenants' backups. Since you don't want to store them on the C:\ drive, you will attach a new disk to your VM. This disk will hold your backup files. There are a few possibilities to create and attach an empty disk.

There are two methods to do this:

- Attach an existing disk
- Attach a new disk

More information can be found here: <https://docs.microsoft.com/en-us/azure/virtual-machines/windows/attach-disk-portal>

In this case, we will attach a new disk to the VM. To do so, select your virtual machine in the Azure portal, go to **Disks** and click on the blue **+ Add data disk** button.

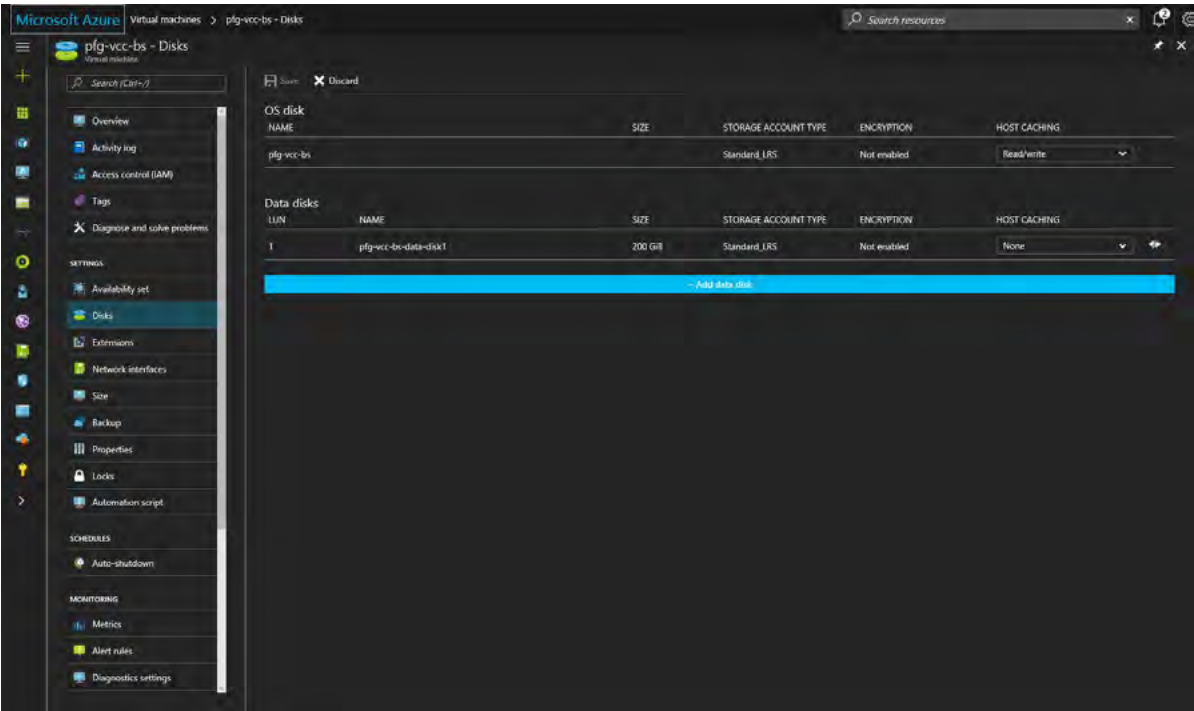


Figure 21: Virtual machine disks

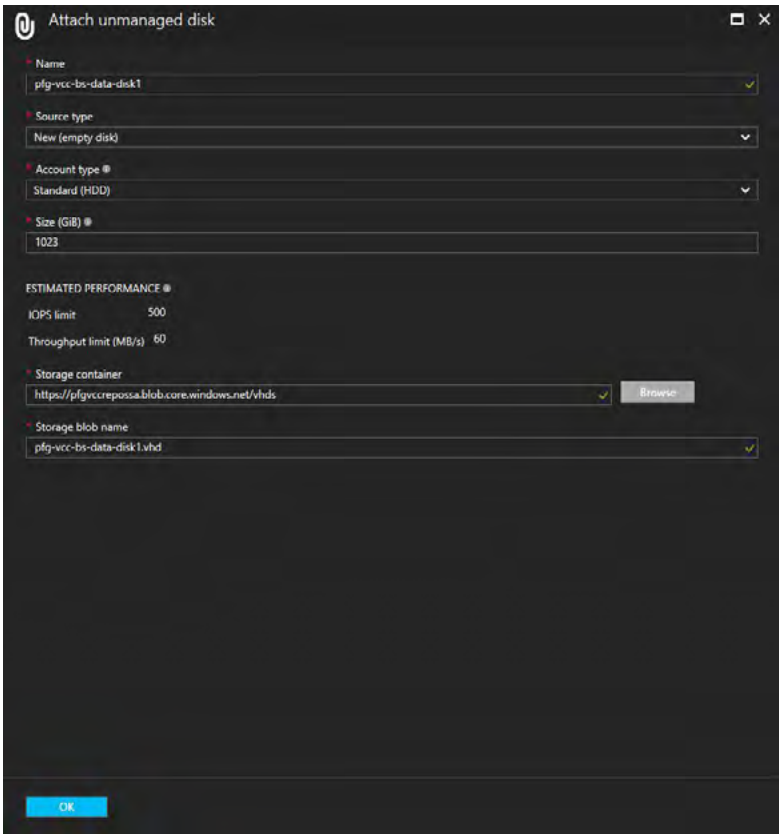
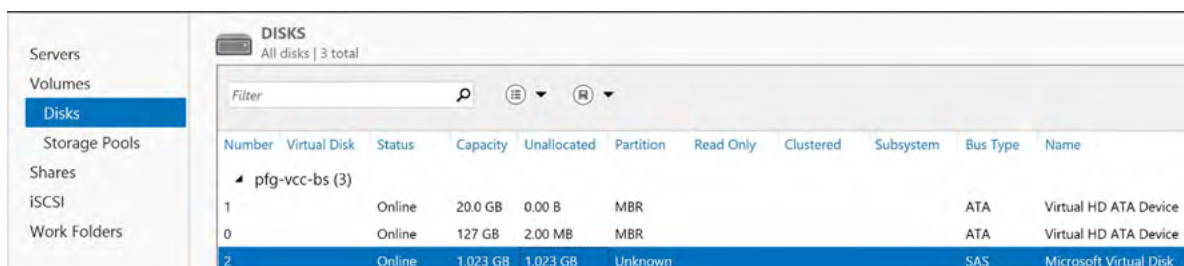


Figure 22: Attach a new data disk

Please note that I have chosen **None** (Figure 21) as the host caching preference.

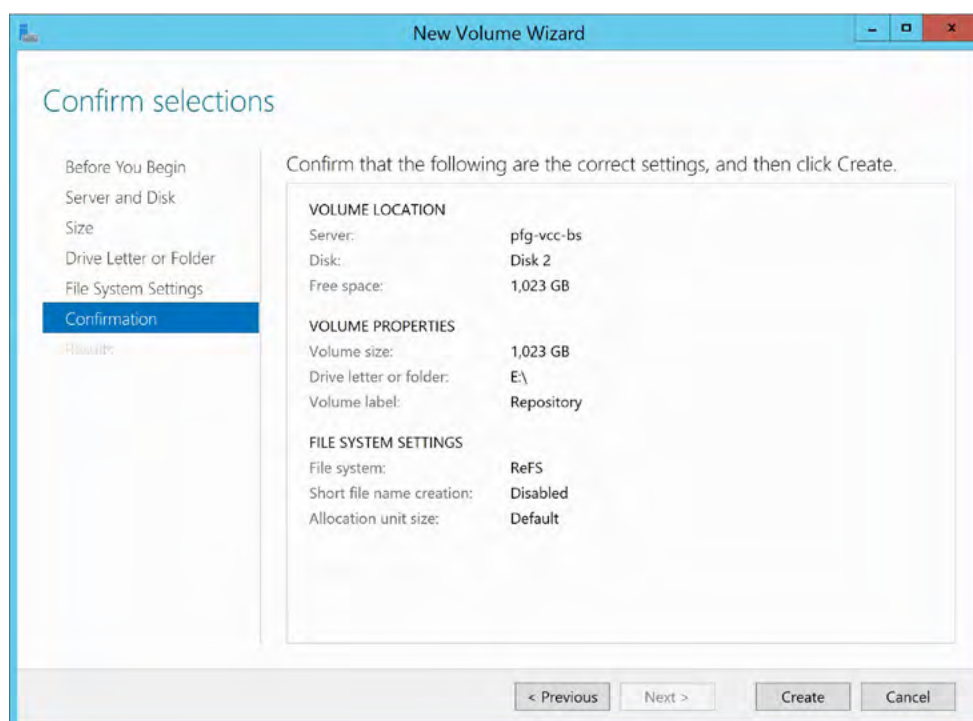
Finally, after the disk is attached to the VM, you will need to connect through RDP and initialize the disk within the guest operating system (see Figure 23 and Figure 24). Don't forget to click on the **Save** button once your new data disks have been added (see Figure 21).



The screenshot shows the 'DISKS' section in Veeam, listing three disks attached to the VM 'pfg-vcc-bs'. The table below represents the data shown in the interface.

Number	Virtual Disk	Status	Capacity	Unallocated	Partition	Read Only	Clustered	Subsystem	Bus Type	Name
1		Online	20.0 GB	0.00 B	MBR				ATA	Virtual HD ATA Device
0		Online	127 GB	2.00 MB	MBR				ATA	Virtual HD ATA Device
2		Online	1,023 GB	1,023 GB	Unknown				SAS	Microsoft Virtual Disk

Figure 23: Configure the disk in your VM



The screenshot shows the 'New Volume Wizard' window at the 'Confirm selections' step. The wizard confirms the following settings for creating a new volume:

- VOLUME LOCATION**
  - Server: pfg-vcc-bs
  - Disk: Disk 2
  - Free space: 1,023 GB
- VOLUME PROPERTIES**
  - Volume size: 1,023 GB
  - Drive letter or folder: E:\
  - Volume label: Repository
- FILE SYSTEM SETTINGS**
  - File system: ReFS
  - Short file name creation: Disabled
  - Allocation unit size: Default

At the bottom of the wizard, there are buttons for '< Previous', 'Next >', 'Create', and 'Cancel'.

Figure 24: Create the volume

## Configure Veeam Backup & Replication

The first part is ready. You have your Azure subscription, you've deployed the Veeam Cloud Connect *for the Enterprise* Azure Marketplace solution and added one or more data disks.

The next steps for the cloud side are:

- Initial configuration
- Repository configuration
- WAN acceleration configuration (optional): See [Appendix A: Using WAN acceleration](#)
- Veeam Cloud Connect configuration

### Initial configuration

Before telling your end users what they need to do, you need to do some small configuration steps.

First, you need to connect to your running VM through RDP. Note that you can find a connect option as described earlier. After you have done that, you can start Veeam Backup & Replication for the first time.

The first time Veeam Backup & Replication starts, it will request a specific license file for Veeam Cloud Connect that you will have received after purchasing new Veeam Cloud Connect *for the Enterprise* licenses. Add your license file and accept the EULA.

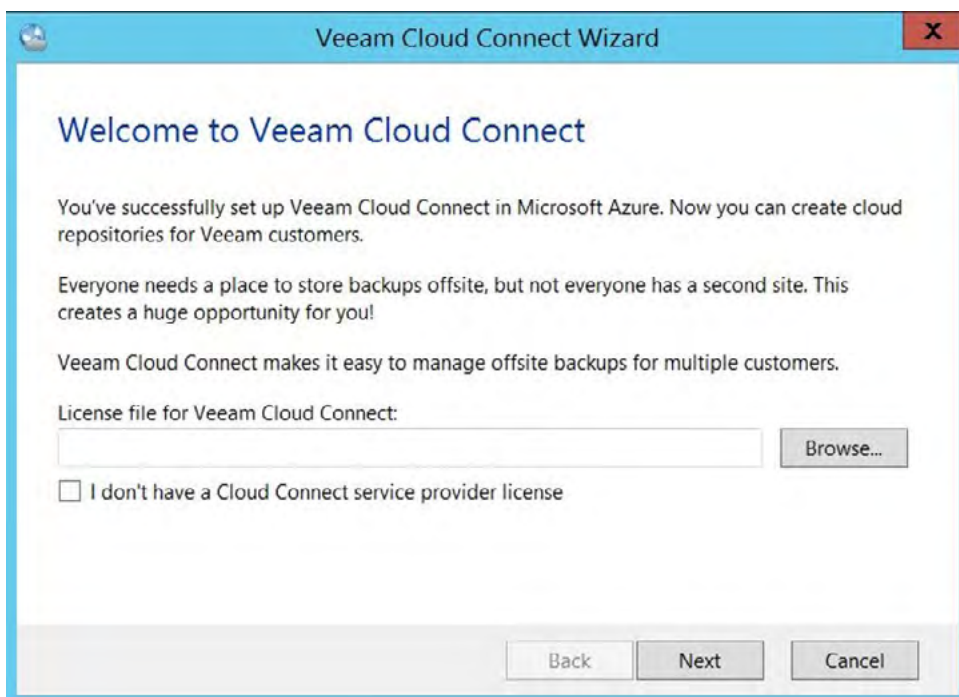


Figure 25: Browse and insert the Veeam Cloud Connect license file

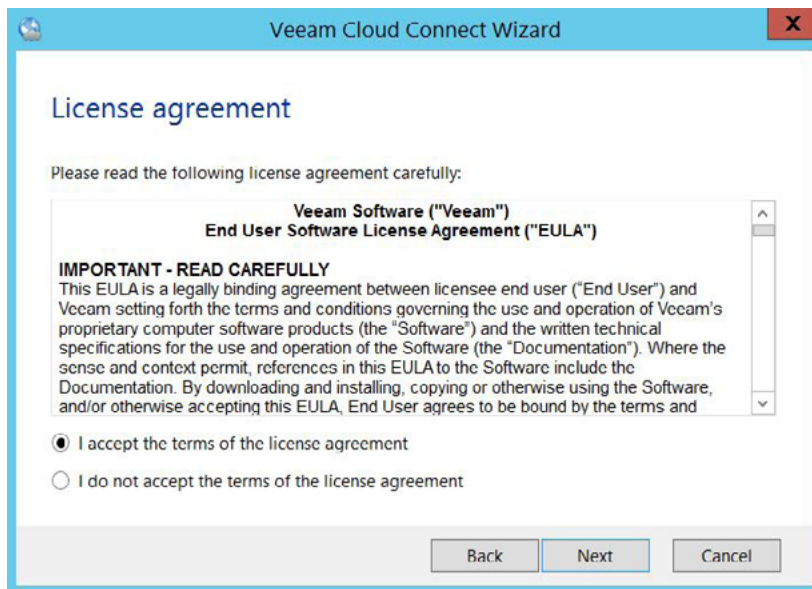


Figure 26: Accept EULA

After accepting EULA, you'll see a couple more screens reminding you to make sure TCP port 6180 is open for Veeam Cloud Connect secure communications, and that you've set a DNS name for your VM. Now you are ready for the final configuration steps described at the end of the wizard and to accept your first customers.

## Configure Veeam Backup & Replication repository

Before you configure the cloud gateway, SSL certificate, tenant and cloud repositories, you need a regular repository. By default, the template already has a backup repository but it is located on the same drive as the OS installation. It is better to create a repository on the additional disk that you attached earlier.

Go to **Backup Infrastructure > Backup Repositories** and click the **Add Repository** button in the ribbon.

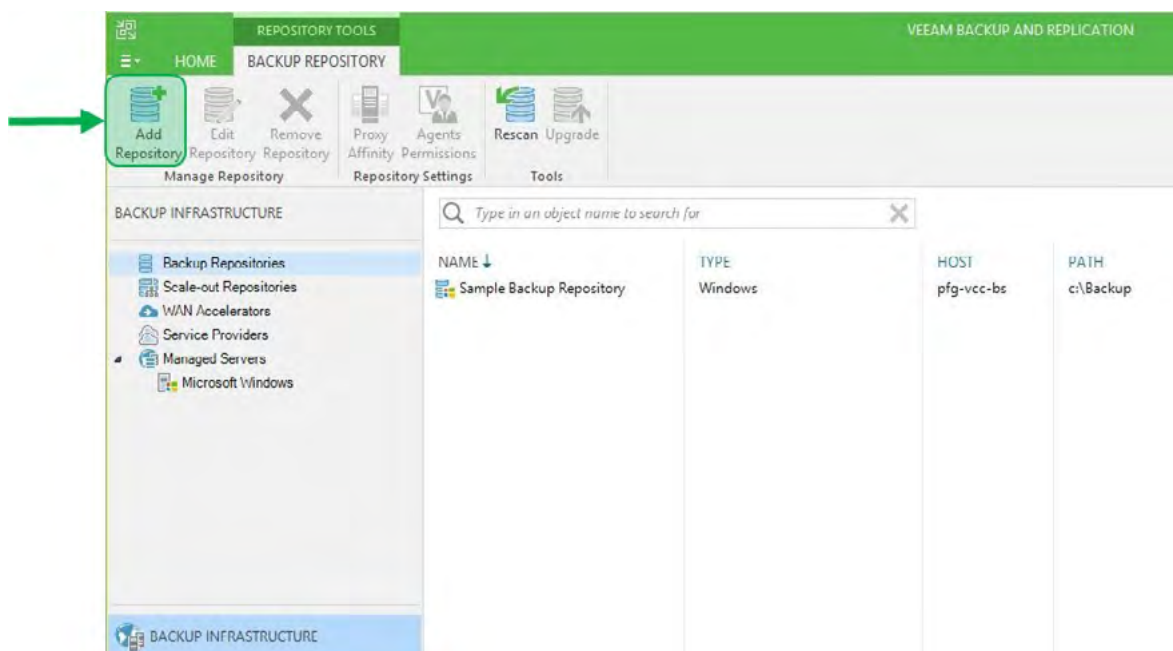
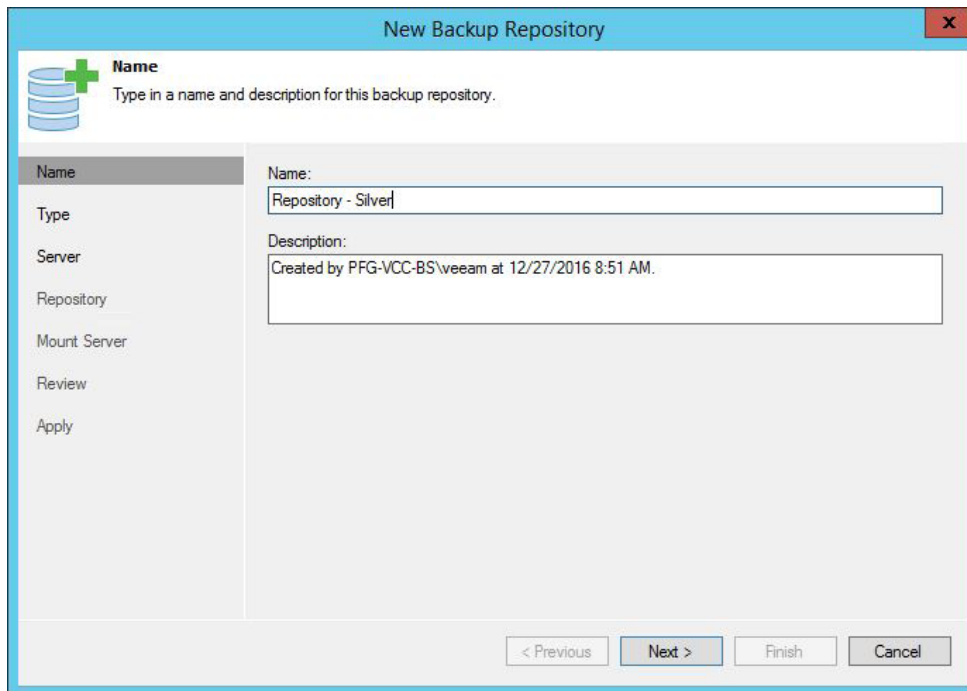


Figure 27: Add Repository



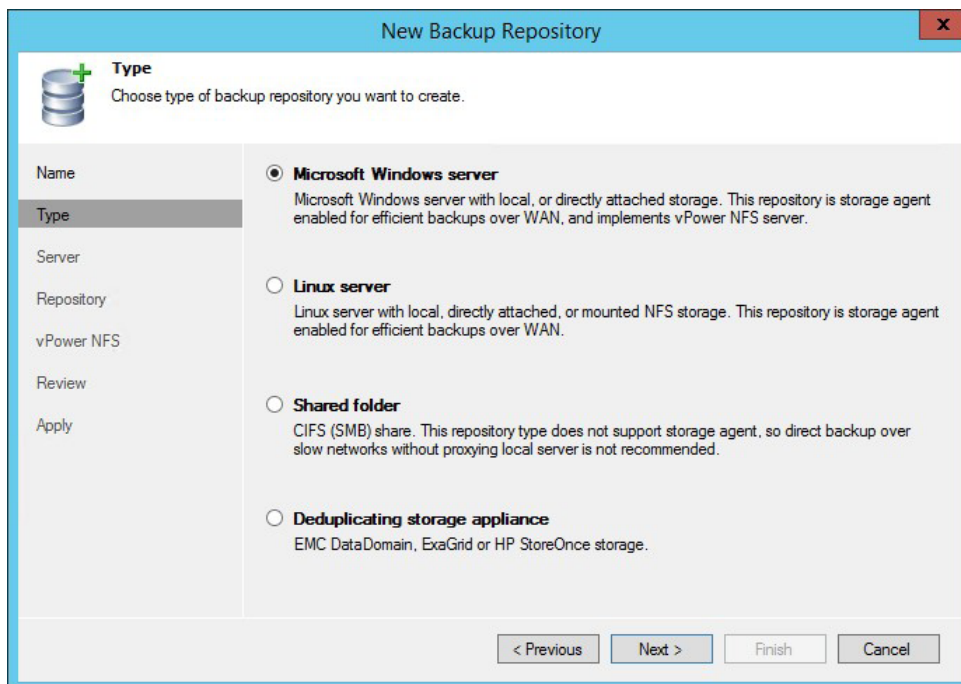
In the **New Backup Repository** wizard, type in the name and description for the repository. Note that this name won't be shown to the tenant. You will give each tenant a specific name that the tenant will see in its infrastructure. Press **Next**.



The screenshot shows the 'New Backup Repository' wizard with the 'Name' page selected. The left sidebar lists the steps: Name, Type, Server, Repository, Mount Server, Review, and Apply. The main area has a title 'Name' with a green plus icon and a database icon, followed by the instruction 'Type in a name and description for this backup repository.' Below this are two text input fields: 'Name:' with the value 'Repository - Silver' and 'Description:' with the value 'Created by PFG-VCC-BS\veeam at 12/27/2016 8:51 AM.' At the bottom are four buttons: '< Previous', 'Next >', 'Finish', and 'Cancel'.

Figure 28: Name your repository

On the **Type** page, select **Microsoft Windows server** and press **Next**.



The screenshot shows the 'New Backup Repository' wizard with the 'Type' page selected. The left sidebar lists the steps: Name, Type, Server, Repository, vPower NFS, Review, and Apply. The main area has a title 'Type' with a green plus icon and a database icon, followed by the instruction 'Choose type of backup repository you want to create.' Below this are four radio button options: 'Microsoft Windows server' (selected), 'Linux server', 'Shared folder', and 'Deduplicating storage appliance'. Each option has a brief description. At the bottom are four buttons: '< Previous', 'Next >', 'Finish', and 'Cancel'.

Figure 29: Choose Microsoft Windows server



On the **Server** page, choose the server that holds the storage and select the path to the data disk.

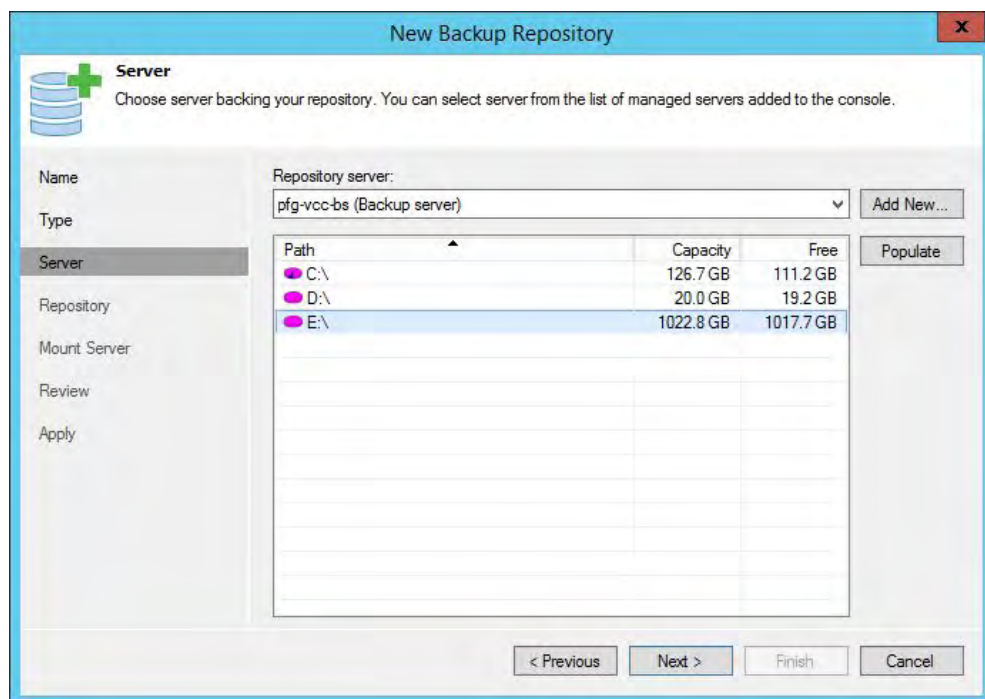


Figure 30: Choose the destination drive

On the **Repository** page, create or type in the path to the specific folder. A Windows Server volume can hold multiple different repositories, so it is always good to agree on a folder structure upfront.

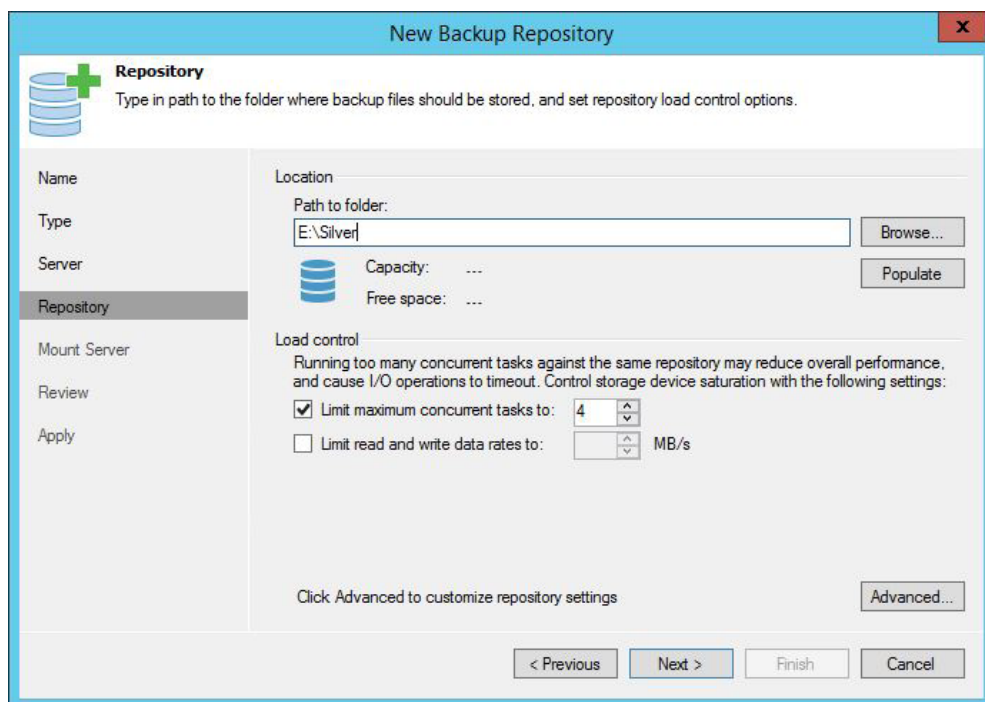


Figure 31: Choose the destination path and advanced settings

On the **Mount server** page, deselect **Enable vPower NFS server** as this cannot be used in Veeam Cloud Connect.

**New Backup Repository**

**Mount Server**

Specify a server to mount backups to for file-level restores. vPower NFS service allows for running virtual machines directly from backup files, enabling advanced functionality such as Instant VM Recovery, SureBackup and On-Demand Sandbox.

Name: Mount server: pfg-vcc-bs (Backup server)

Type: [Dropdown]

Server: ☐ Enable vPower NFS service on the mount server (recommended)

Repository: Specify vPower NFS write cache location on the mount server. Make sure the selected volume has enough free disk space available to store changed disk blocks of instantly recovered VMs.

Mount Server: Folder: C:\ProgramData\Veeam\Backup\NfsDatastore [Browse...]

Review: Click Ports to change NFS server and backup mount listener ports [Ports...]

Apply: [Ports...]

< Previous Next > Finish Cancel

Figure 32: Disable vPower NFS

Review your settings and create the backup repository.

**New Backup Repository**

**Review**

Please review the settings, and click Next to continue.

Name: Backup repository properties:

Type: Repository type: **Windows**

Server: Mount host: **pfg-vcc-bs**

Repository: Account: **Not set**

Mount Server: Backup folder: **E:\Silver**

Review: Write throughput: **Not limited**

Apply: Max parallel tasks: **4**

The following components will be processed on server pfg-vcc-bs:

Transport: **already exists**

Mount Server: **already exists**

☐ Import existing backups automatically

☐ Import guest file system index

< Previous Next > Finish Cancel

Figure 33: Review your settings

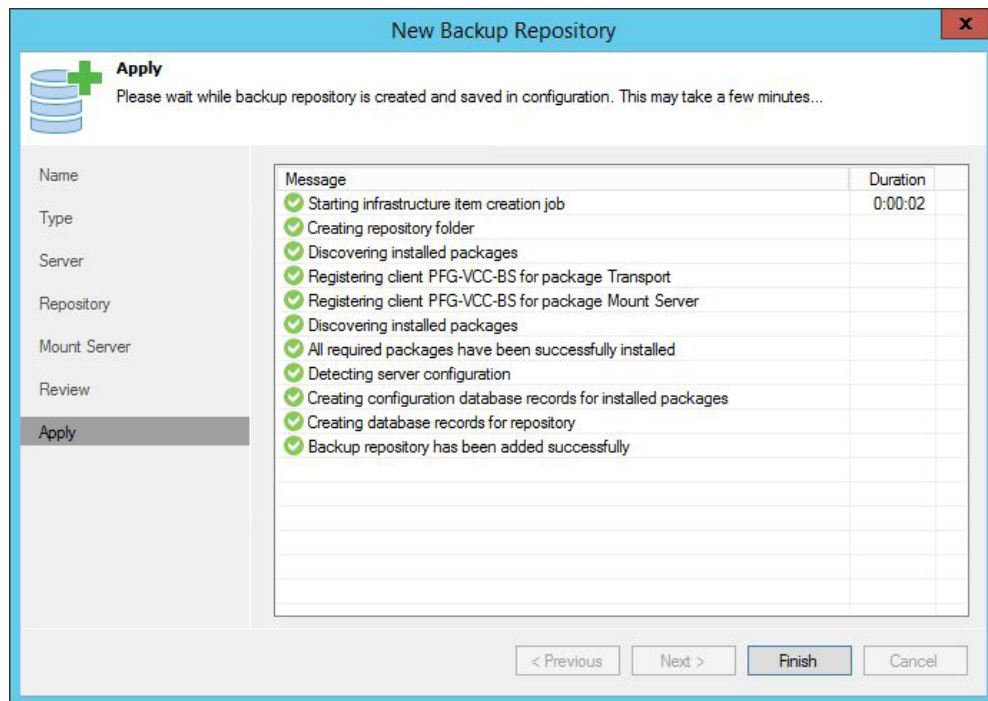


Figure 34: Successful creation of your repository

## Configure Veeam Cloud Connect

Theoretically, you can start creating a tenant now and skip the first two parts described below. However, we want to give you additional information as you might not want to work with a self-signed certificate in production or when you want to change options to the cloud gateway afterward.

### Manage a certificate

First, make sure that you have a certificate. Communication between components in the Veeam Cloud Connect infrastructure is carried out over an SSL connection secured with an SSL certificate. The SSL certificate is used for both authentication and tunnel encryption. It helps the cloud infrastructure and tenants identify themselves and ensures that parties taking part in data transfer are the ones they claim to be<sup>1</sup>.

There are two types of certificates that you can use:

- SSL certificate verified by a CA<sup>2</sup>: IT administrators can import this certificate through the UI and use that for verification purposes between the different components.
- Self-signed certificates: IT administrators can create a self-signed certificate with Veeam Backup & Replication. Veeam Backup & Replication uses the RSA full cryptographic service provider in Windows Server, but generating a self-signed certificate with any third-party solution is also an option.

Go to **Cloud Connect Infrastructure > Manage Certificates** to review, generate new, select or import a certificate.

<sup>1</sup> Note: SSL certificates are not being used for encrypting data stored on the cloud repository. If the tenant wants to encrypt data, he or she needs to enable encryption in Veeam Backup & Replication

<sup>2</sup> CA: Certificate Authority

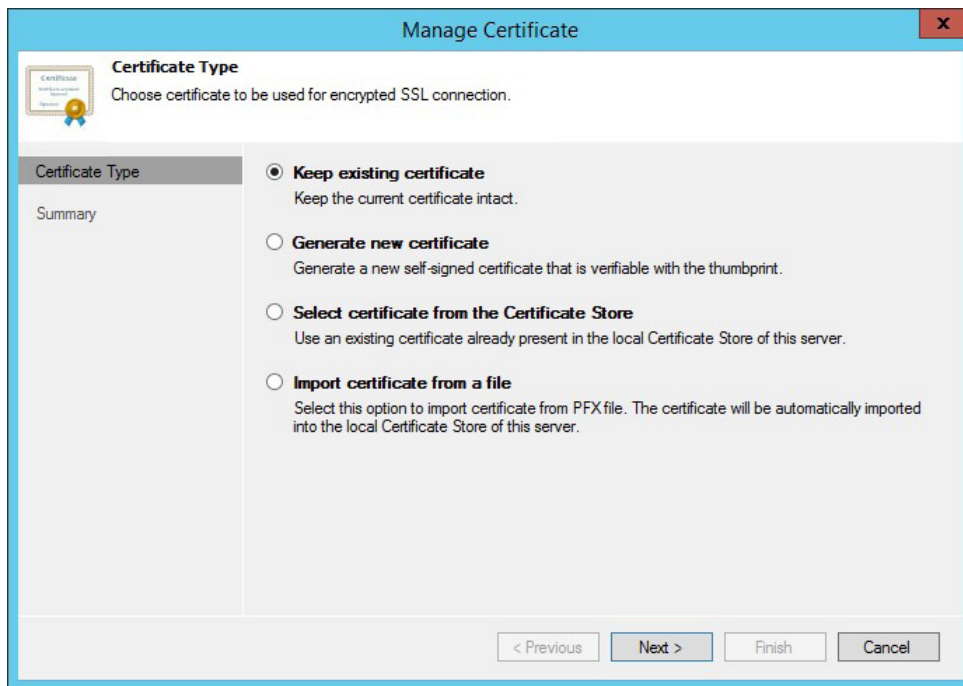


Figure 35: Certificate possibilities

## Cloud gateway

The cloud gateway is a service that resides on the cloud side and is the communication endpoint for tenants. It routes traffic and commands between Microsoft Azure, tenants and the cloud repository.

The template preconfigures the cloud gateway and you don't need to do anything. However, you might want to deploy additional cloud gateways later or change the settings when changes occur. For a description of the procedure to create or edit an existing cloud gateway, see [Deploy and configure additional cloud gateways](#).

## Create tenant

At this point, everything is ready for you to start serving your first tenant. Before you start, you should know the quota<sup>3</sup> in GB or TB that the tenant is allowed and whether the tenant's contract should expire at some date or not<sup>4</sup>. If you have configured the WAN acceleration, you need to know which WAN accelerator the tenant is allowed to use.

Go to **Cloud Connect** and press the **Add Tenant** button in the ribbon.

<sup>3</sup> Quota: The amount of storage assigned to one tenant on one cloud repository. It is a chunk of storage resources that the tenant can use for storing backups within the cloud repository. IT administrators can assign quotas on different cloud repositories to one tenant.

<sup>4</sup> Automatic expiration: This is a date at which the tenants access to their quotas within the cloud repository will expire. The lease settings help the IT administrators restrict for how long the tenant should be able to use cloud repository resources. However, it might not always be necessary, especially when both sides are managed by the same IT team internally.

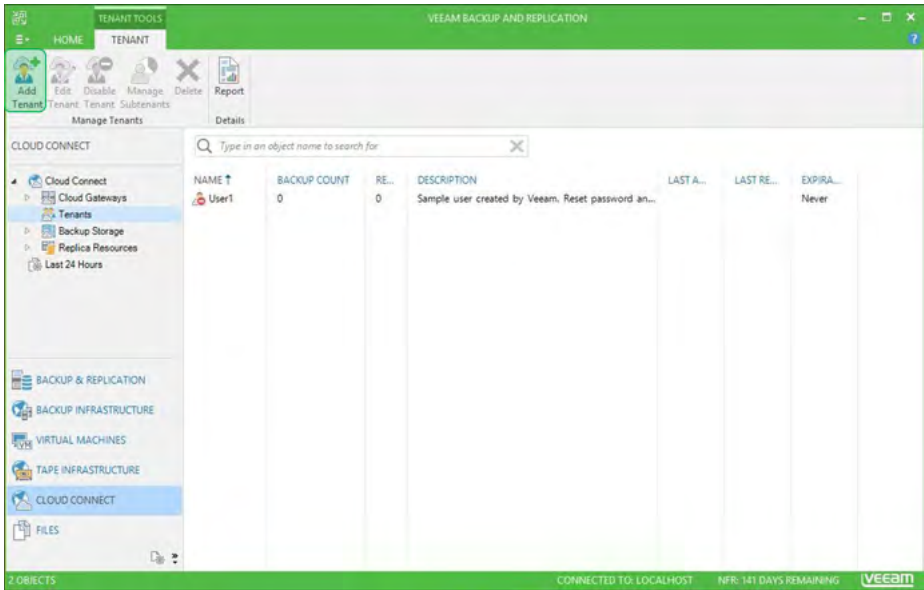


Figure 36: Add Tenant

Type in the username for that tenant and create a password (or use the **Generate new** button) for the tenant. Select **Backup storage in Assigned resources** and select a date for the lease ending period (if any).

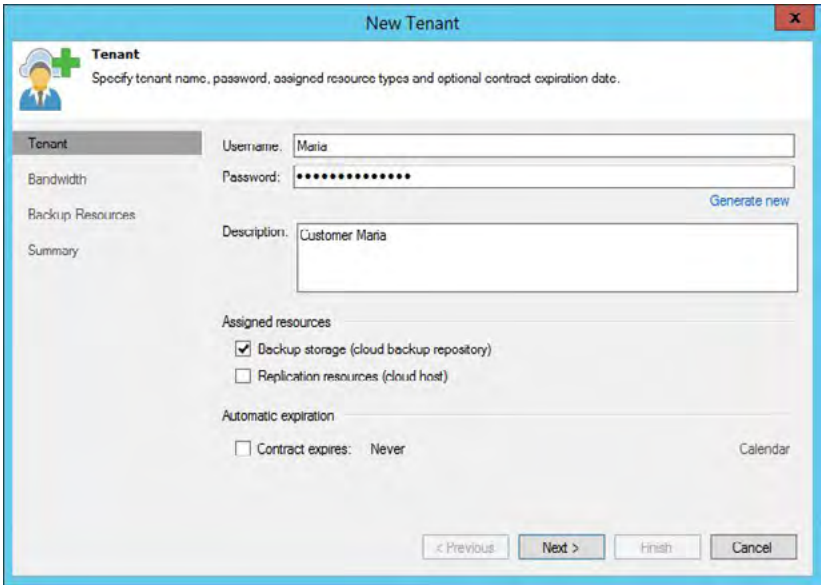


Figure 37: Create a new tenant

In the **Bandwidth** page, IT administrators can configure how many virtual disks can be copied in parallel to the tenants' cloud repository, and whether the inbound bandwidth for this tenant should be limited or not.

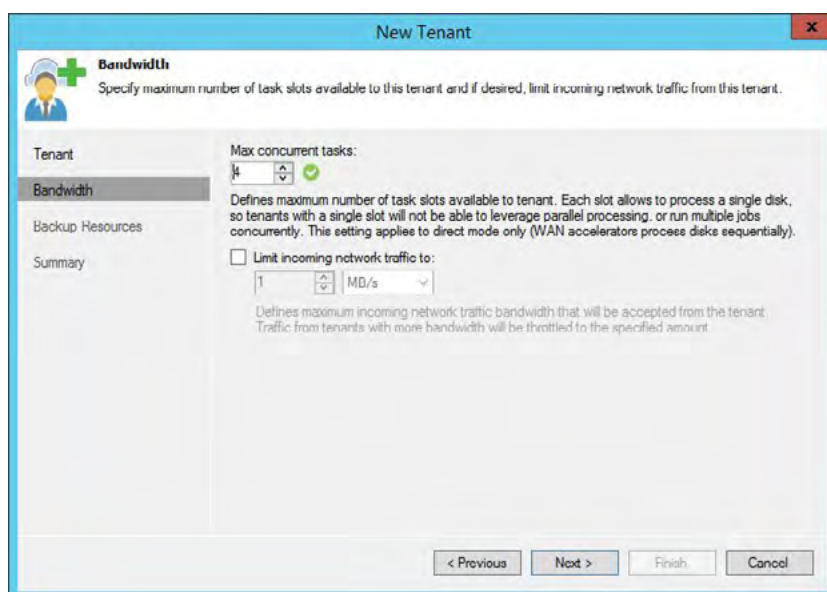


Figure 38: Bandwidth configuration for the tenant

The next page shows the **Backup Resources** page. This is where you will add the allowed quota and repository (or multiple). Press the **Add** button.

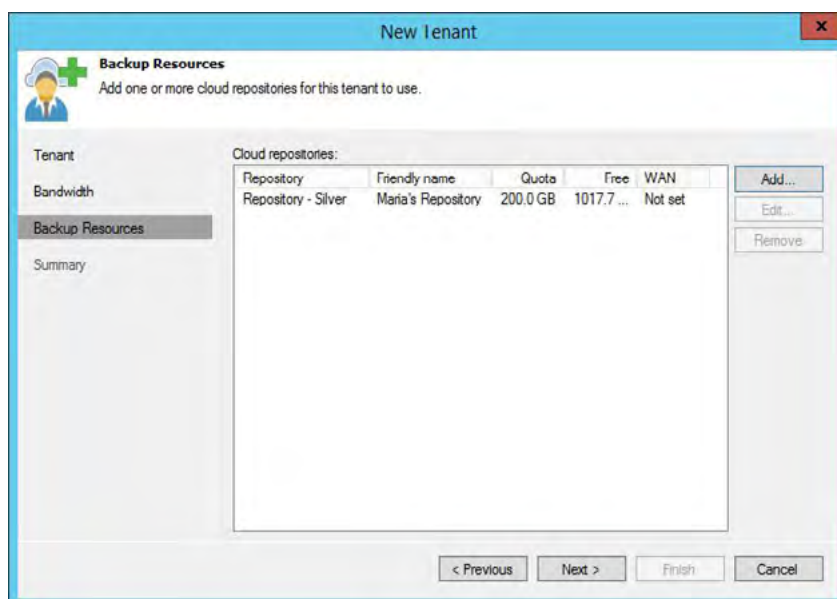


Figure 39: Add resources for the tenant

Type in the friendly name for the cloud repository — this name will be seen by the tenant. Next, select the effective backup repository you previously created and fill in the user quota. Optionally, choose the WAN accelerator that the tenant is permitted to use.

As you can see, the tenant now has a resource available within your environment. In the case where the IT administrators have multiple resources for the tenant, they can be added as well.



Review your summary, and you are ready.

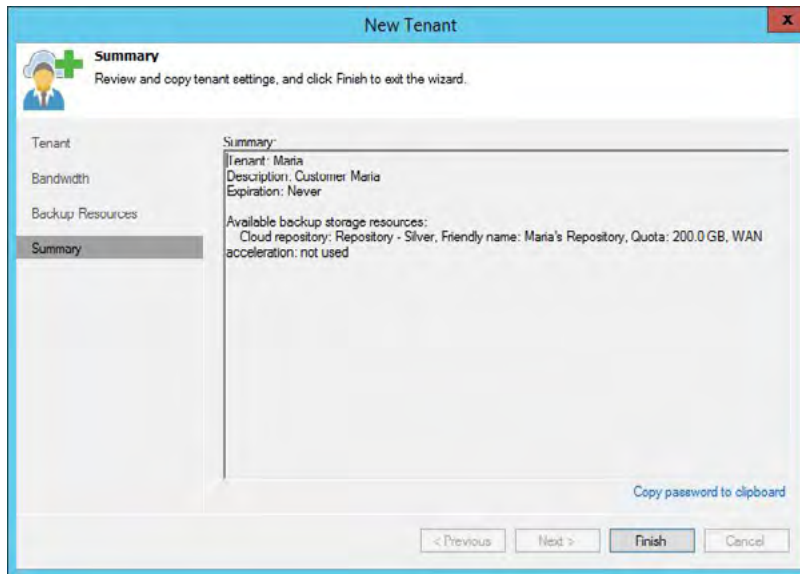
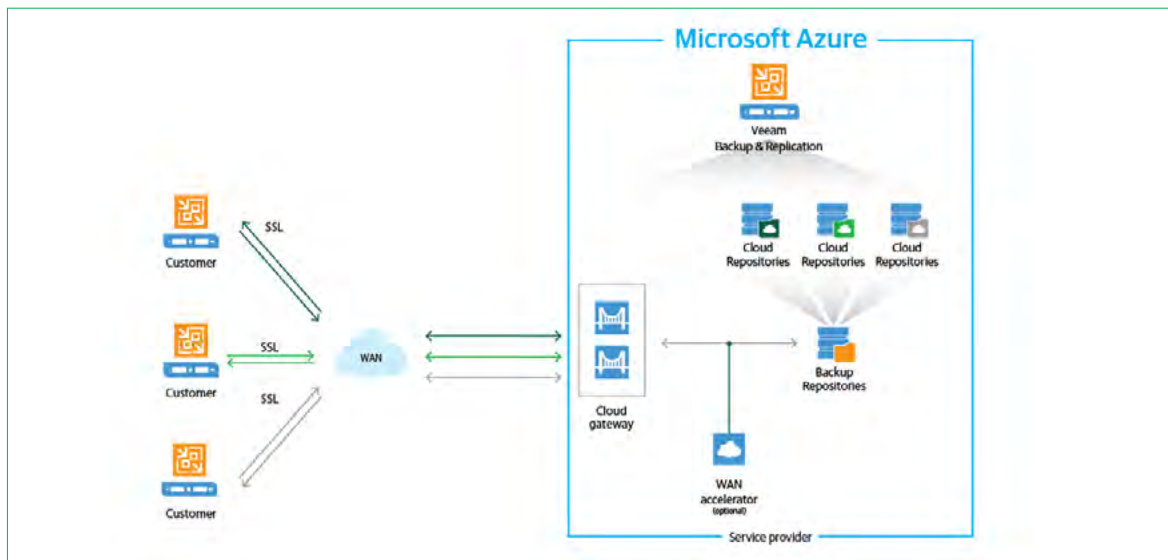


Figure 40: Tenant is ready for use

At this point, if a single-VM deployment is what's right for you, then jump straight to [The end-user side!](#) Otherwise, read on and learn how to deploy a more advanced, distributed and scale-out Veeam Cloud Connect in Azure!

## Distributed Veeam Cloud Connect infrastructure

Expanding the capacity from a simple, single-VM setup is easy to do thanks to the distributed model of Veeam Backup & Replication and the Azure's rapid resource provisioning capabilities. The diagram below illustrates a distributed model.



As the on-premises infrastructures, whether they are a data center, departments or subsidiaries, send more data to Azure, IT administrators can easily separate the roles on different VMs running in Azure. All of the components described previously can be installed on different servers. IT administrators can also have multiple servers running repositories alone, multiple servers running the cloud gateway role (with a single DNS name for the pool) and even dedicate servers for WAN accelerators. The important thing to keep in mind is that the different roles will talk to each other through specific ports, so IT administrators will need to configure Azure to allow communication between the running VMs.



## Scenario

Now that we have our base Veeam Cloud Connect infrastructure up and running, we'll reuse it as a starting point for our distributed architecture. But we will separate the roles and add some dedicated virtual machines for the repository and cloud gateway roles. In order to get closer to a real-life environment, we'll add more than one cloud gateway and configure load balancing and Availability for them, and more than one repository server and configure them under an unlimited Scale-out Backup Repository™.

For the purpose of this white paper, here is what the distributed architecture will look like:

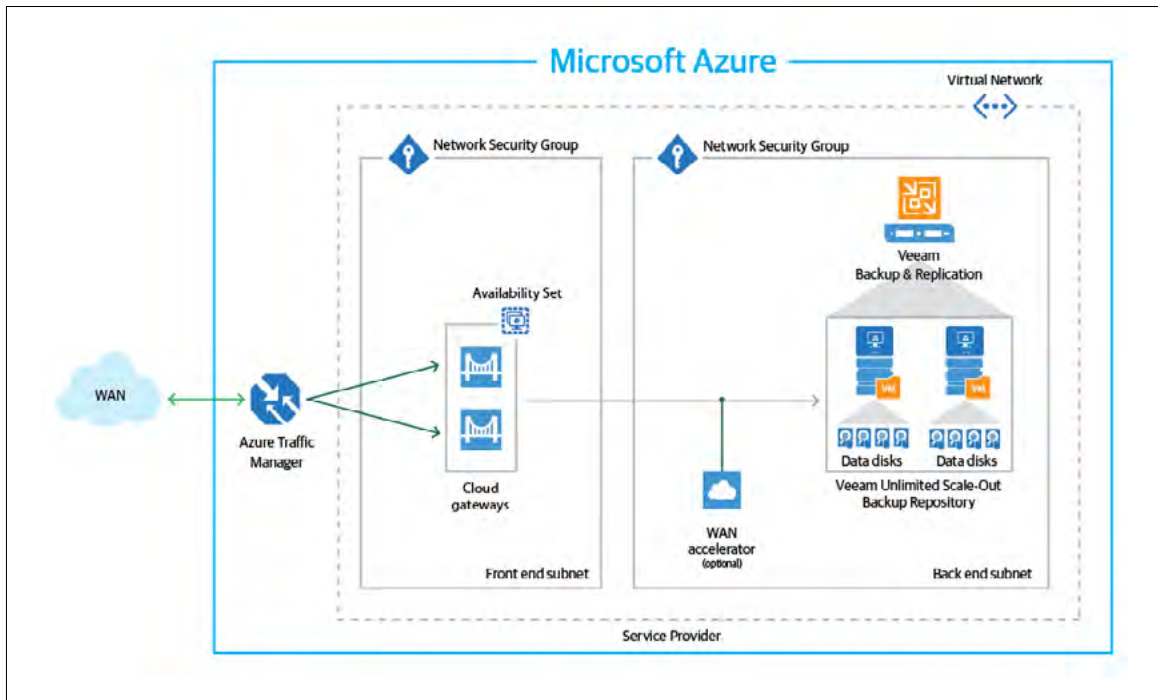


Figure 41: Example of a distributed Veeam Cloud Connect infrastructure in Azure

It will be composed of:

- One Veeam Backup & Replication server, deployed from Azure Marketplace, and configured in the first part of this white paper
- Two VMs with the repository role, each of them having four data disks attached. These will be configured under a single Scale-out Backup Repository
- Two VMs with the cloud gateway role. They will be configured in an Availability set to ensure High Availability of the service
- Two subnets with one network security group each to ensure only necessary incoming communications are allowed in the right perimeters
- One Azure Traffic Manager profile to provide a unique entry point to the cloud gateways pool

## Deploy and configure additional repositories

To make sure our already deployed backup server will use its resources only to manage the Veeam Cloud Connect infrastructure, we'll move unnecessary roles to other dedicated VMs. Let's start with the repository role.

Back to the Azure portal, click on the plus symbol button to deploy a new resource, and search for Windows Server. As Windows Server 2016 is now generally available — and in Azure — we'll choose it which will allow us to leverage new capabilities of ReFS 3.1 and the BlockClone API which Veeam Backup & Replication 9.5 and later integrate with. If this new version of Windows is not yet validated in your company to be utilized, you can fallback to Windows Server 2012 R2.

To learn more about the Veeam Backup & Replication 9.5 integration with ReFS 3.1, see: <https://www.veeam.com/blog/advanced-refs-integration-coming-veeam-availability-suite.html>

Select the version of Windows that works best for you and click **Create**.

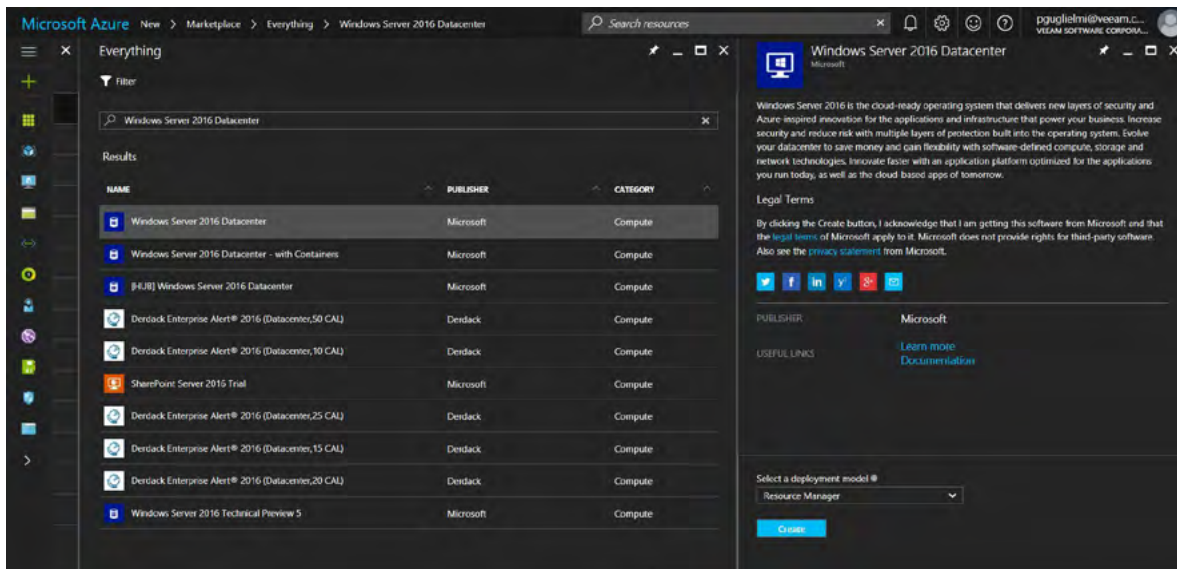


Figure 42: Search Windows Server in Azure Marketplace

Just like we did for the first VM, we now have settings to configure, starting with a name for the VM (**pfg-vcc-repo1** in this example), a username and a password, and a subscription. For the resource group, we recommend to choose the same as the first VM, and the future ones, as they are all part of the same infrastructure and share the same lifecycle for the duration of these tests. And finally, we choose the same Azure data center – which is North Europe in this case – and click **OK**.

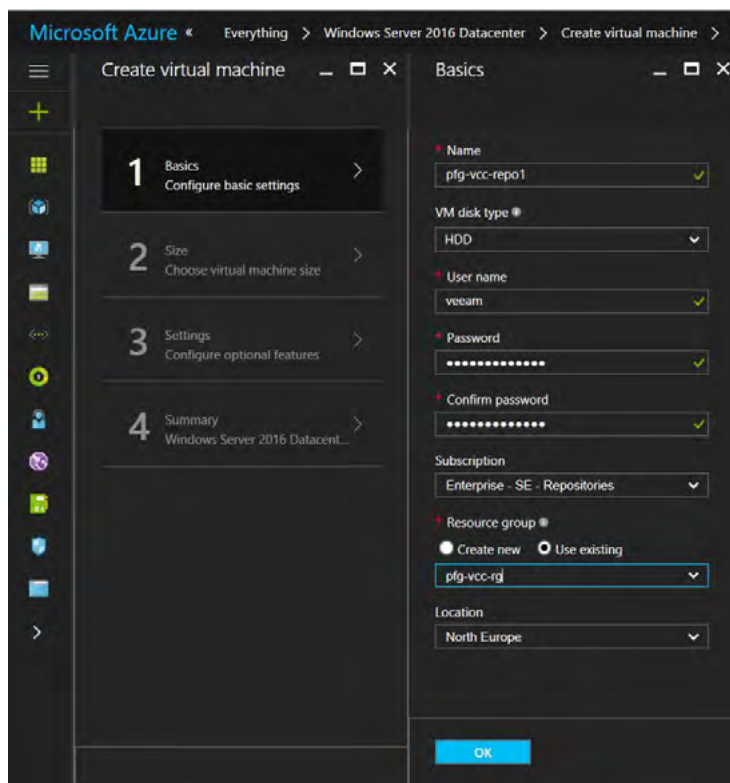


Figure 43: Configure new VM basic settings

Next, we have to choose a VM size. Besides a repository server's requirements, there are a few things to consider when choosing the VM size in Azure. We need storage capacity, but have to take into consideration that today the maximum size of a virtual hard disk in Azure is 4095 GB. With this in mind, we can add multiple virtual disks to our repositories. The larger the VM size, the more data disks you can add, but larger VM configurations come also with more cores and memory as well as the VM cost will increase accordingly. From cost and resources usage perspectives, it can be interesting to choose a smaller VM size for repository servers with a lower number of maximum data disks, and multiply such VMs, especially since we will be able to add them all as extents in an Unlimited Scale-out Backup Repository. Therefore, we will choose the A2\_V2 VM size for the repositories and click Select. Each repository server will have a maximum of four data disks, along with two cores and 4 GB of RAM.

**Note:** The VM size chosen for the repository server VMs is good enough for the purpose of this white paper, and to start with a rather small environment. Consider choosing VMs with more resources – especially RAM – when many parallel tasks are run by the tenants and when the **Use per-VM backup files** option is selected.

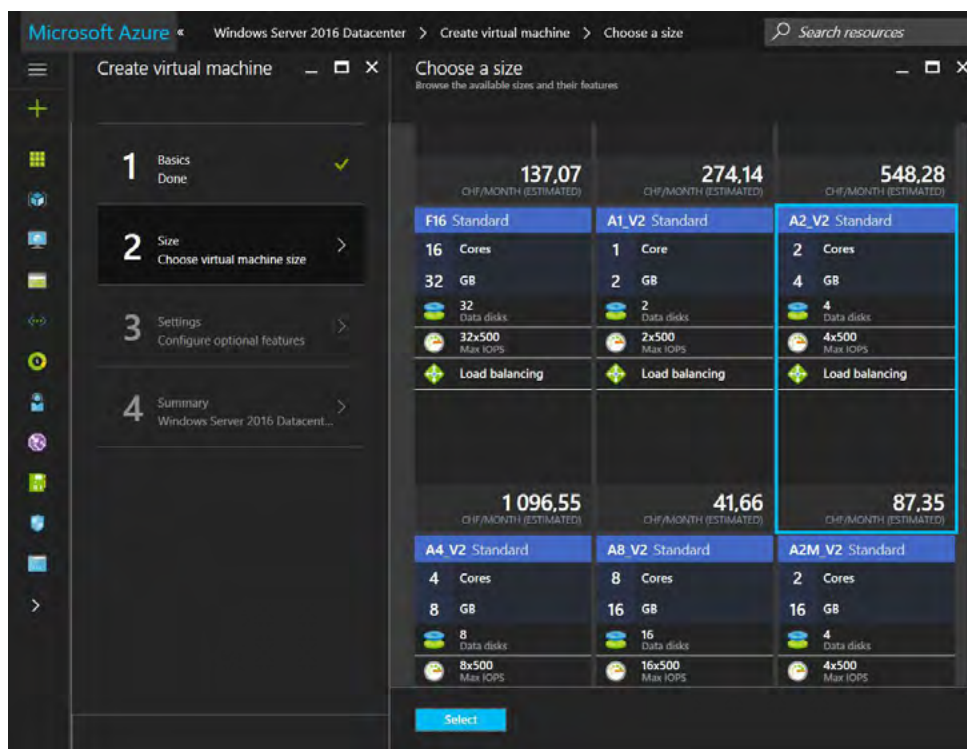


Figure 44: Choose a VM size

For the repository servers, we choose the same storage account for the OS disk as the Veeam Backup & Replication server deployed earlier (**pfgvccvhdssa**), the same virtual network, **Internal** subnet, network security group and other settings. Click **OK**.

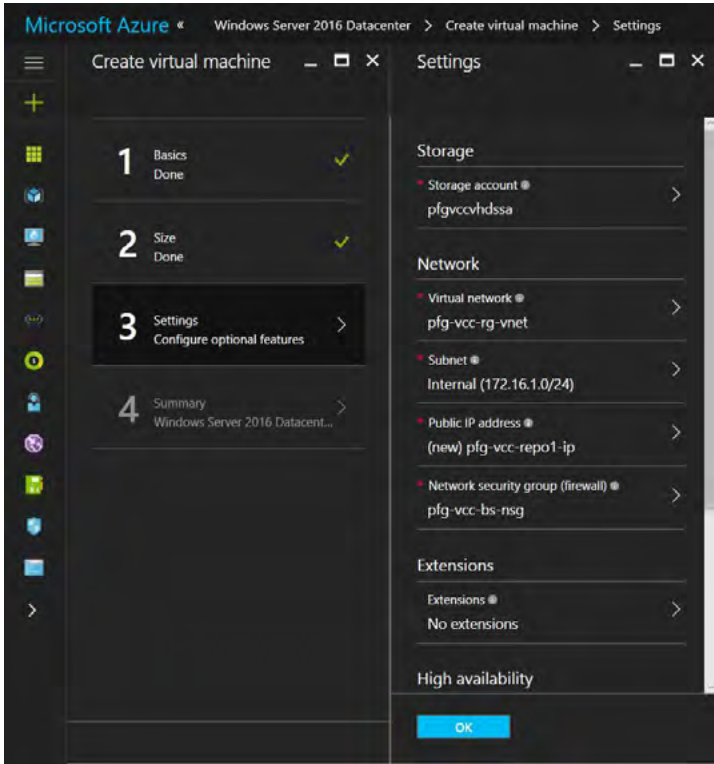


Figure 45: Settings and optional features for the new repository server

Review the summary of the settings and click **OK**.

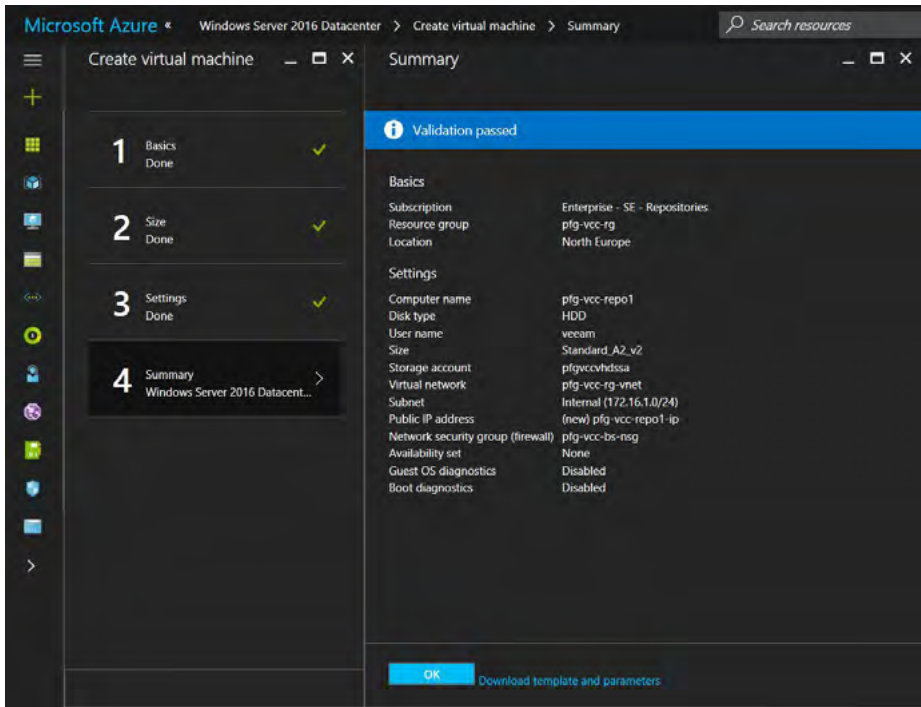


Figure 46: Summary of the new repository server creation

After a few minutes, your new repository VM will be ready for configuration.

The first configuration step is to set a DNS name label to the new VM because the default IP address is dynamic. To do so, select the new repository virtual machine in the Azure portal, within the **Overview** section, click on **None** under **DNS name label**.

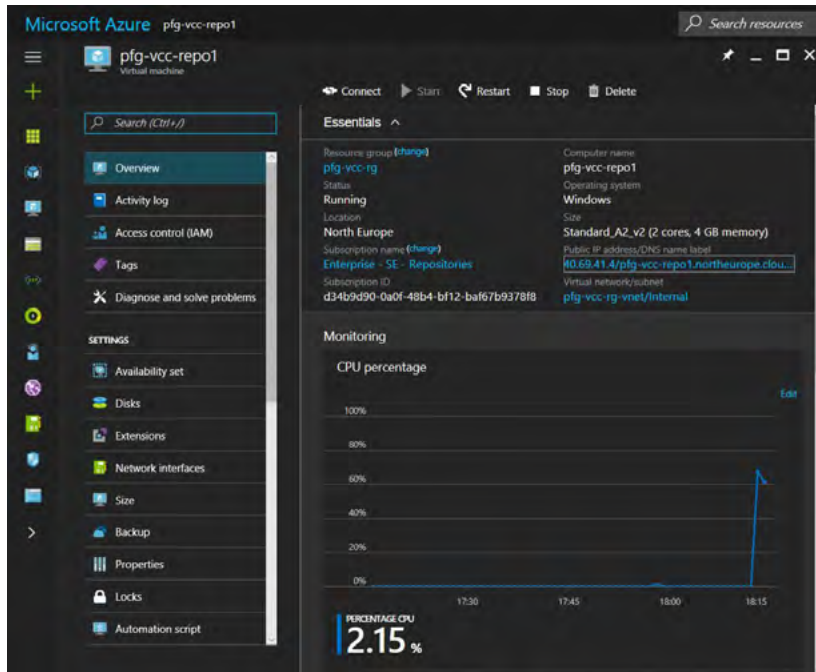


Figure 47: DNS name label configured on the new repository server

Because this new VM will be a Veeam repository server, create and attach new data disks which will store the backup copies sent by the tenant. In this example VM, we're using an A2\_V2. We'll add the maximum number of data disks, which is four.

Click on **Disks** and then **Attach new**.

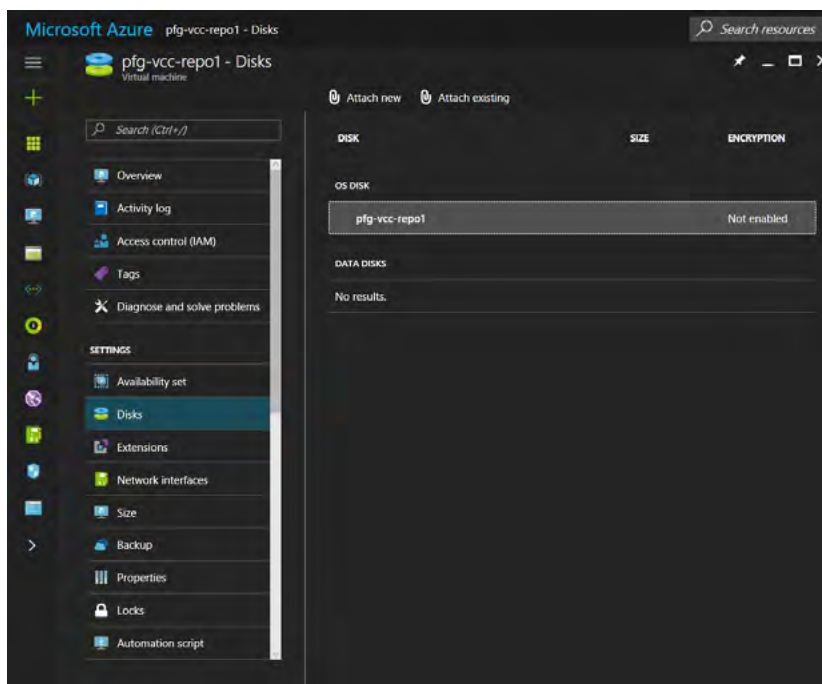


Figure 48: Repository VM disks configuration



While creating the first repository disk (**pfg-vcc-repo1-disk1**) to this first repository VM (**pfg-vcc-repo1**), I'm also creating a new storage account (**pfgvccreposit**) that will store exclusively the repository data disks. It is recommended to plan your storage accounts upfront, according to your own needs and constraints.

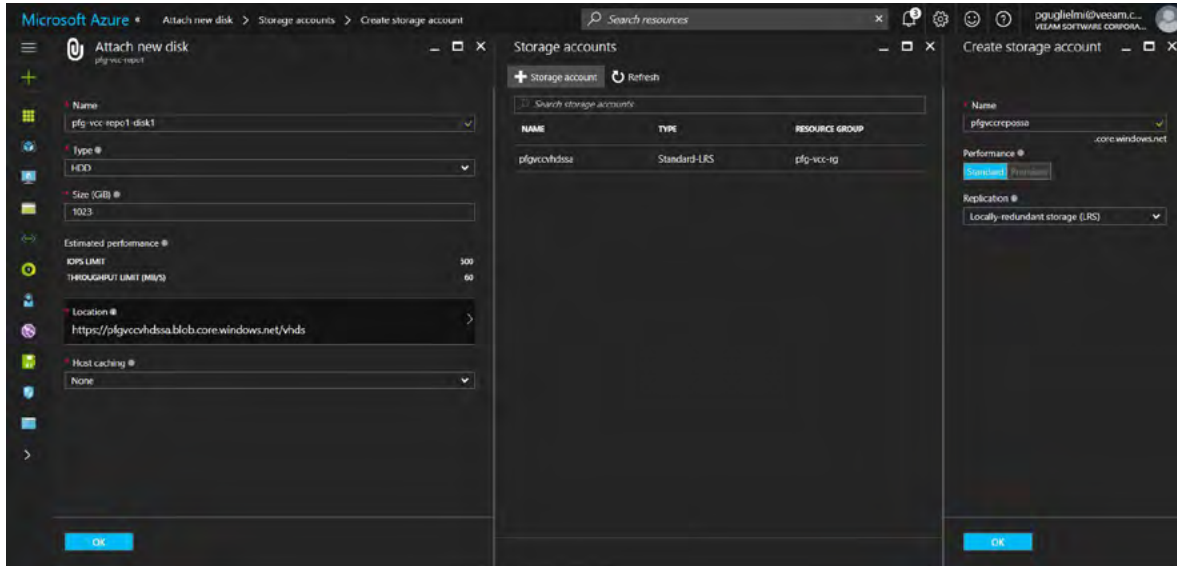


Figure 49: New storage account for repository data disks

Review the settings of the new data disk and click **OK**.

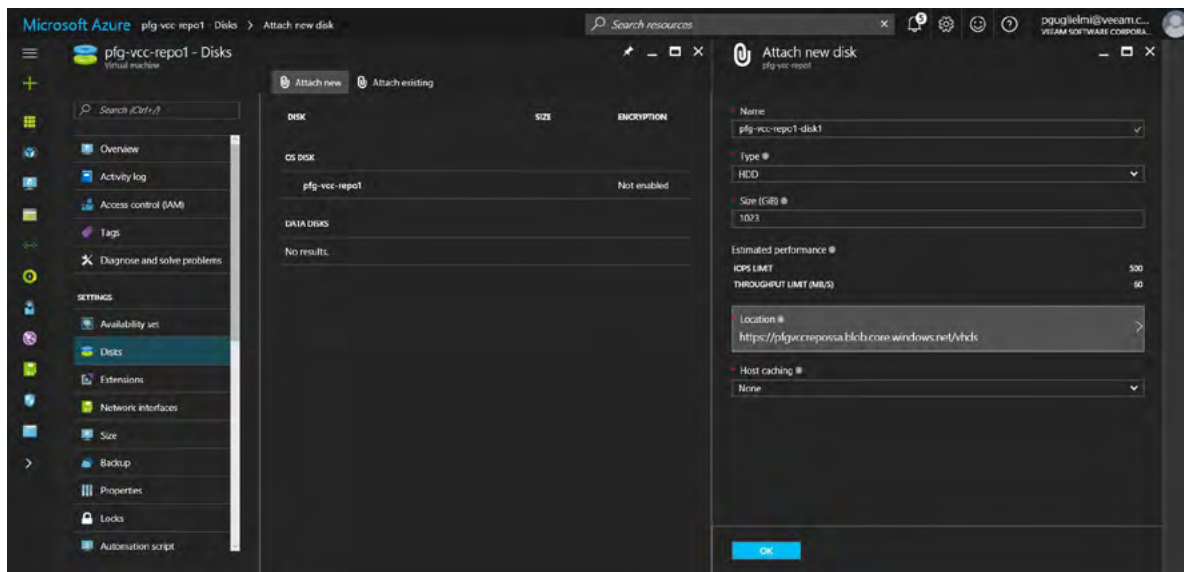


Figure 50: Attach new disk settings



Repeat the steps to attach more data disks until you have what you want. In this case, we end up with four new 1,023 GB data disks.

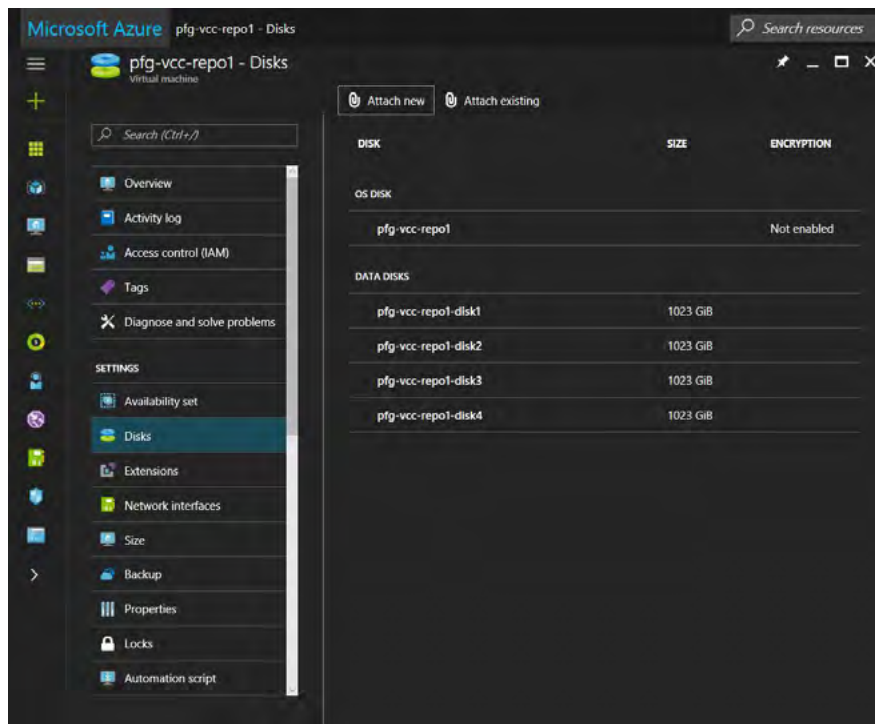


Figure 51: All four new data disks have been attached

Before configuring the settings within the guest OS, deploy the other additional repository servers with the exact same settings. In this example, we have two repository VMs, **pfg-vcc-repo1** and **pfg-vcc-repo2**, each with the maximum — four data disks.

Now that we have our repository VMs deployed with data disks attached, we need to connect to each of them to configure a few things like the Windows firewall and the new disks.

In the Azure portal, select the repository server you want to configure and go back to **Overview** and click **Connect**. Once connected, open the Windows firewall .

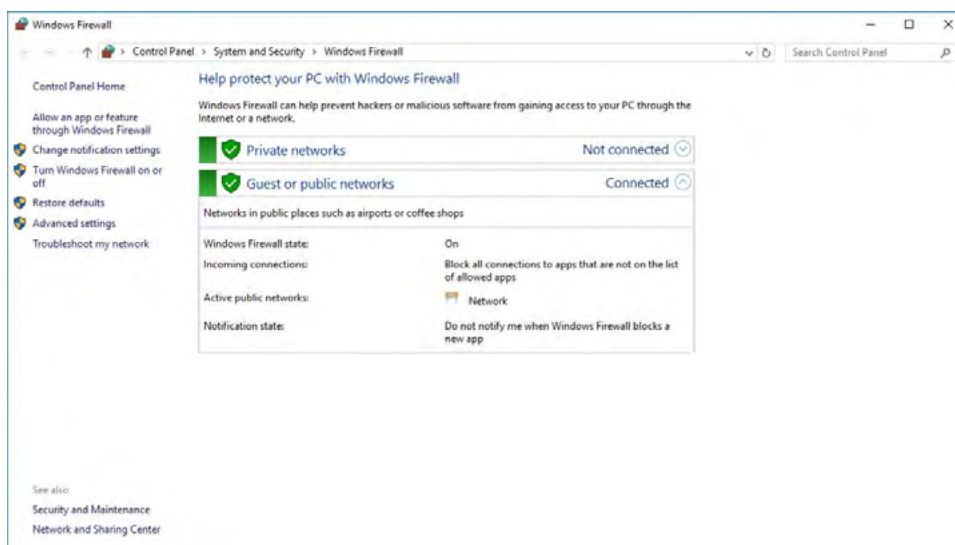


Figure 52: Windows firewall

We'll need to make this VM an actual Veeam Backup & Replication repository. To achieve that goal, the necessary Windows services are installed, and the configuration is done from the Veeam Backup & Replication console. The services are installed by pushing .msi packages through the Windows admin shares, which is why we need to temporarily allow the **File and Printer Sharing** app in the Windows firewall.

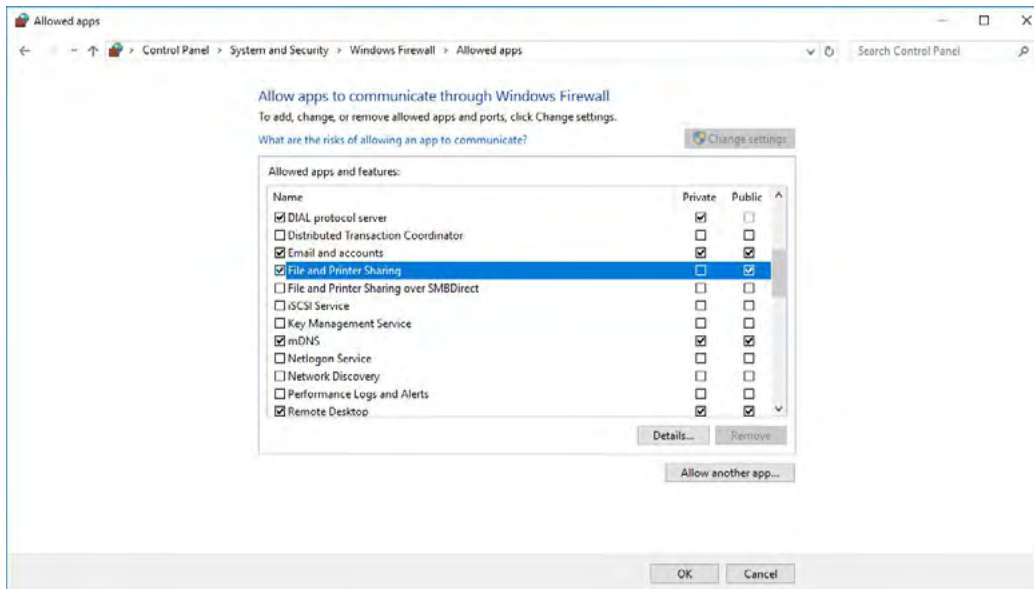


Figure 53: Allow File and Printer Sharing in the Windows firewall

Next, create the new volumes on the data disks we've attached previously. Let's open the **Disk Management** console.

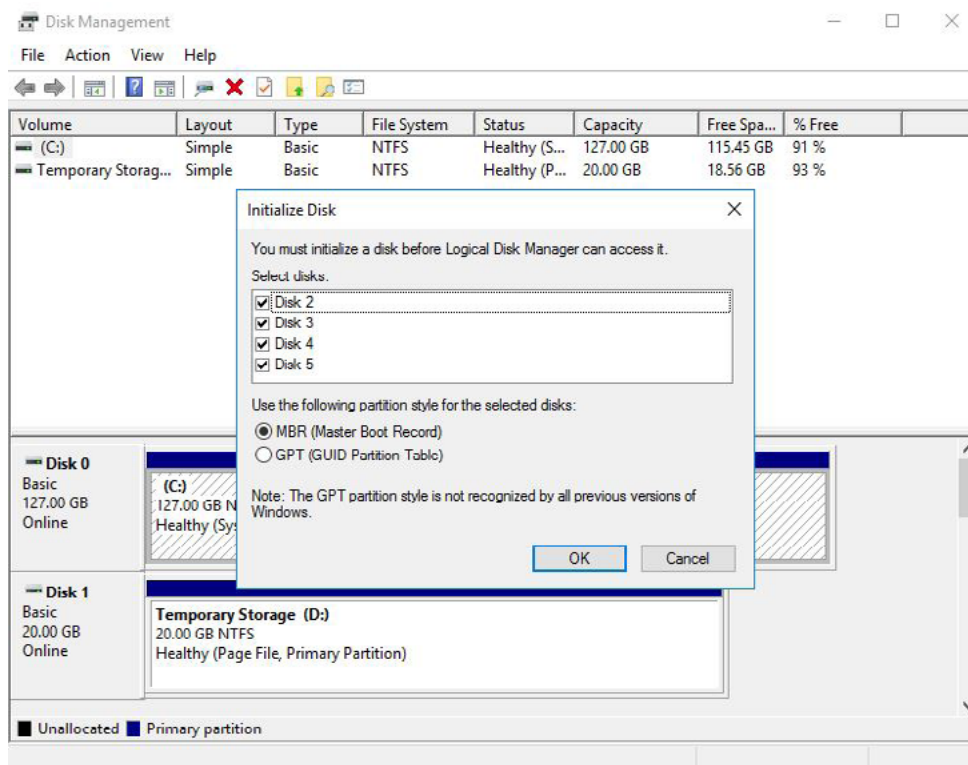


Figure 54: Windows Disk Management console

Windows detects the new data disks automatically and offers to initialize all disks at once. Click the **OK** button. Now it's time to create the volumes. Right-click on the first new disk and select **New simple volume**.

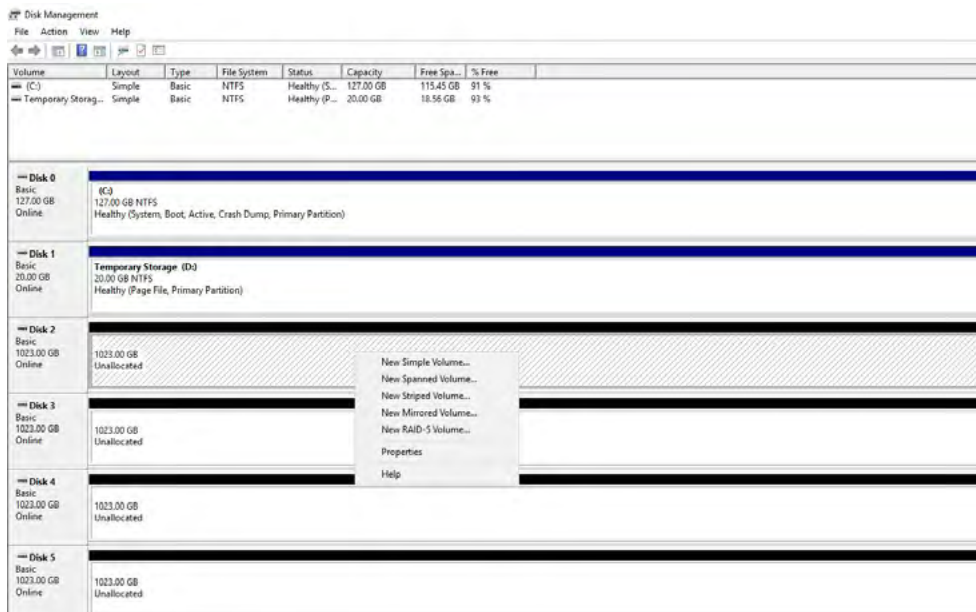


Figure 55: Launch the New simple volume wizard

Follow the standard Windows volume creation wizard. In this example, at the **Format partition** page, we give a name to the volume (**Repo1-Disk1**), and select **ReFS** as a file system.

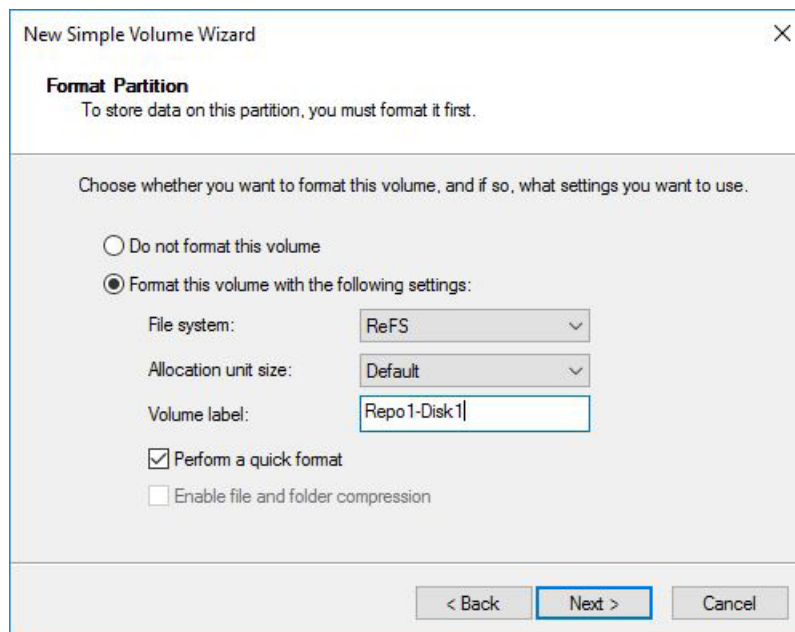


Figure 56: File system selection

Complete the wizard, and repeat the steps to create the additional volumes on the other data disks. In this case, we have

four new ReFS volumes per repository server.

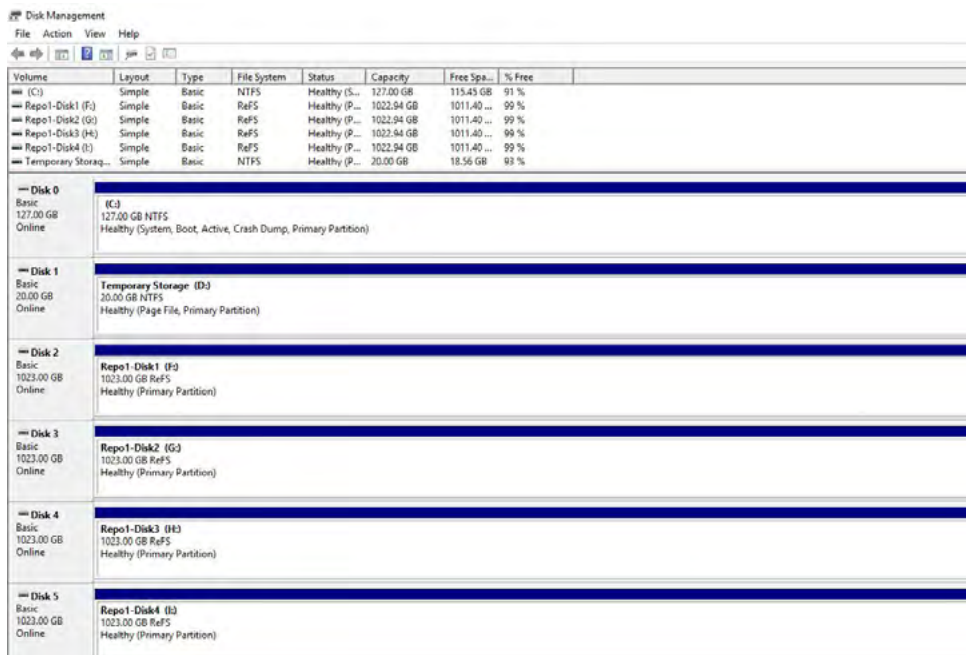


Figure 57: Windows Disk Management showing four new ReFS volumes

Sign out of the virtual machine, and connect to the other VMs and repeat the configuration steps. Once complete, connect to the Veeam Backup & Replication server (the first virtual machine we've deployed from Azure Marketplace), to add the repositories to the console.

From the Azure portal, select your Veeam Backup & Replication server and click **Connect**. When connected, double-click on the Veeam Backup & Replication icon on the desktop. The console will launch and prompt for server to connect to and credentials – click **Connect**. When the console is open, go to **Backup Infrastructure** and select **Backup Repositories** in the left pane. Click on the **Add Repository** button in the ribbon.

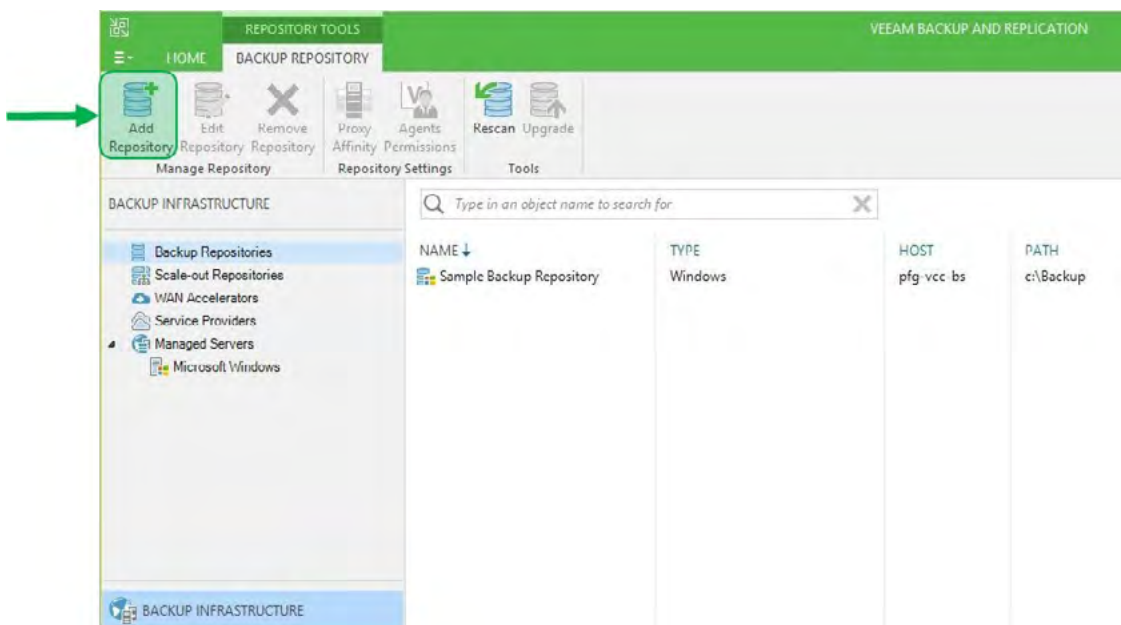


Figure 58: Add Repository

Similar to the steps in [Initial configuration](#), we need to add our eight data disks attached to our two repository servers as Veeam Backup & Replication repositories. The difference being that we will have to add the two repository VMs to the console during the operation.

In the **New Backup Repository** wizard, type in the name and description for the repository. Click **Next**.

Figure 59: Name the repository

On the **Type** page, select **Microsoft Windows** server and press **Next**.

Figure 60: Choose Microsoft Windows server

The repository disks we need to add are attached to a separate VM, add the server via FQDN / DNS hostname to the console. On the **Server** page, click **Add New**.

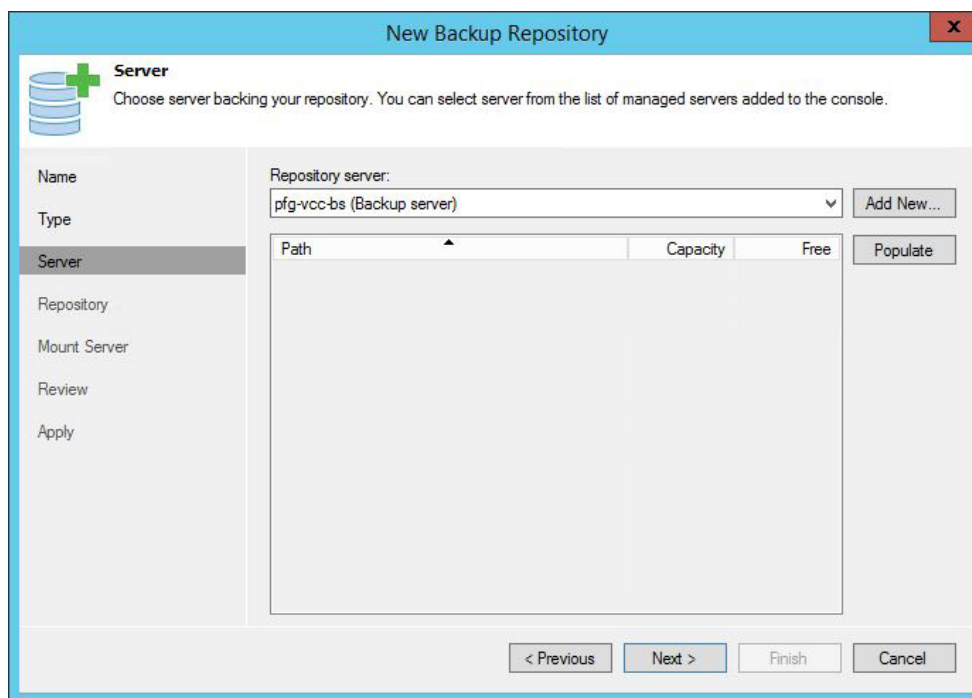


Figure 61: Add a new server

The **New Windows Server** wizard opens. Enter the hostname or IP address of your repository VM. In this example, the hostname is **pfg-vcc-repo1**.




Figure 62: Enter the name of the new repository server



This wizard allows to add a new machine as a Veeam Backup & Replication repository server by deploying the necessary components, as explained earlier. To achieve that, we've prepared the guest OS by allowing the **File and Printer Sharing** in the Windows firewall (Figure 53), and now we need to enter admin credentials to be able to push the .msi packages and install them.

On the **Credentials** page, click **Add**.




Figure 63: Add credentials

Enter a username and password for the new VM, and an optional description. The user has to be a local administrator of the VM. Click **OK**.

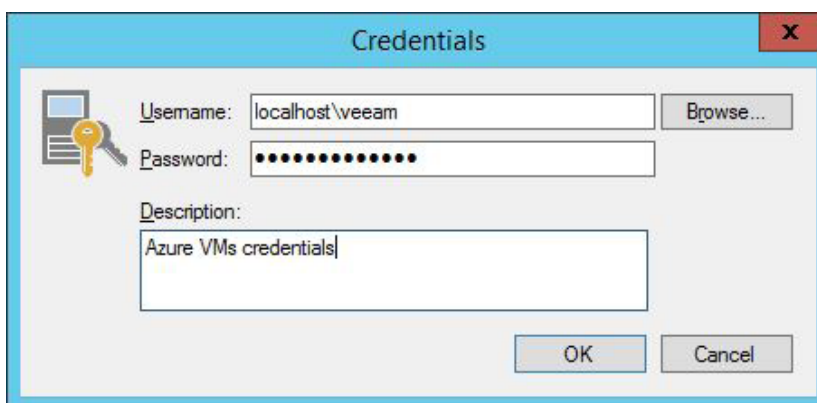


Figure 64: Enter credentials

On the Credentials page of the wizard ([Figure 63](#)), now you can see your newly entered credentials selected. Click **Next**.

Review the settings and click **Next**.

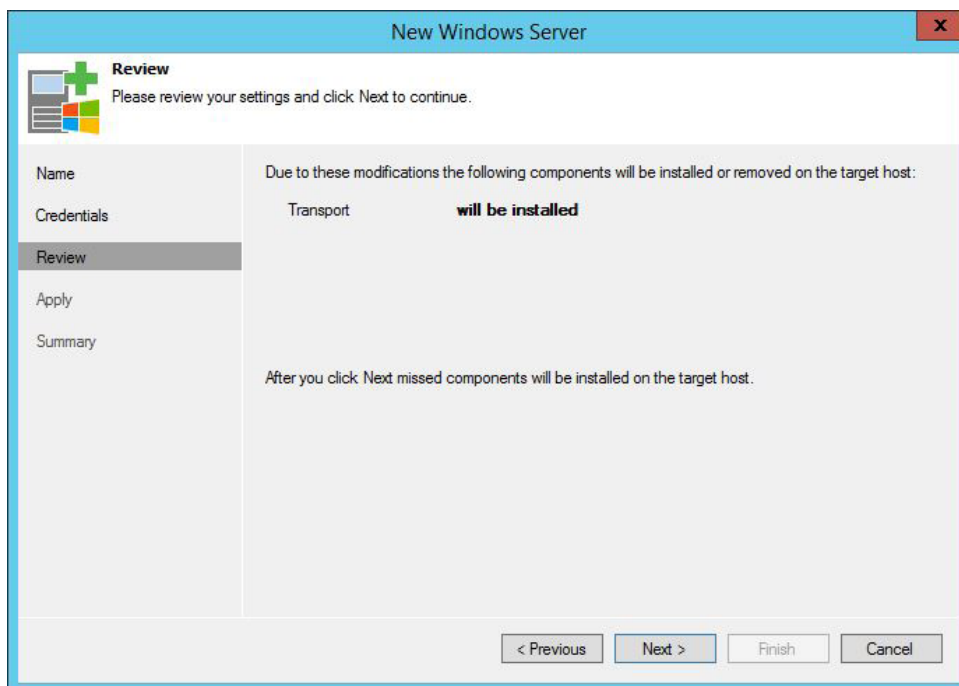


Figure 65: Review page

Within the next step, you can see the actions taken automatically by the Backup & Replication to register a new Windows server to the Veeam infrastructure. Click **Next**.

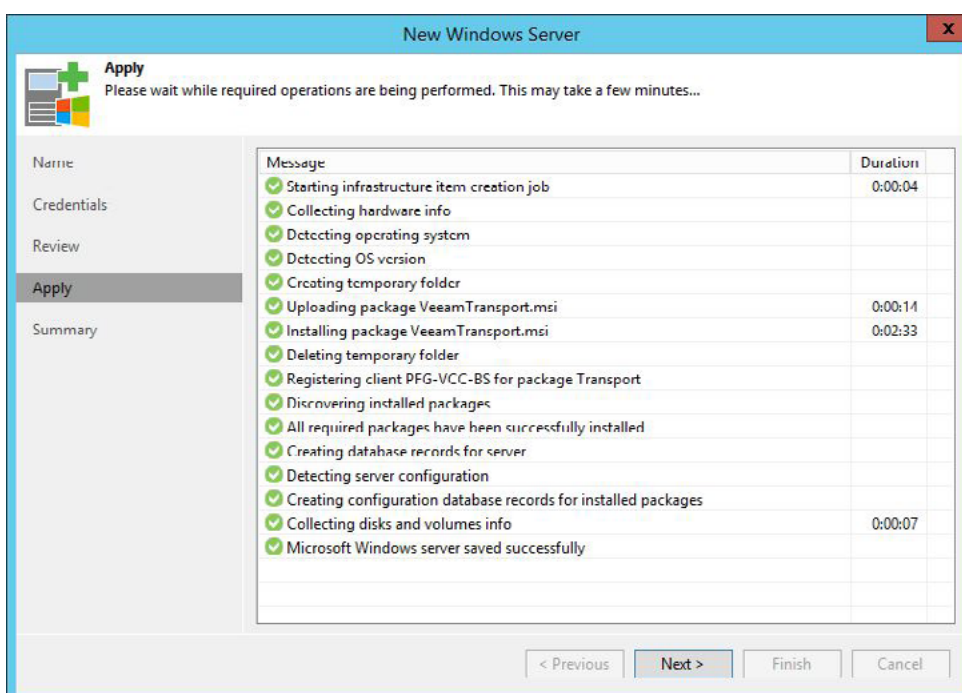


Figure 66: A new Windows server is being added

On the last screen of the **New Windows Server** wizard, you can review the summary and click **Finish**.

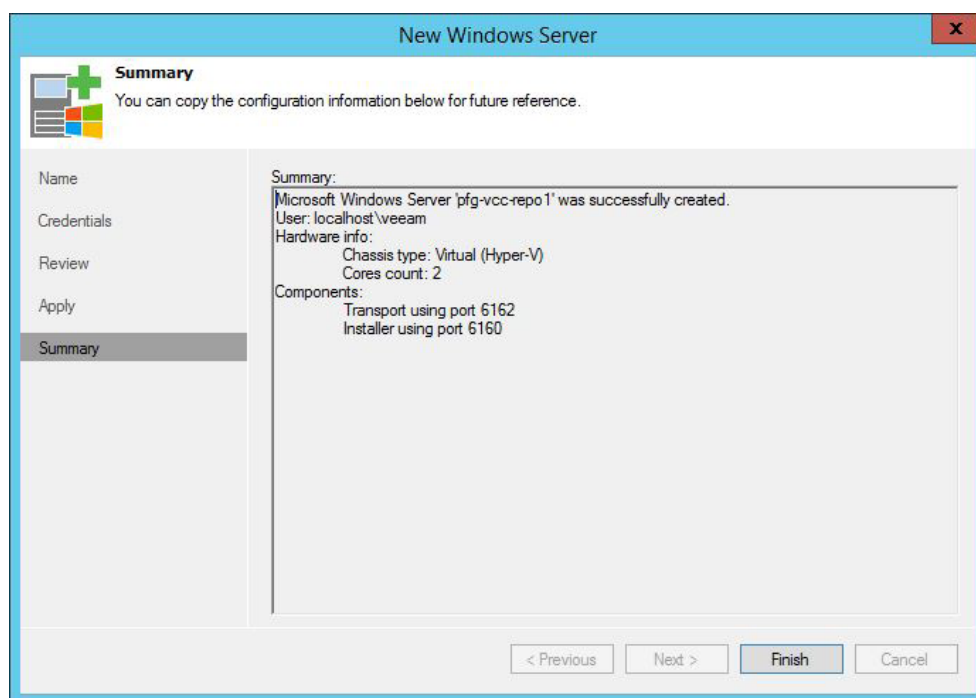


Figure 67: Summary

In the New Backup Repository wizard, we can see the newly added repository server with available volumes attached to it. Select the first volume created earlier on the first data disk added to the repository virtual machine. In this example, it's the 1,023 GB F: drive. Click **Next**.

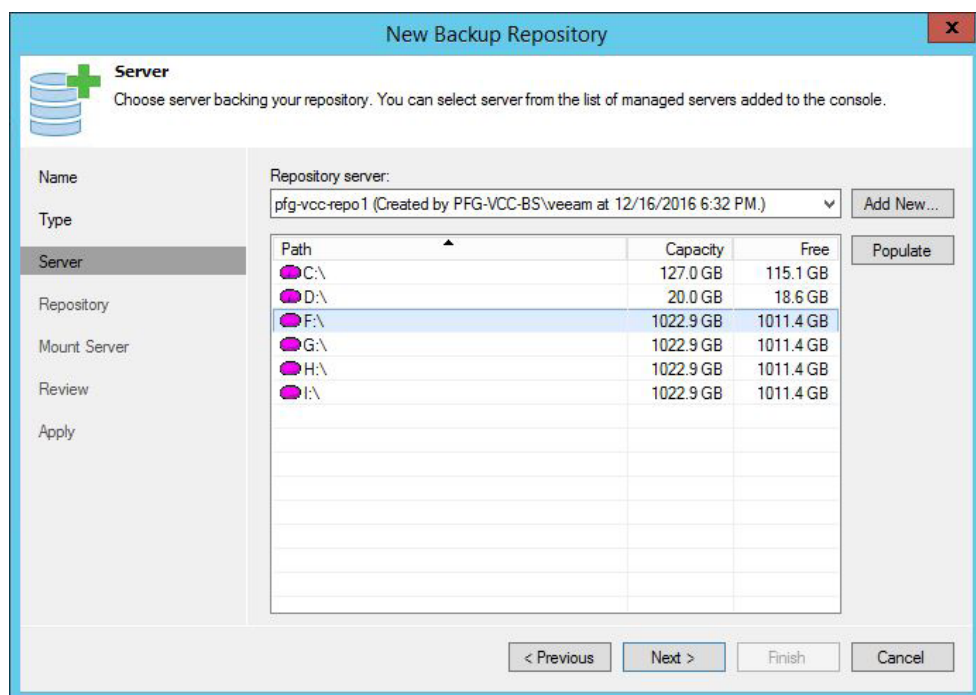


Figure 68: Choose the destination drive

On the Repository page, type in a path to a folder to store the backup files. Here you can also control the load this repository can handle, either by number of concurrent tasks or by read and write data rates. For the purpose of this white paper, we just keep the default settings and click **Next**.

Figure 69: Choose the destination path and advanced settings

On the **Mount server** page, deselect **Enable vPower NFS server** as this cannot be used in Veeam Cloud Connect.

Figure 70: Disable vPower NFS

Review your settings and create the backup repository by clicking **Next** and then **Finish**.

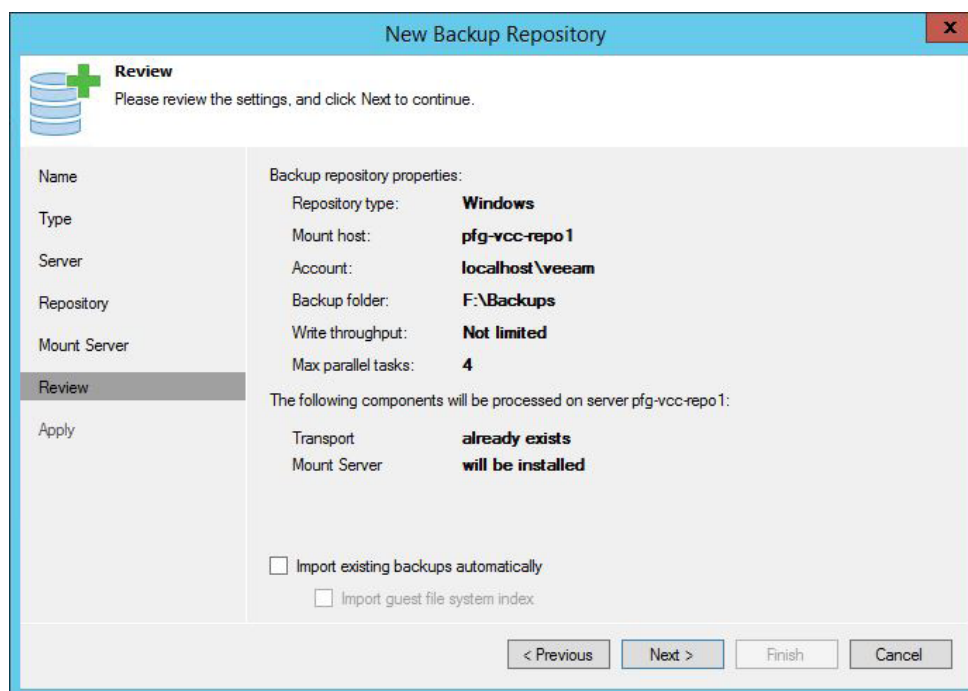


Figure 71: Review settings

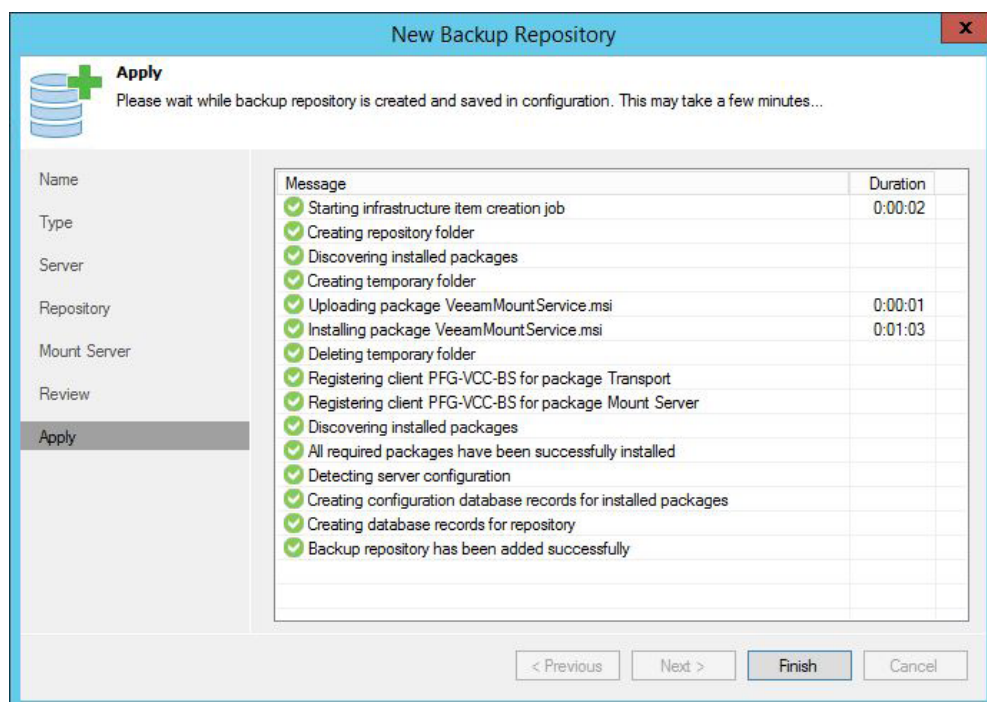


Figure 72: Successful creation of the backup repository

Repeat the steps described above to add all other repositories to the Veeam infrastructure. In this example, we have a total of eight repository disks spread across two virtual machines .

NAME	TYPE	HOST	PATH	CAPACITY	FREE
pfg-vcc-repo1-disk1	Windows	pfg-vcc-repo1	F:\Backups	1022.9 GB	1011.4 GB
pfg-vcc-repo1-disk2	Windows	pfg-vcc-repo1	G:\Backups	1022.9 GB	1011.4 GB
pfg-vcc-repo1-disk3	Windows	pfg-vcc-repo1	H:\Backups	1022.9 GB	1011.4 GB
pfg-vcc-repo1-disk4	Windows	pfg-vcc-repo1	I:\Backups	1022.9 GB	1011.4 GB
pfg-vcc-repo2-disk1	Windows	pfg-vcc-repo2	F:\Backups	1022.9 GB	1011.4 GB
pfg-vcc-repo2-disk2	Windows	pfg-vcc-repo2	G:\Backups	1022.9 GB	1011.4 GB
pfg-vcc-repo2-disk3	Windows	pfg-vcc-repo2	H:\Backups	1022.9 GB	1011.4 GB
pfg-vcc-repo2-disk4	Windows	pfg-vcc-repo2	I:\Backups	1022.9 GB	1011.4 GB
Sample Backup Repository	Windows	pfg-vcc-bs	c:\Backup	126.7 GB	112.8 GB

Figure 73: List of Veeam backup repositories

The goal of adding multiple backup repositories is to add more disk capacity to store more backups, and also to scale-out easily by adding more repositories later. But as explained before, the maximum size of data disks attached to an Azure Virtual Machine is 4,095 GB, and we've added them individually to the Veeam infrastructure. Using them like this would be cumbersome and difficult to manage. Worse off, the backup file chains could not be larger than 4 TB. This would cause unwanted limitations.

In order to achieve true scalability, several options can be considered such as Storage Spaces, Storage Spaces Direct or third-party solutions like StoneFly Scale-Out Cloud Storage. In this white paper we will leverage the Veeam Unlimited Scale-Out Backup Repository feature, which is supported to be used with Cloud Connect since version 9.5. We will utilize the added standard repositories as extents of the new Scale-out Backup Repository.

However, note that individual deduped and compressed backup files must remain smaller than the extents. Should larger backup files need to be stored, then consider configuring Storage Spaces or Storage Spaces Direct as an option to bypass the 4TB file size limitation.

To learn more about Veeam's Unlimited Scale-out Backup Repository, see: <https://www.veeam.com/wp-scale-out-backup-repository.html>



Go to **Backup Infrastructure > Scale-out Repositories**, and click the **Add Scale-out Repository** button in the ribbon.

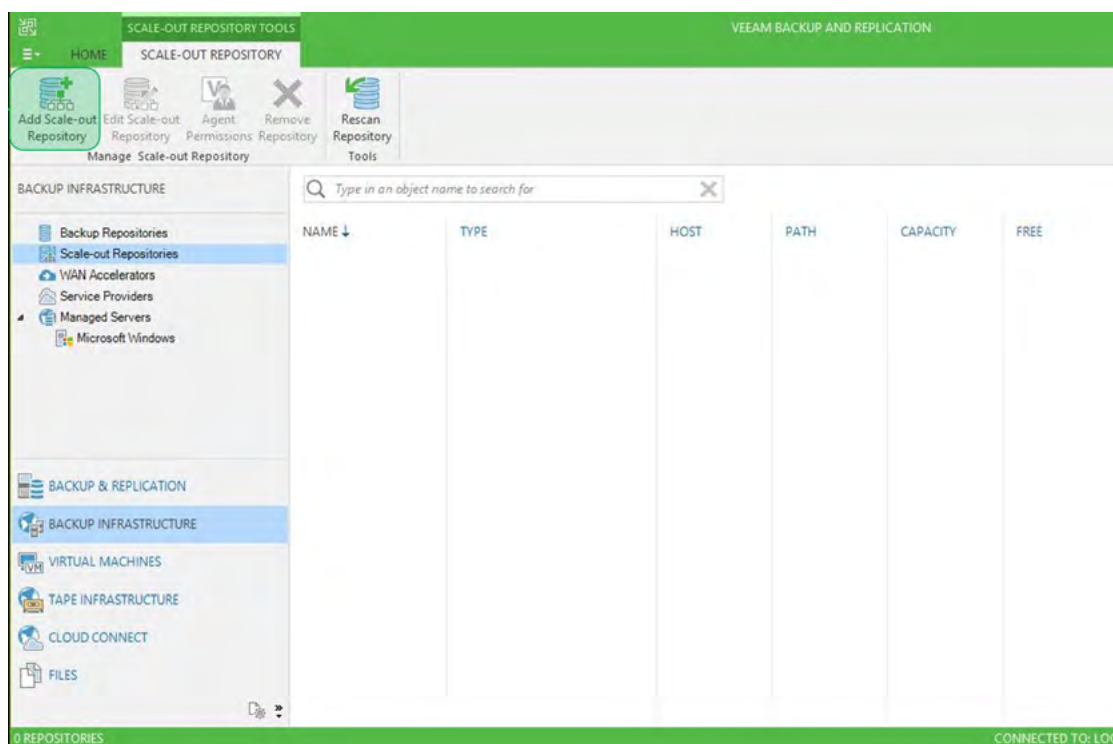


Figure 74: Add Scale-out Repository

Provide a useful name as well as the optional description to the Scale-out Backup Repository and click **Next**.




Figure 75: Name the Scale-out Backup Repository

On the **Extents** page, click **Add**.

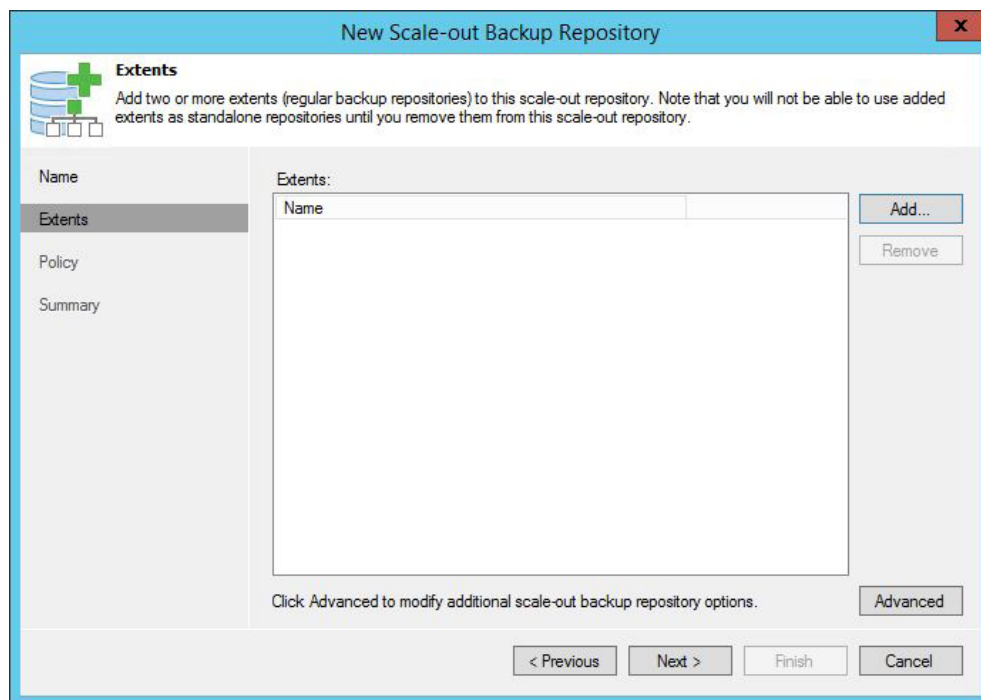


Figure 76: Add extents

Select the repositories you want to build your scale-out one with and click **OK**.




Figure 77: Choose repositories

Back to the **Extents** page, verify the selected repositories and click **Advanced**.

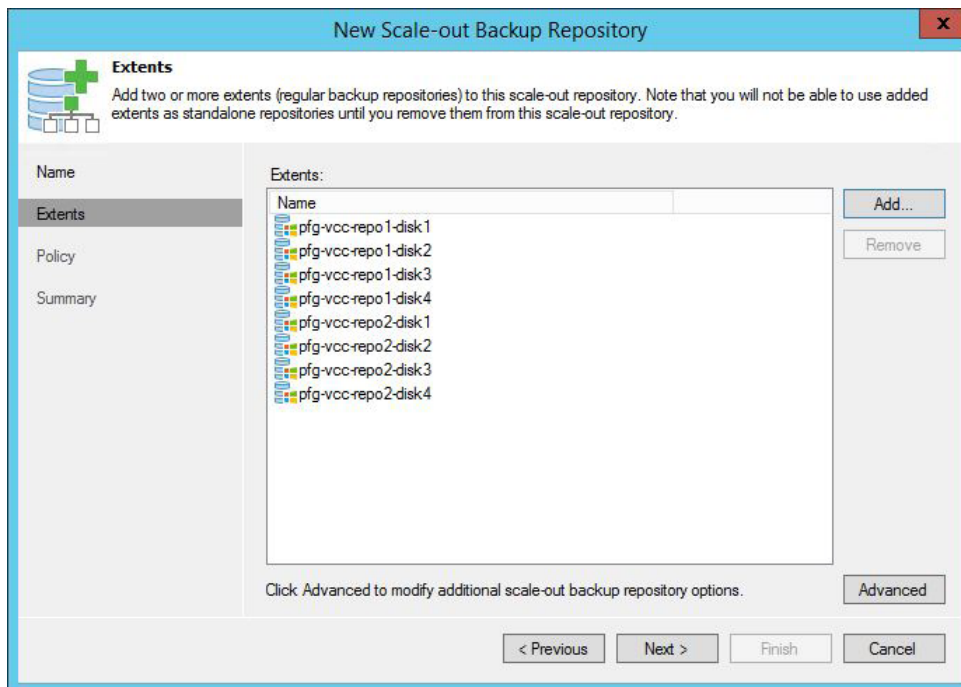


Figure 78: Review selected extents

Select the **Use per-VM backup files** option which is recommended to better utilize disk space on the extents — as well as to make backup files smaller — and click **OK**.

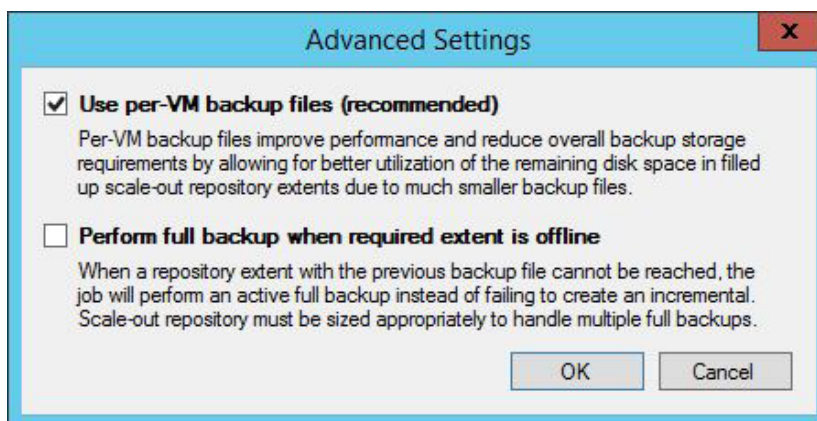


Figure 79: Configure advanced settings

On the **Policy** page, select **Data locality** which better fits an environment where all extents are the same, and click **Next**.

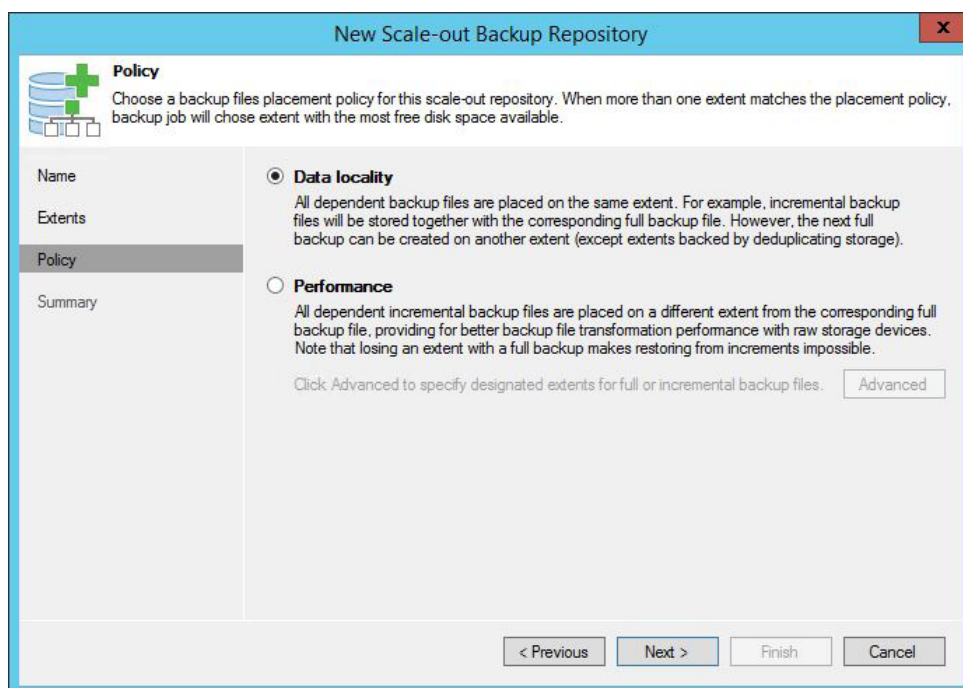


Figure 80: Choose policy

On the **Summary** page, click **Finish**.

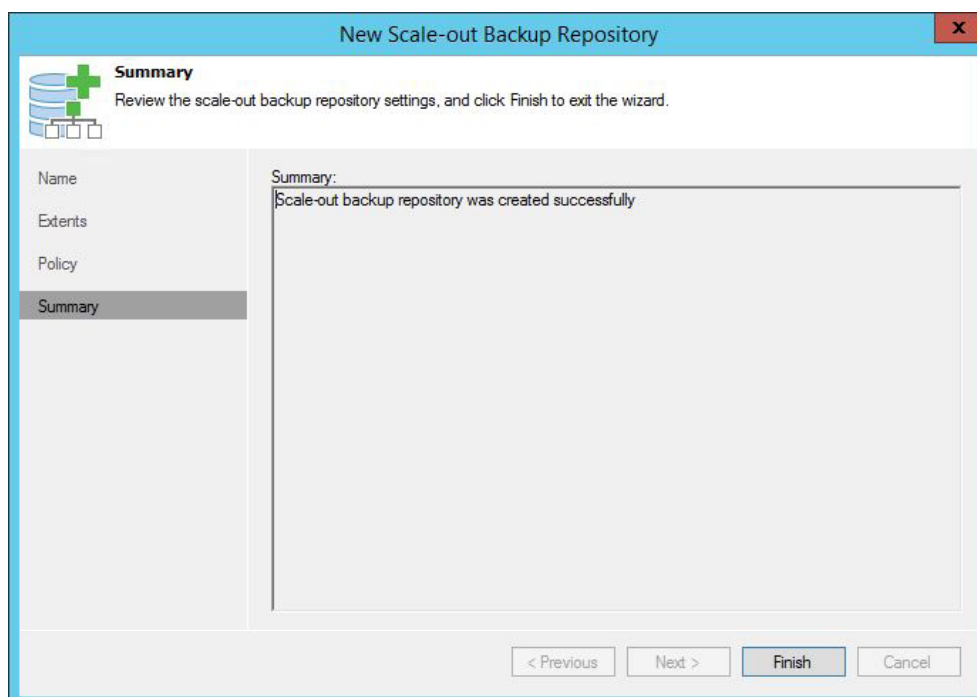


Figure 81: Summary

Your new Scale-out Backup Repository is now ready for use! In this example, we have an 8 TB repository to store customers' backup copies, but you can create much larger repositories if necessary.

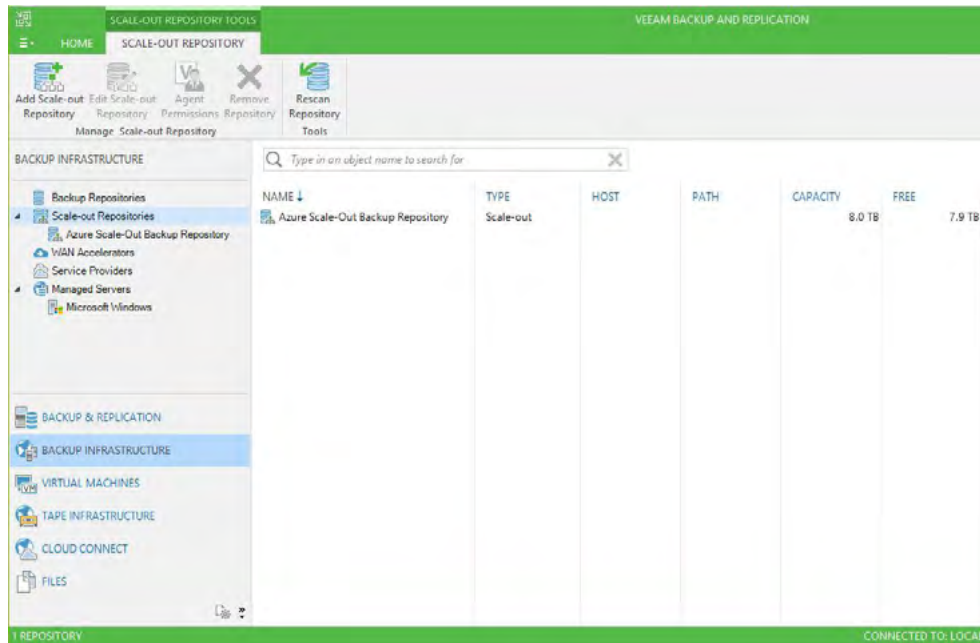


Figure 82: List of available Scale-out Backup Repositories

## Configure virtual networking in Azure

As described in [Figure 41](#), for security reasons, we now need to separate VMs that need to be accessed from the internet from those that don't. We need to create a front-end subnet and a back-end subnet in the same virtual network to avoid VNet to VNet peering<sup>5</sup> — which is not free of charge.

On top of that, we will leverage Azure network security groups to control the access to our Azure VMs. Network security groups are comprised of inbound and outbound rules. In this case, we want the Veeam cloud gateways to be accessible from the internet on TCP port 6180, the Veeam Backup & Replication server on TCP port 3389 for remote management over RDP, and optionally on TCP ports 9392 and 10003 for the Veeam remote console and management.

**Note:** It is not recommended to expose TCP ports 9392 and 10003 to the internet without additional protection such as a VPN, unless source filters are applied to only allow specific source connections. The same rule applies in case of low-bandwidth and/or high-latency links where RDP would perform better.

In this example, the virtual network and subnets configuration will be as follows:

Virtual network:

- Name: **pfg-vcc-rg-vnet**
- Location: **North Europe**
- Address space: **172.16.0.0/16**

<sup>5</sup>Virtual network peering enables you to connect two VNets in the same region through the Azure backbone network. In terms of pricing, there is a nominal charge for ingress and egress traffic. To learn more about VNet peering, visit: <https://docs.microsoft.com/en-us/azure/virtual-network/virtual-network-peering-overview>

#### Subnets:

- Back-end:
  - Name: **Internal**
  - Address range: **172.16.1.0/24**
  - Network security group: **pfg-vcc-internal-nsg**
- Front-end:
  - Name: **External**
  - Address range: **172.16.10.0/24**
  - Network security group: **pfg-vcc-external-nsg**

The virtual network named **pfg-vcc-rg-vnet** has already been created during the deployment of the first virtual machine (Figure 11), along with the back-end subnet named **Internal**. A first network security group called **pfg-vcc-bs-nsg** has also been created at the same time, with ports TCP 3389 and TCP 6180 open inbound (Figure 14), and it is applied to the virtual machine network interface and not to a subnet.

In ARM, network security groups can be either attached to subnets or to virtual network interfaces. To learn more about network security groups and how rules are applied, visit: <https://docs.microsoft.com/en-us/azure/virtual-network/virtual-networks-nsg>

We now have to create a new **External** subnet in the **pfg-vcc-rg-vnet** virtual network, create a new network security group and attach it to the **External** subnet.

Log in to the Azure portal, go to **Virtual networks** and select the virtual network created during the first VM deployment. In this case, it's **pfg-vcc-rg-vnet**. Next, click on **Subnets**.

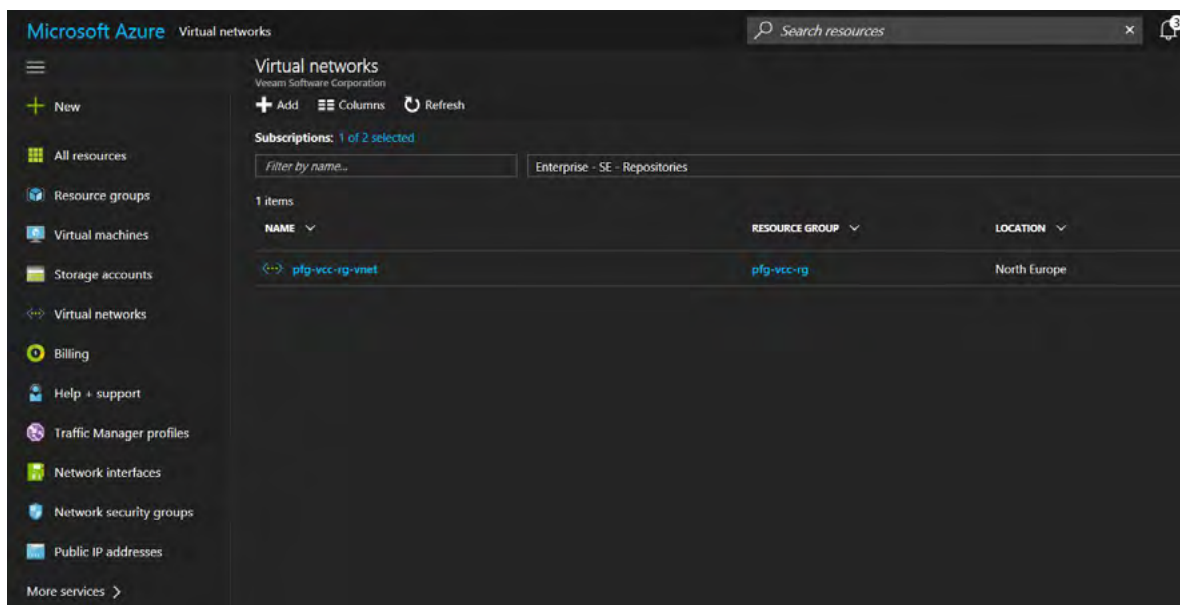


Figure 83: Virtual networks



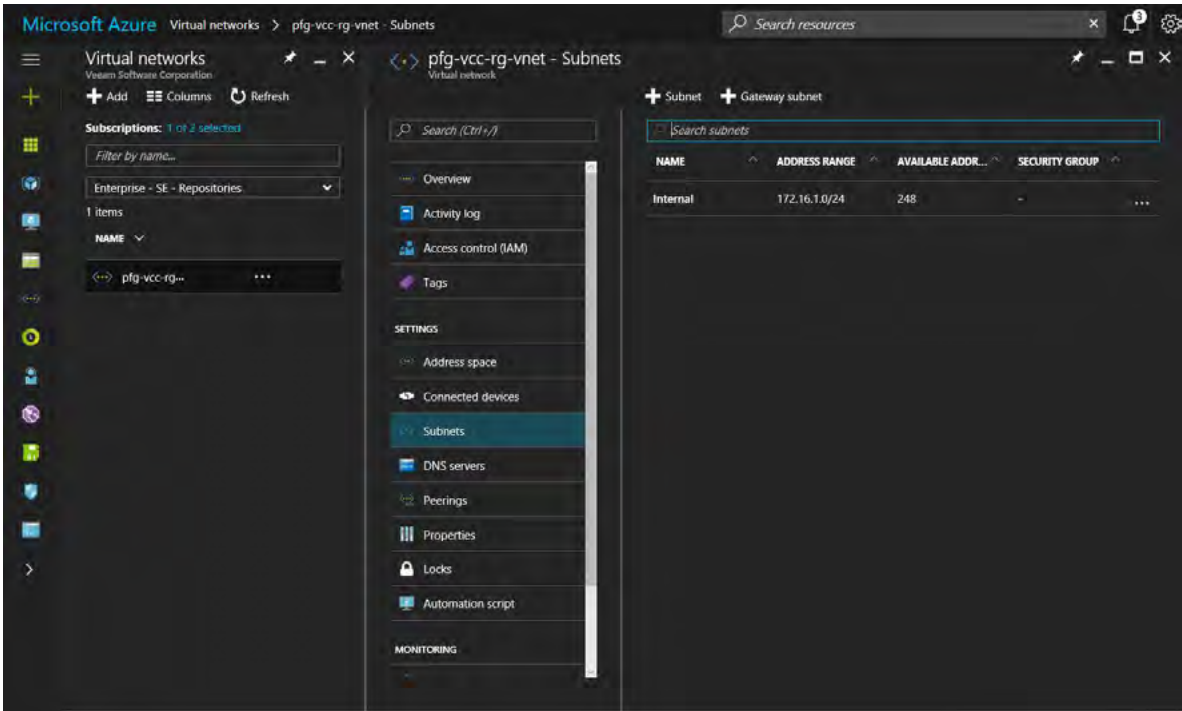


Figure 84: Subnets

We can see the existing **Internal** subnet. Click the **+ Subnet** button. Give a name to the new subnet, and configure the **Address Range**. In this example, it's **External** and **172.16.10.0/24**. Click **OK**.

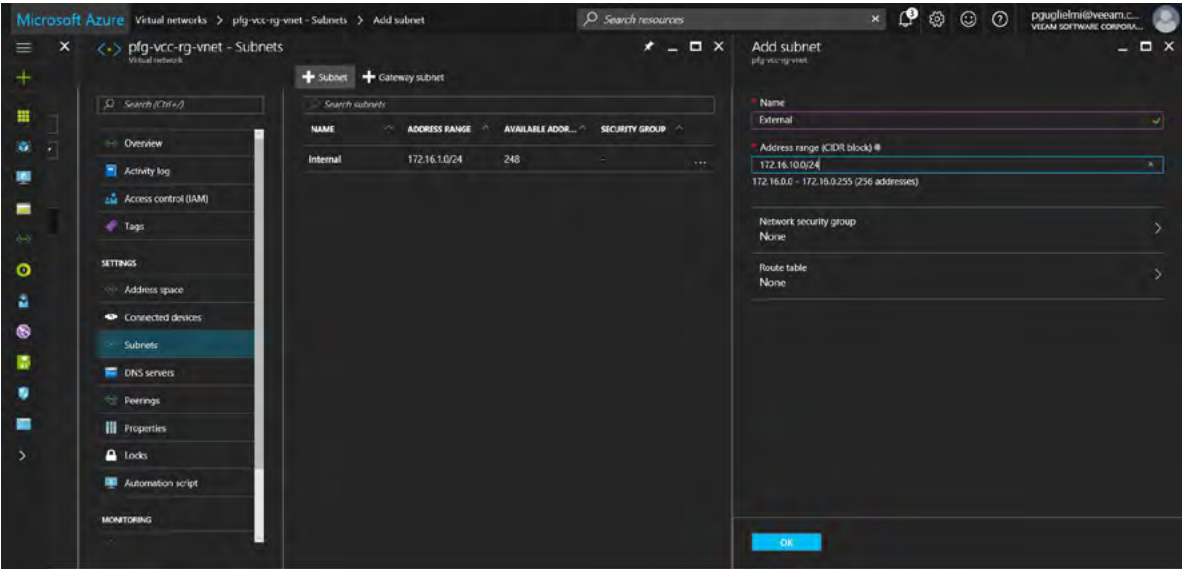


Figure 85: Add subnet

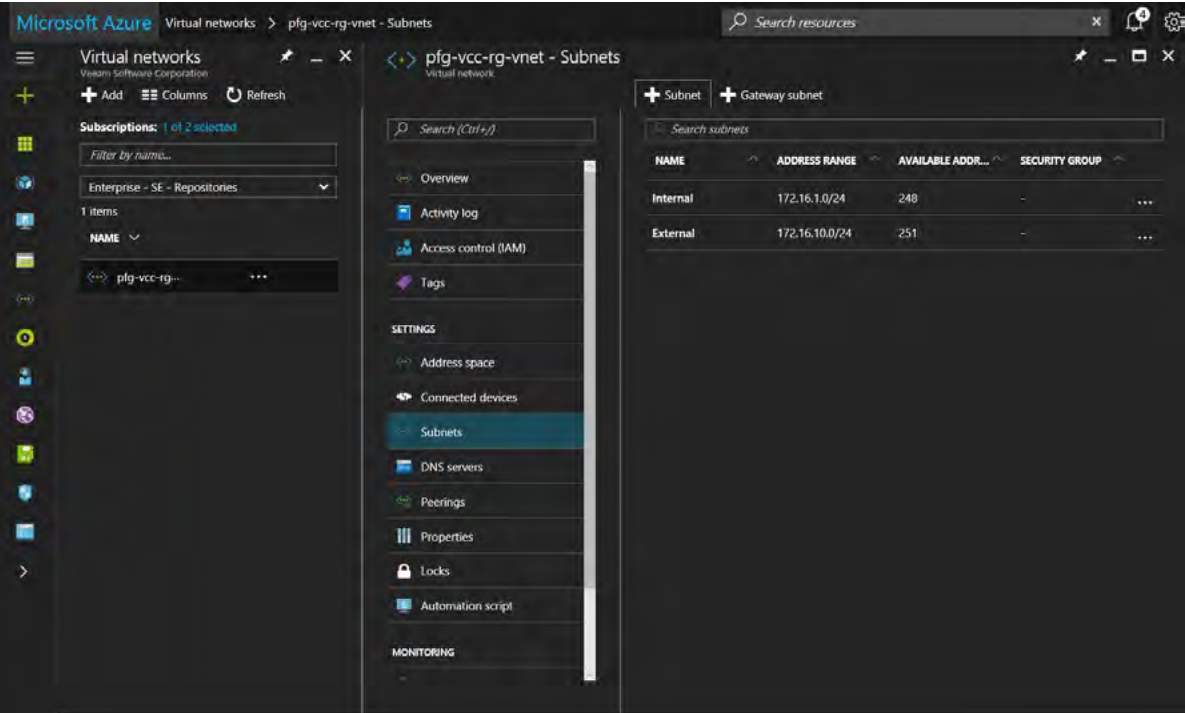


Figure 86 : Successful creation of the new External subnet

Create the new network security group that will be attached to the **External** subnet. In the Azure portal, go to **Network security groups** and click the **+Add** button.




Figure 87: Network security groups

Give a name to the new network security group and choose a subscription, a resource group and a location, and click **Create**.

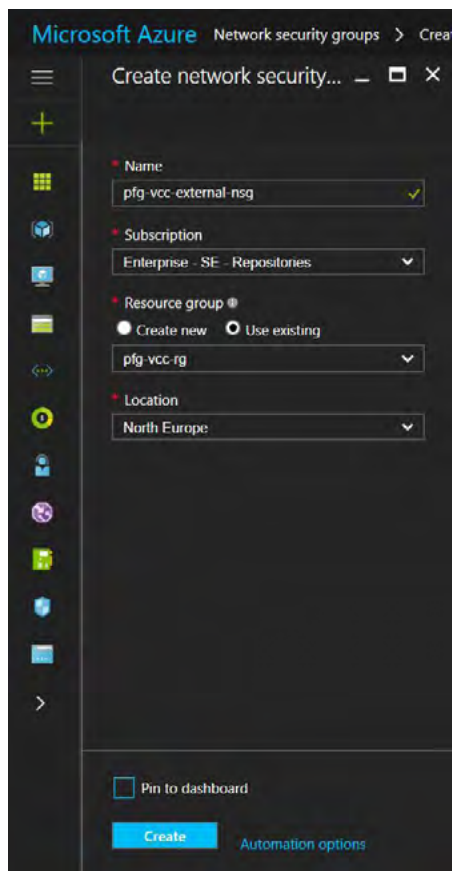


Figure 88: Create network security group

Go to **Inbound security rules** and click the **+Add** button.

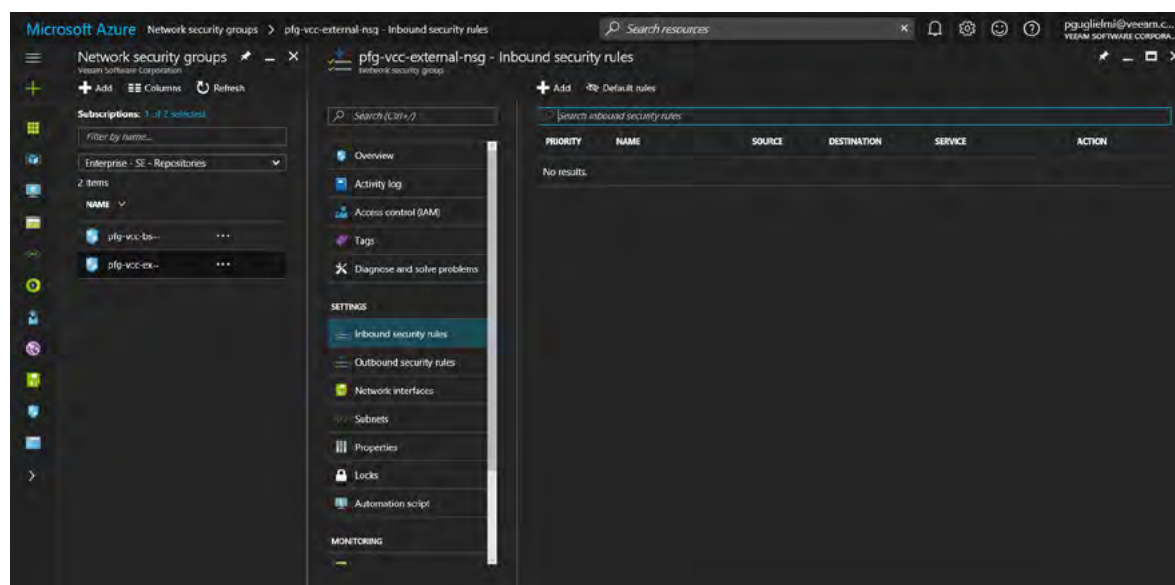


Figure 89: Add inbound security rules

As explained earlier, a rule to allow the incoming Veeam Cloud Connect communications, and a temporary one to allow RDP for the initial configuration of the cloud gateway virtual machines is needed.

Parameters for the new inbound rules:

- Name: **CloudConnect**
- Priority: **1010**
- Source: **Any**
- Service: **Custom**
- Protocol: **TCP**
- Port: **6180**
- Action: **Allow**

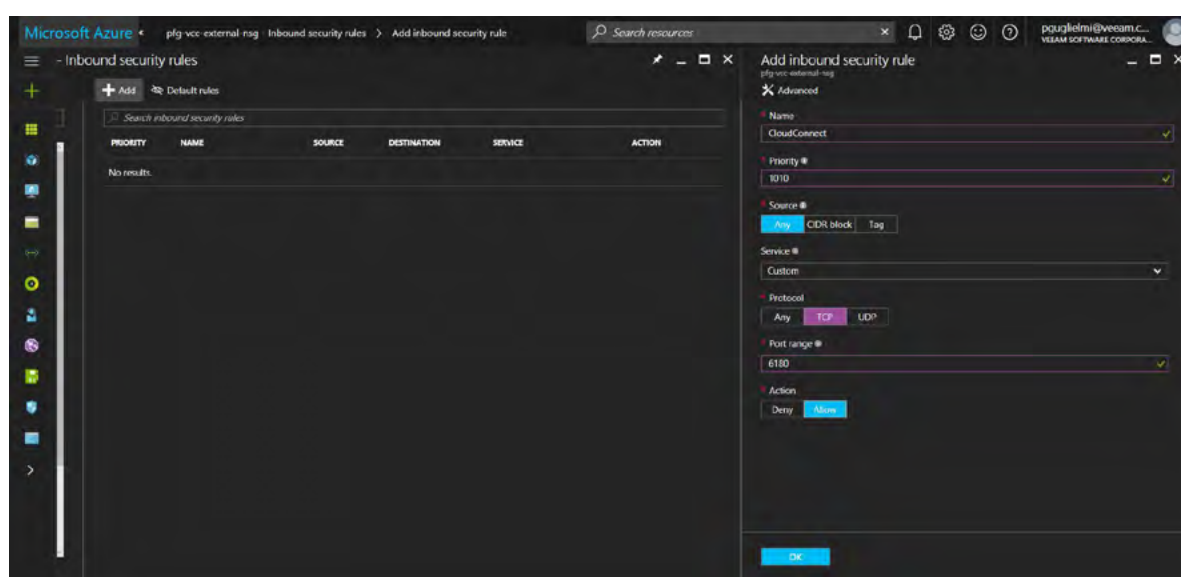


Figure 90: New inbound security rule for Cloud Connect

Repeat the steps above with the following parameters to allow RDP:

- Name: **default-allow-rdp**
- Priority: **1020**
- Source: **Any**
- Service: **RDP**
- Action: **Allow**

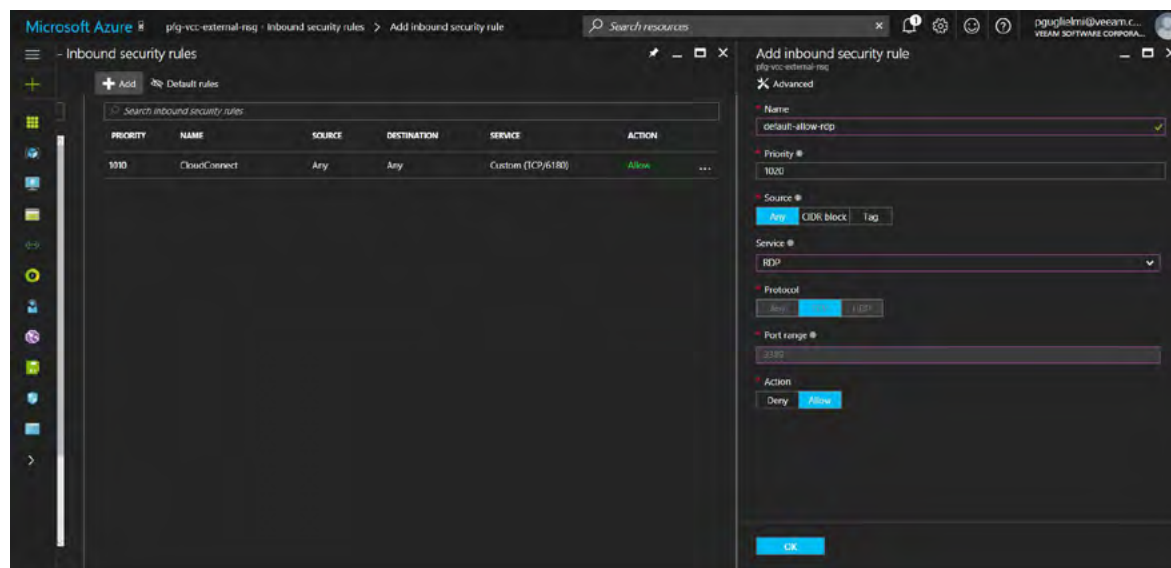


Figure 91: New inbound security rule for RDP

Confirm that you now have the two required inbound rules.

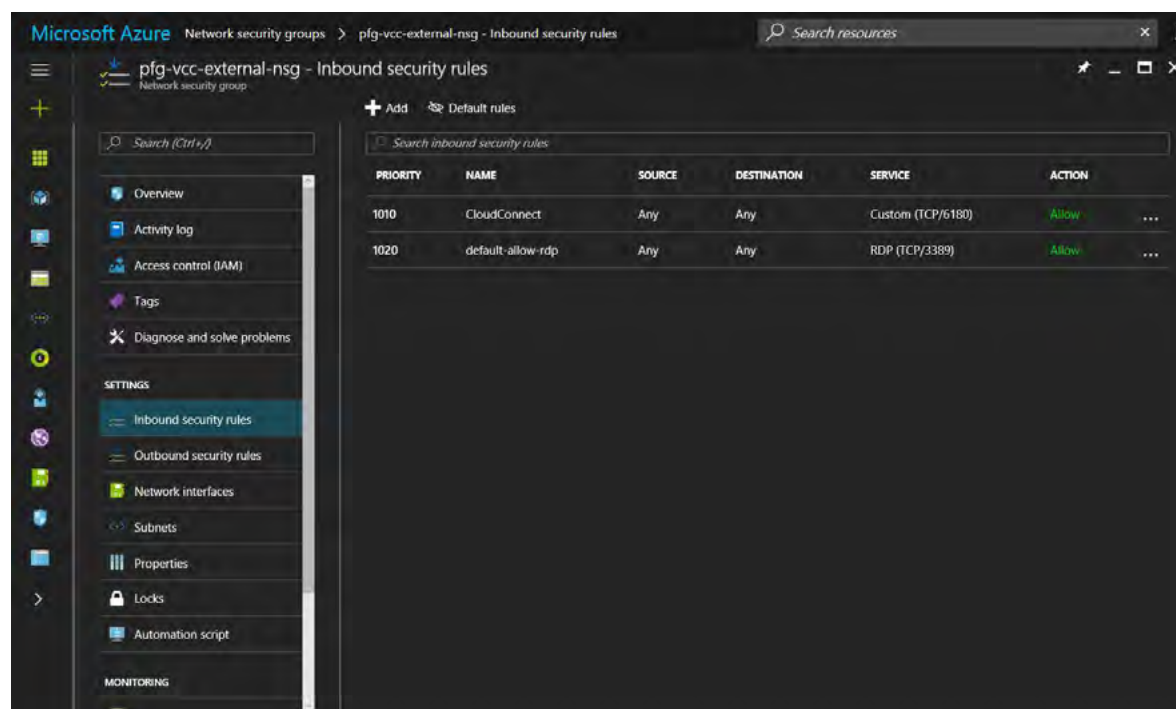


Figure 92: Inbound security rules for CloudConnect and RDP

Now, attach this network security group to the **External** subnet. In the Azure portal, go to **Virtual networks > Subnets**, select the **External** subnet, then click on **Network security group** and select the one created previously. In this case, it's called **pfg-vcc-external-nsg**. Finally, click on the **Save** button.

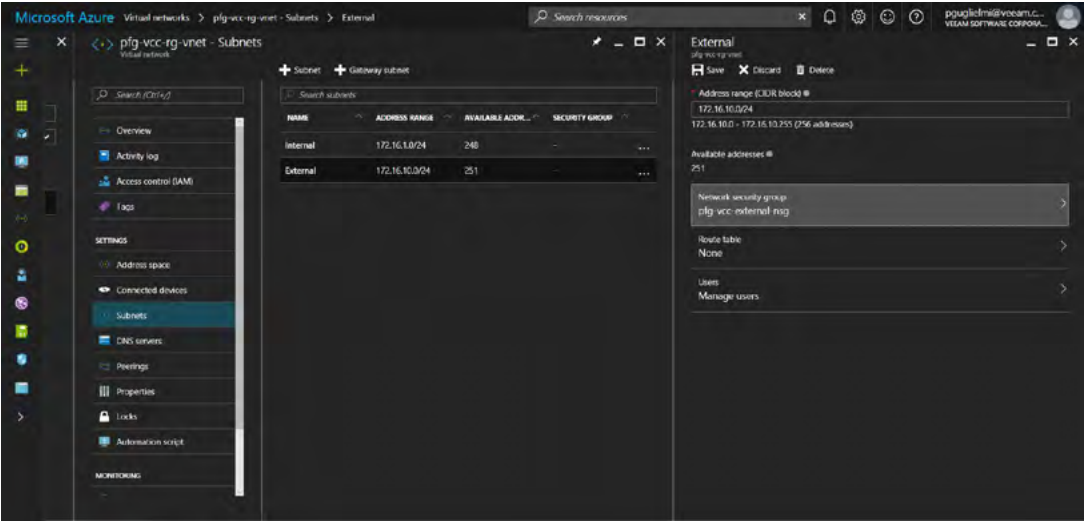


Figure 93: Attach a network security group to a subnet

To complete the network configuration in Azure, apply the right access control rules to the VMs in the **Internal** perimeter. In fact, accessing the Veeam Backup & Replication server, either with RDP or the Veeam console and not directly other VMs like repository servers or WAN accelerators, is needed. Do not apply a network security group to the entire internal subnet, but rather to the Veeam backup server NIC. Utilize the previously created network security group that was deployed from Azure Marketplace (**pfg-vcc-bs-nsg**). However, a few minor reconfiguration steps are required.

In the Azure portal, go to your network security group (**pfg-vcc-bs-nsg**), and click **Inbound security rules**.

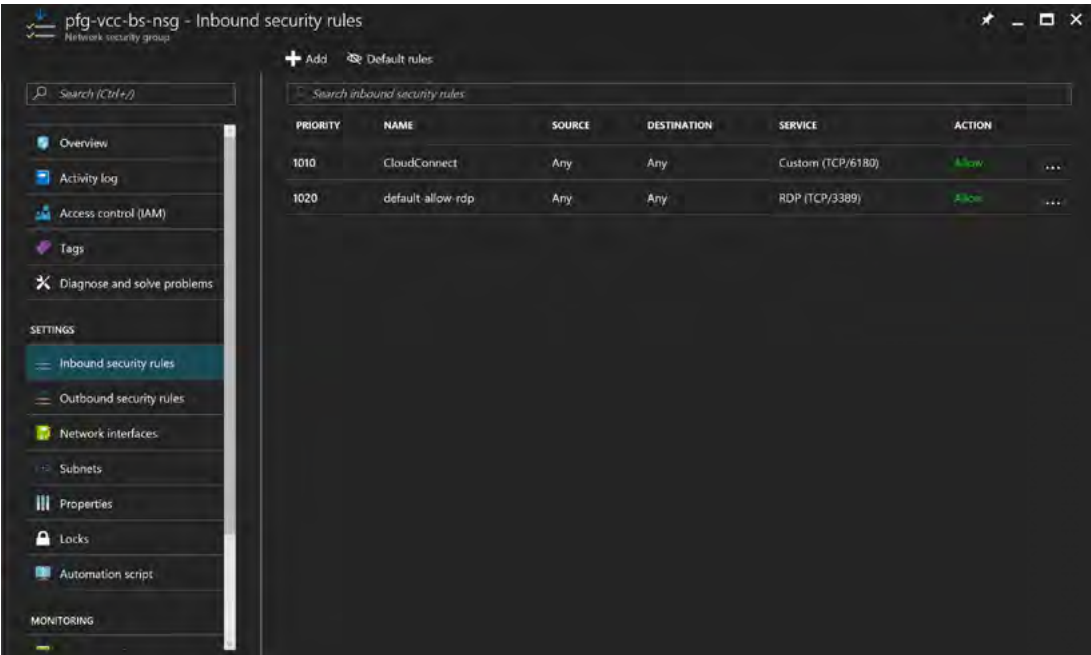


Figure 94: Existing inbound security rules



The **CloudConnect** rule on the network security group for the Veeam Backup & Replication server is no longer required. Click on it and then on the **Delete** button. Confirm by clicking **Yes**.

The screenshot shows the 'CloudConnect' rule configuration window in Azure. The window title is 'CloudConnect' with a subtitle 'pfg-vee-rg'. The toolbar includes 'Save', 'Discard', 'Delete', and 'Advanced' buttons. The configuration fields are as follows:

- Name:** CloudConnect
- Priority:** 1010
- Source:** Any (selected), CIDR block, Tag
- Service:** Custom (selected)
- Protocol:** Any, TCP (selected), UDP
- Port range:** 6180
- Action:** Deny, Allow (selected)

Figure 95: Edit or delete a rule

**Note:** It is not recommended to expose TCP ports 9392 and 10003 to the internet without additional protection, such as a VPN, unless source filters are applied to only allow specific source connections. The same rule applies in case of low-bandwidth and/or high-latency links where RDP would perform better.

Now, create a new rule to allow remote Veeam consoles to connect from the internet. Click **Add** and use the following parameters to create the rule, then click **OK**:

- Name: **VeeamConsole**
- Priority: **1010**
- Source: **Any**
- Service: **Custom**
- Protocol: **TCP**
- Port: **9392**
- Action: **Allow**

The screenshot shows a Windows-style dialog box titled "Add inbound security rule" with a subtitle "pfq-vcc-bs-nsg". It has a close button (X) in the top right corner. Below the title bar, there is a tab labeled "Advanced". The dialog contains several fields with red asterisks indicating required fields:

- Name:** A text box containing "VeeamConsole" with a green checkmark on the right.
- Priority:** A text box containing "1010" with a green checkmark on the right.
- Source:** A group box containing three buttons: "Any" (highlighted in blue), "CIDR block", and "Tag".
- Service:** A dropdown menu showing "Custom" with a downward arrow.
- Protocol:** A group box containing three buttons: "Any", "TCP" (highlighted in blue), and "UDP".
- Port range:** A text box containing "9392" with a green checkmark on the right.
- Action:** A group box containing two buttons: "Deny" and "Allow" (highlighted in blue).

At the bottom left, there is a blue button labeled "OK".

Figure 96: Add an inbound security rule for the Veeam console

And finally, create the last inbound security rule to allow some management tasks, such as adding tenants or Veeam components, when connected remotely with the Veeam console. Click **Add** and use the following parameters to create the rule, then click **OK**:

- Name: **RemoteManagement**
- Priority: **1030**
- Source: **Any**
- Service: **Custom**
- Protocol: **TCP**
- Port: **10003**
- Action: **Allow**

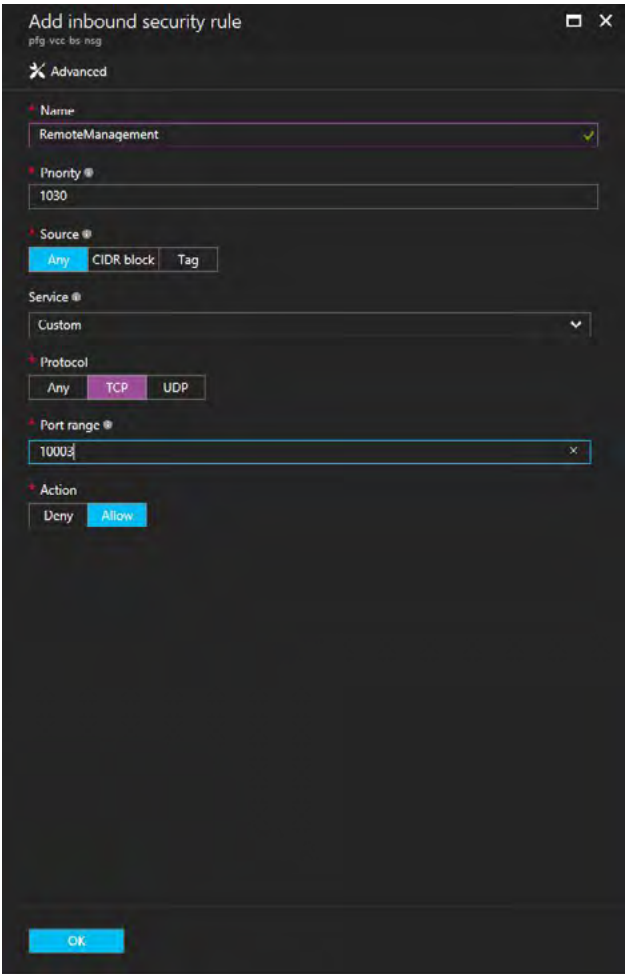


Figure 97: Add an inbound security rule for remote management tasks

+ Add 102 Default rules						
Search inbound security rules						
PRIORITY	NAME	SOURCE	DESTINATION	SERVICE	ACTION	
1010	VeeamConsole	Any	Any	Custom (TCP/9392)	Allow	...
1020	default-allow-rdp	Any	Any	RDP (TCP/3389)	Allow	...
1030	RemoteManagement	Any	Any	Custom (TCP/10003)	Allow	...

Figure 98: Inbound security rules configured

## Deploy and configure additional cloud gateways

Now that the virtual network is configured properly both for back-end and front-end, deploy the additional Veeam cloud gateways in the front-end perimeter, as they are the components in charge of allowing secure communication between the Veeam Cloud Connect for the Enterprise infrastructure running within Azure and the tenants' ones thanks to an SSL tunnel.

Additionally, we'll cover how to implement High Availability and load balancing for the cloud gateway VMs.

Start by deploying a few new Windows VMs. Having two cloud gateway VMs is the minimum to ensure Availability of the service. A minimum of two cloud gateways is considered to be a Veeam best practice. However, it's best to deploy three. This allows for an N+1 configuration and always provide complete redundancy. For the purpose of this white paper, we'll deploy two.

To deploy the cloud gateway, repeat the steps illustrated within [Deploy and configure additional repositories](#). The main differences here are that resources requirements for cloud gateways are rather low. As stated in the [Veeam Cloud Connect Guide](#), a single connection from a tenant consumes around 512 KB of memory, which means that 1 GB of memory allows to receive up to 2,000 concurrent connections! For this reason, we'll choose a smaller Azure VM size for cloud gateway servers, which is A1\_V2 with one core and 2 GB of memory. Secondly, as we want to make the service highly available, we'll leverage a native Azure service called Availability Set. Putting virtual machines in an Availability Set ensures that the service they deliver will always be available in case of maintenance events, whether they are planned such as updates, or unplanned such as hardware failures.

To ensure Availability during planned maintenance, Azure uses update domains to make sure the VMs in the same Availability Set are not updated all at the same time.

Similarly, Azure uses fault domains to make sure VMs of the same Availability Set are not all on the same underlying hardware, preventing the hosted services to become completely unavailable during unplanned maintenance events.

Due to the fact that adding a VM in an Availability Set is an action that can be done only with the help of PowerShell and that it requires a reboot, we choose to create the Availability Set and to add VMs in it while deploying them through the Azure portal.

To learn more about Availability Sets, update domains and fault domains, visit: <https://docs.microsoft.com/en-us/azure/virtual-machines/virtual-machines-windows-manage-availability> and <https://docs.microsoft.com/en-us/azure/virtual-machines/virtual-machines-windows-create-availability-set>

In the Azure portal, search for Windows Server. Again, for the purposes of this document, we choose Windows Server 2016, but Veeam Backup & Replication also supports Windows Server 2012 R2.

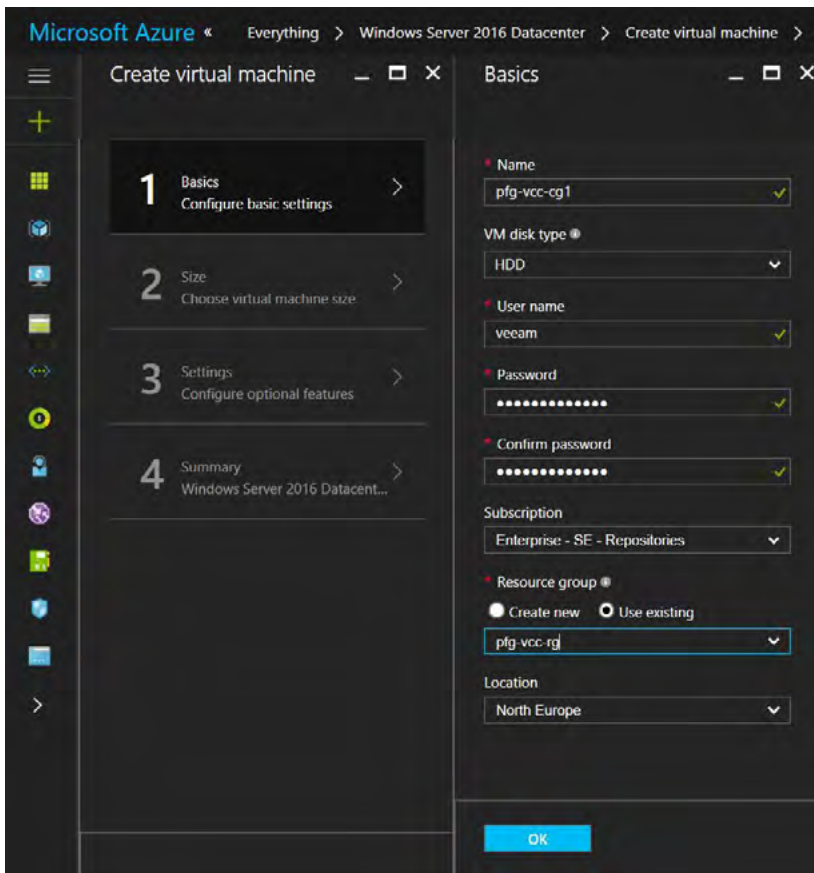
The image shows the 'Create virtual machine' wizard in the Microsoft Azure portal. The breadcrumb navigation at the top indicates the path: 'Everything > Windows Server 2016 Datacenter > Create virtual machine > Basics'. The left sidebar contains icons for various Azure services. The main content area is divided into two panels. The left panel shows a four-step progress bar: 1. Basics (Configure basic settings), 2. Size (Choose virtual machine size), 3. Settings (Configure optional features), and 4. Summary (Windows Server 2016 Datacent...). The right panel, titled 'Basics', contains the following fields: 'Name' (pfg-vcc-eg1), 'VM disk type' (HDD), 'User name' (veeam), 'Password' (masked with dots), 'Confirm password' (masked with dots), 'Subscription' (Enterprise - SE - Repositories), 'Resource group' (pfg-vcc-rg), and 'Location' (North Europe). At the bottom right of the 'Basics' panel is an 'OK' button.

Figure 99: Basic settings for cloud gateway VM creation

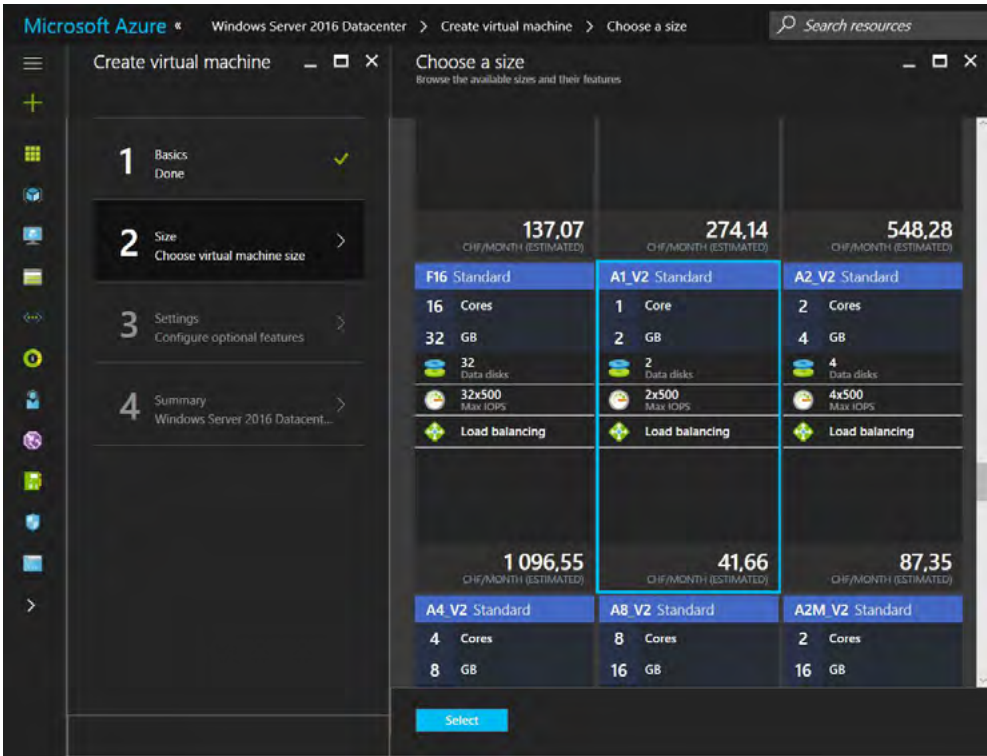


Figure 100: Choose a VM size for cloud gateways servers

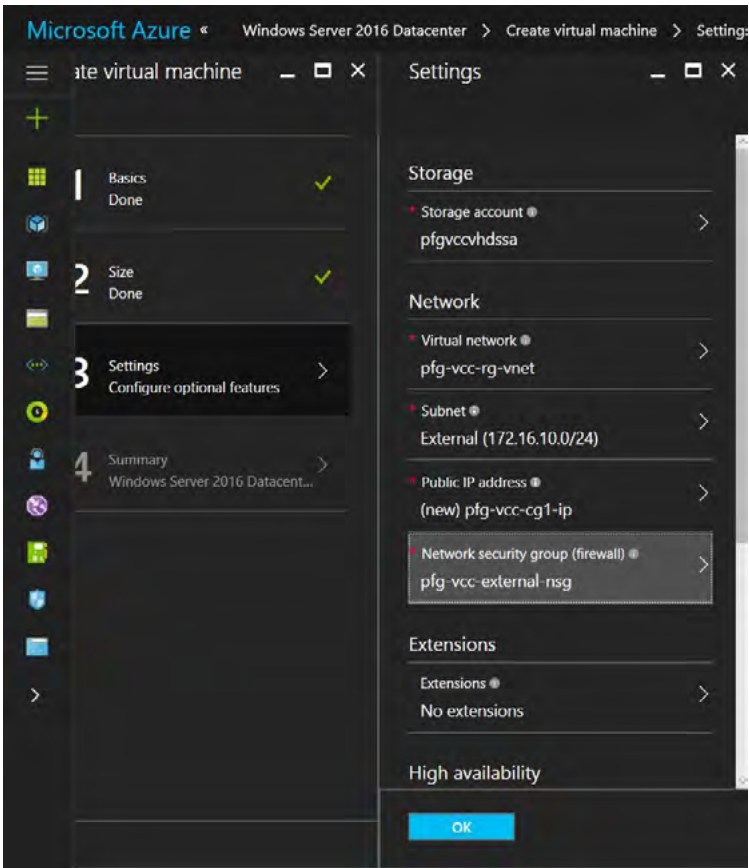


Figure 101: Choose the front-end subnet and network security group

Under **High availability**, click **Availability set — None**.

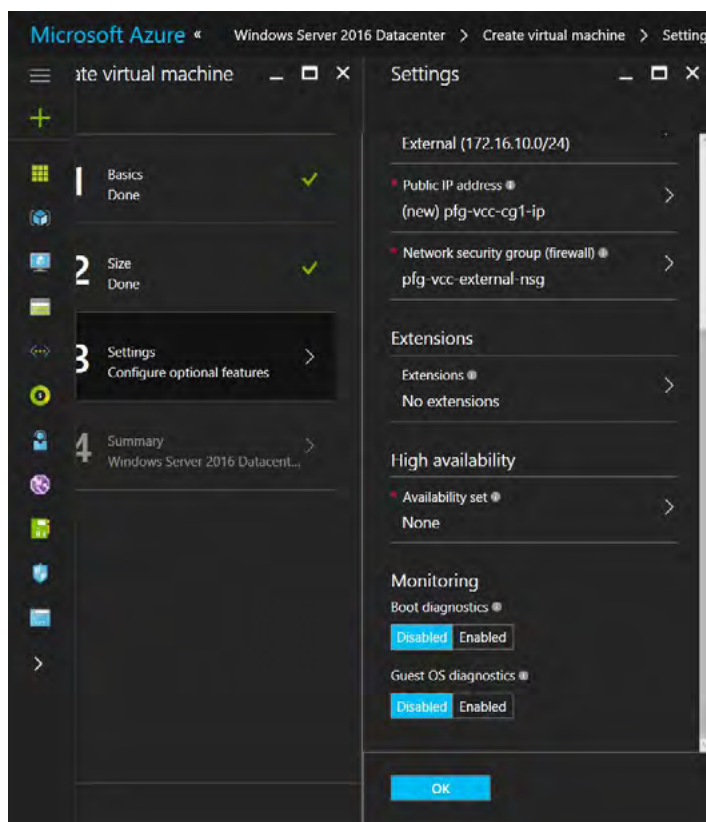


Figure 102: Click on Availability Set

On the new VM blade, click on **Create new**, give a name to the new Availability Set and choose the number of update domains and fault domains. In this case, it will be called **pfg-vcc-cg-as**, and we accept the default settings of five update domains and three fault domains.

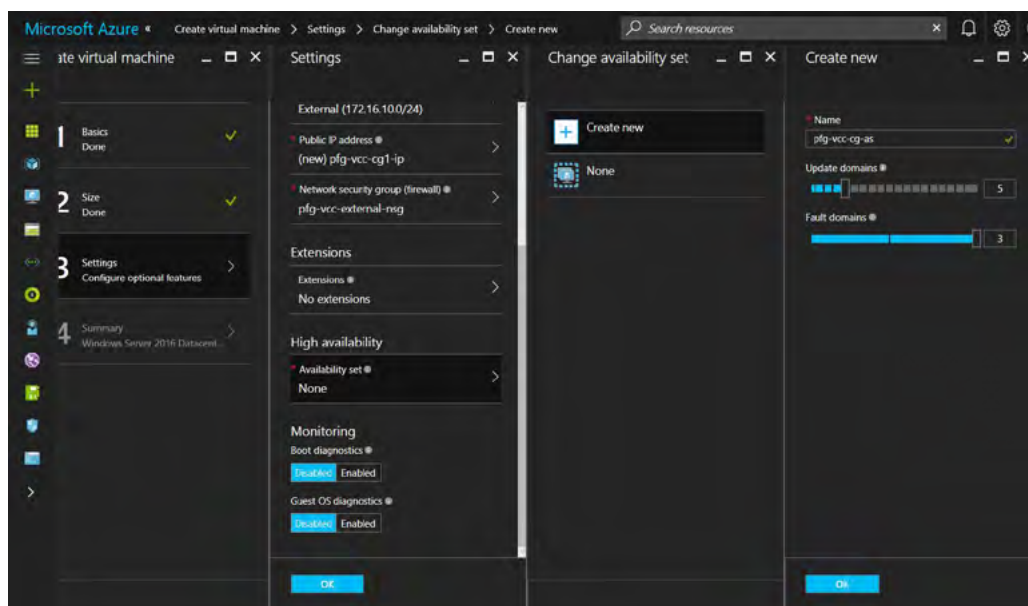


Figure 103: Create new Availability Set



Review the settings and click **OK**. After a few minutes, the new VM will be up and running.

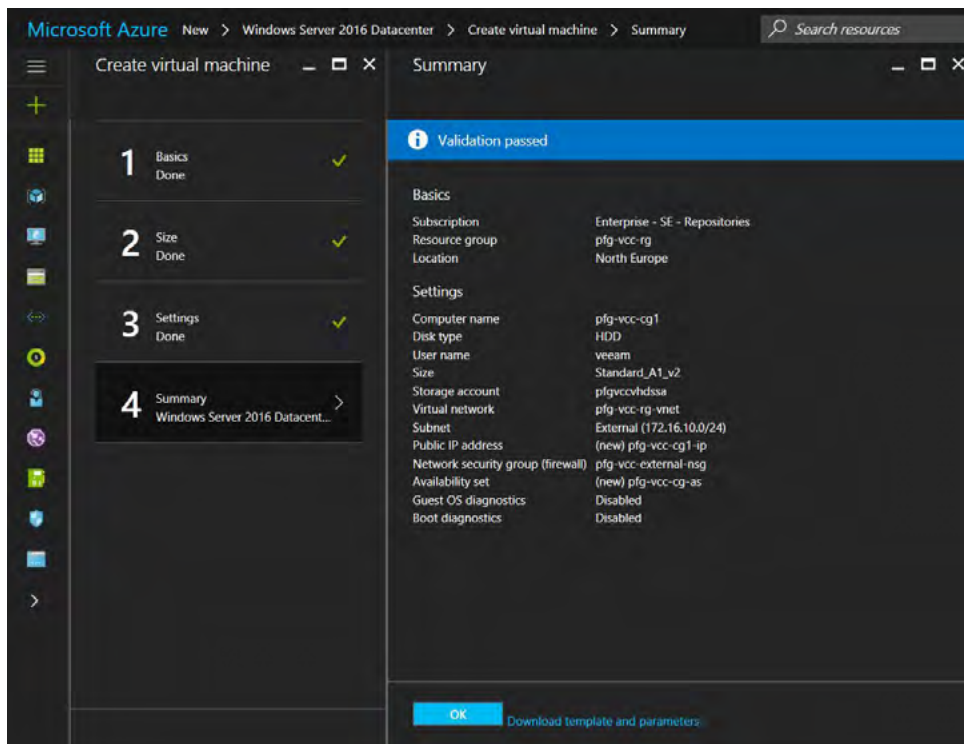


Figure 104: Create virtual machine summary

Like we did for the previous VMs, we also need for the cloud gateway servers to set a DNS name label because IP addresses are allocated dynamically. Moreover, we'll rely on DNS name to configure a unique entry point later on.

In the Azure portal, select the first cloud gateway VM and in **Overview**, click on **None** under **DNS name label**, set a name and click on the **Save** button. In this case the name is **pfq-vcc-cg1**.

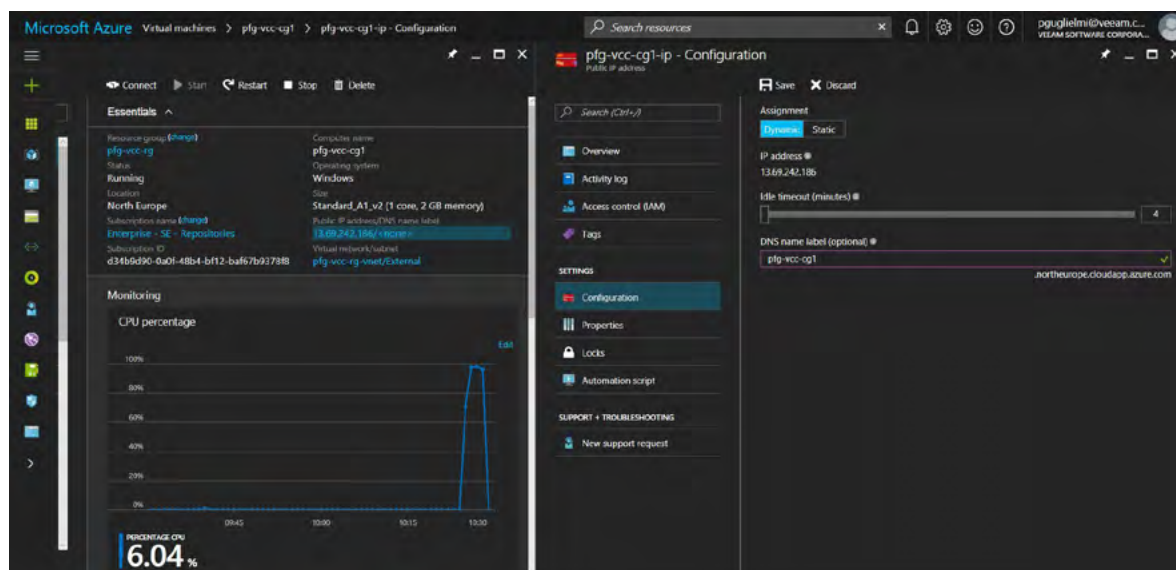


Figure 105: Set a DNS name for the cloud gateway VM

We can continue and deploy the second cloud gateway VM with similar settings. The difference compared to the first one is that we don't need to create a new Availability Set. Instead, we can select the one created previously.

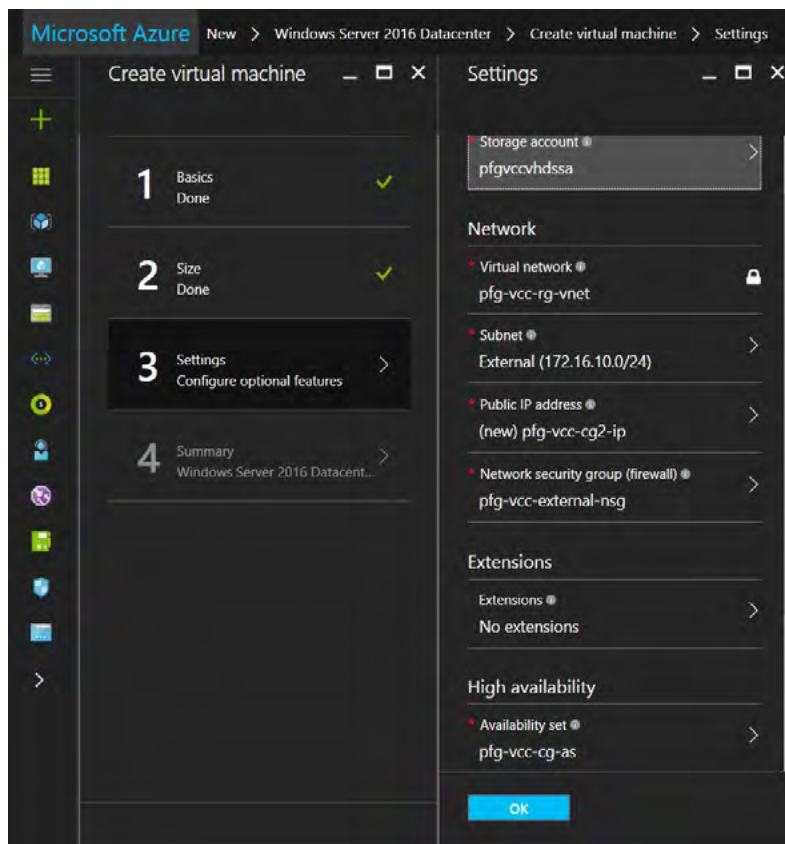


Figure 106: Settings for the second cloud gateway VM

Complete the **Create virtual machine wizard** and wait for the VM to be ready. Repeat the process as many times as required.

The cloud gateway servers are almost ready to be configured and used as actual Veeam cloud gateways. The last thing we need to configure is some sort of load balancing ahead of them. Different services exist in Azure to implement load balancing depending on requirements and on what level you need it to work. The available options are:

- **Azure Load Balancer:** Works at the transport layer (Layer 4 of OSI model)
- **Application Gateway:** Works at the application layer (Layer 7 of OSI model)
- **Traffic Manager:** Works at the DNS level

To make the right decision, it is important to understand that the Veeam cloud gateways have their own native load balancing. It means that a real load balancer before them is not necessary. Optionally, a unique entry point to the cloud gateways pool can be created. This is achieved by working at the DNS level, which is why **Traffic Manager** will be a good candidate.

To learn more about Azure load balancing services, see: <https://docs.microsoft.com/en-us/azure/load-balancer/load-balancer-overview>

To learn more about Veeam cloud gateways load balancing, read the **Load balancers** section of the Veeam Cloud Connect reference architecture guide here: <https://www.veeam.com/wp-cloud-connect-reference-architecture-v9.html>

Now let's go ahead and create a Traffic Manager profile. In the Azure portal, go to **Traffic Manager profiles** and click **Add**.

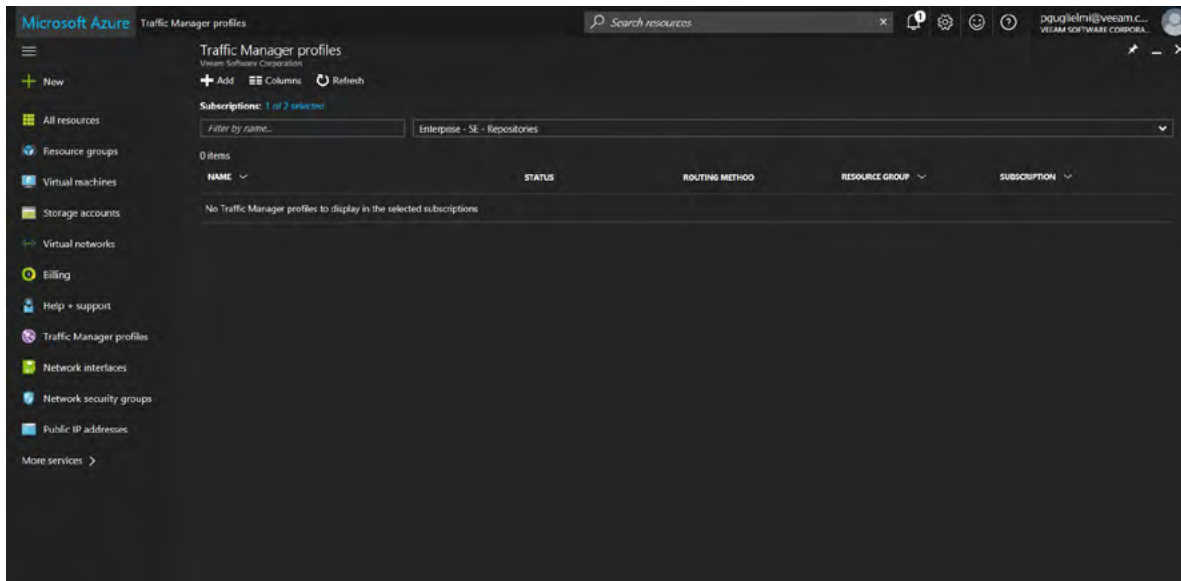


Figure 107: Traffic Manager profiles

Give a name to the new Traffic Manager profile, choose a subscription and a resource group. You also have to choose a **Routing method**. There are four different ones:

- Priority
- Weighted
- Performance
- Geographic

The second one distributes traffic across a set of endpoints, either evenly or according to weights, and thus is the one that is the closest to DNS round robin, which is exactly what we need!

To learn more about Azure Traffic Manager traffic-routing methods, go to: <https://docs.microsoft.com/en-us/azure/traffic-manager/traffic-manager-routing-methods>

To create our Traffic Manager profile, we use the following parameters, then click **Create**:

- Name: **veeamcloudconnect**
- Routing method: **Weighted**
- Resource group: **Use existing — pfg-vcc-rg**

The screenshot shows the 'Create Traffic Manager profile' form in the Microsoft Azure portal. The sidebar on the left contains navigation links for various Azure services. The main form area is titled 'Create Traffic Manager...' and contains the following fields:

- Name:** veeamcloudconnect (with a checkmark icon indicating it's valid)
- Routing method:** Weighted (selected from a dropdown menu)
- Subscription:** Enterprise - SE - Repositories (selected from a dropdown menu)
- Resource group:** pfg-vcc-rg (selected from a dropdown menu, with radio buttons for 'Create new' and 'Use existing')
- Resource group location:** North Europe (selected from a dropdown menu)

At the bottom of the form, there is a checkbox for 'Pin to dashboard' and a blue 'Create' button. A link for 'Automation options' is also visible.

Figure 108: Create Traffic Manager profile

As you can see in [Figure 108](#), the fully qualified domain name of the Traffic Manager profile is **veeamcloudconnect.trafficmanager.net**. It uses the default Azure Traffic Manager domain name but you might want to use your own. To achieve this, you can simply add a **CNAME** (alias) record at your domain registrar to map a name using your own domain name to the one in .trafficmanager.net.

Once created, select the **veeamcloudconnect** Traffic Manager profile and click **Endpoints**. Endpoints are the services for which the traffic manager will balance connections. In our case, the endpoints are the cloud gateway VMs.

Click **Add** to add the first endpoint.

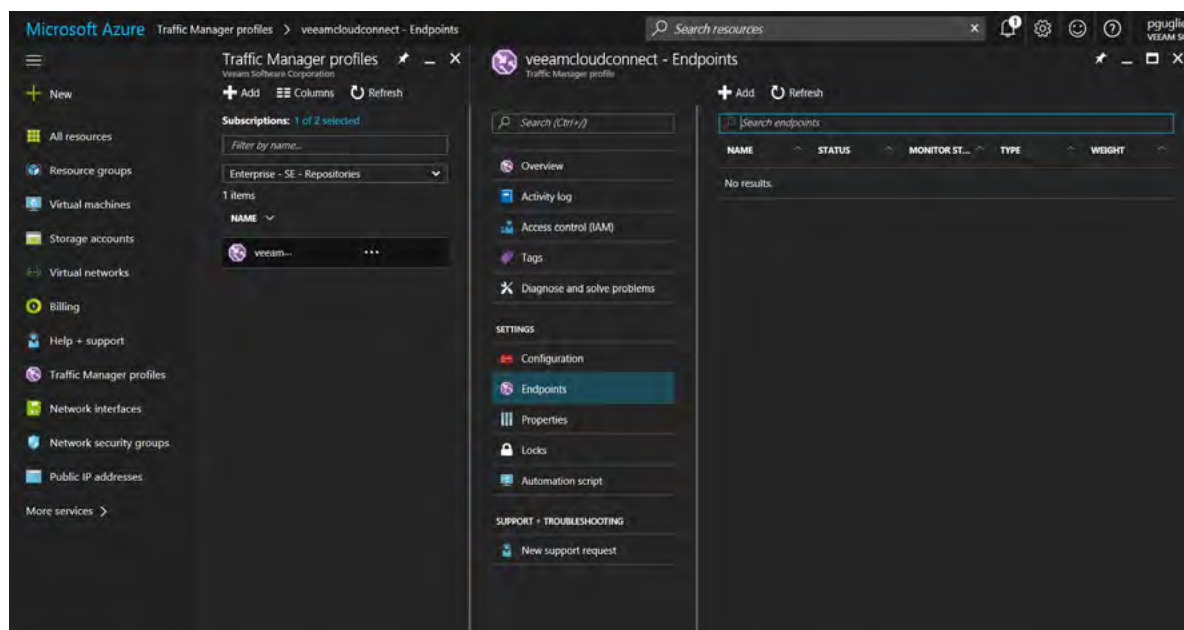


Figure 109: Add endpoints

We'll use the following parameters for the first endpoint:

- Type: **Azure endpoint**
- Name: **pfg-vcc-cg1**
- Target resource type: **Public IP address**
- Target resource: **pfg-vcc-cg1-ip**
- Weight: **1**

Figure 110: Add traffic manager endpoint

Repeat the steps above to add all your cloud gateway VMs as endpoints in the Traffic Manager profile. In this case, we have two endpoints in the Traffic Manager profile.

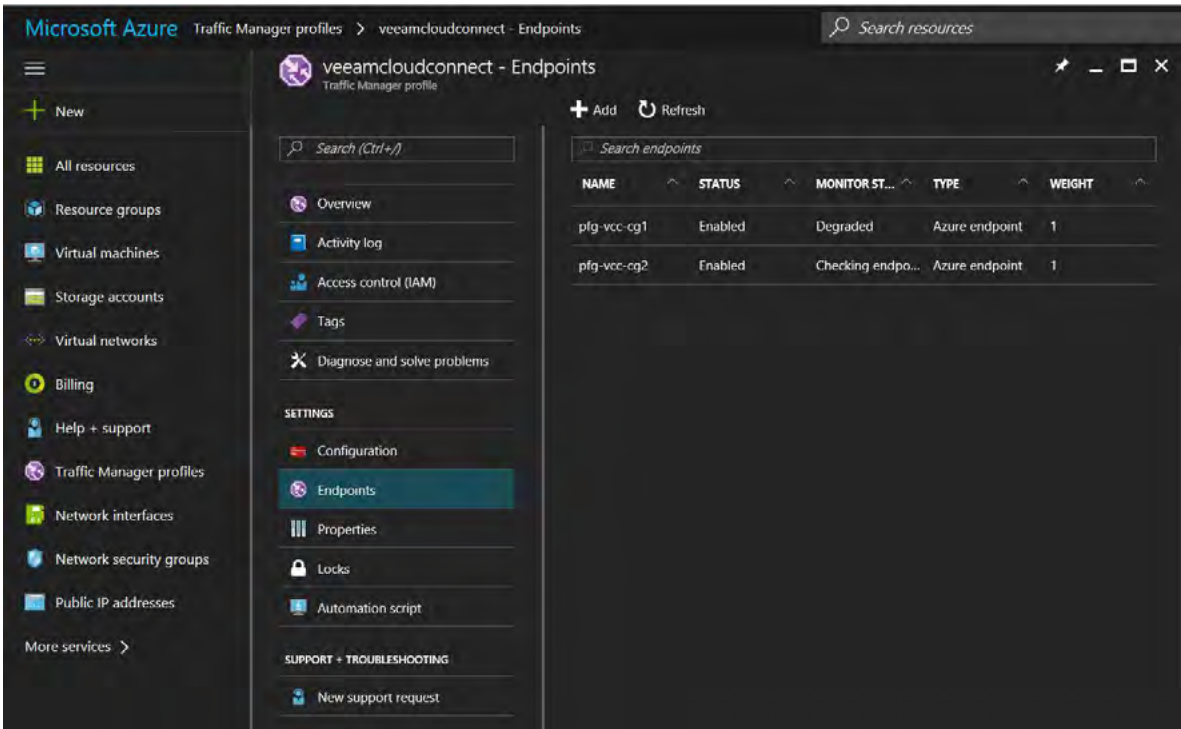


Figure 111: Traffic Manager profile with two endpoints configured

Go back to the Veeam Backup & Replication console to add these two VMs and make them actual Veeam cloud gateways. Connect to the Veeam Backup & Replication server either with RDP or the Veeam console directly using the public DNS name, and go to **Cloud Connect > Cloud gateways**.

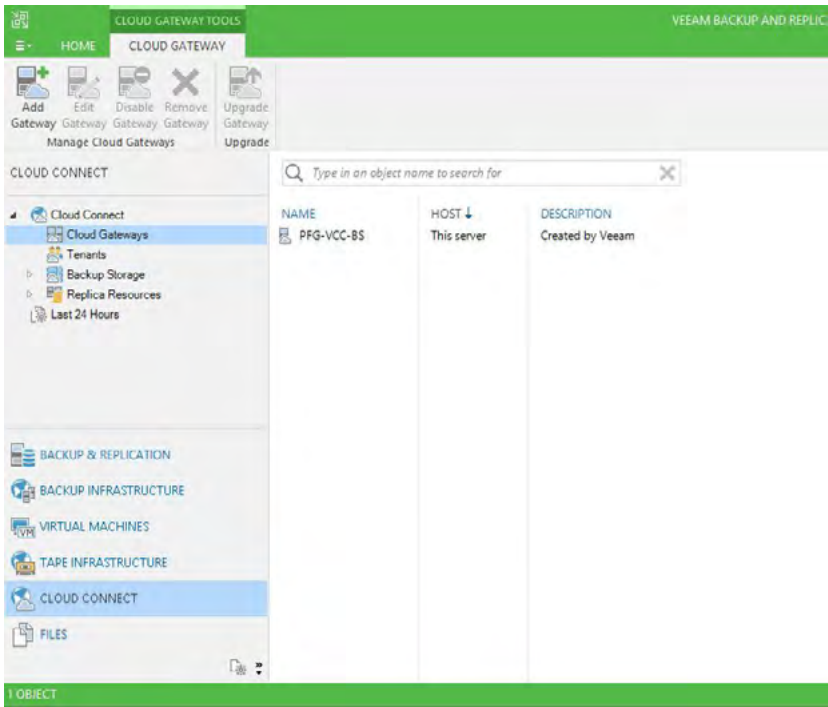


Figure 112: List of available cloud gateways



As you can see in [Figure 112](#), there is already one cloud gateway deployed and available, which is running on the Veeam Backup & Replication server itself. This is because we've deployed the Veeam Cloud Connect backup server from the all-in-one preconfigured template available in Azure Marketplace. We need to add the two additional VMs we've prepared, and afterward we'll be able to remove the already existing one.

Click on the **Add Gateway** button.

Figure 113: New cloud gateway wizard

Since we want to make a separate VM a new cloud gateway, click on **Add New**. Enter the name of the VM you want to add – in this example, the first one is called **pfg-vcc-cg1** – type in an optional description and click **Next**.

Figure 114: Add a new Windows Server

Next, enter credentials with a user that have administrative permissions locally on the VM being added. You can either select credentials in the drop-down list if they've already been added, or click **Add**. In this case, we can select existing credentials in the drop-down list, then click **Next** twice.

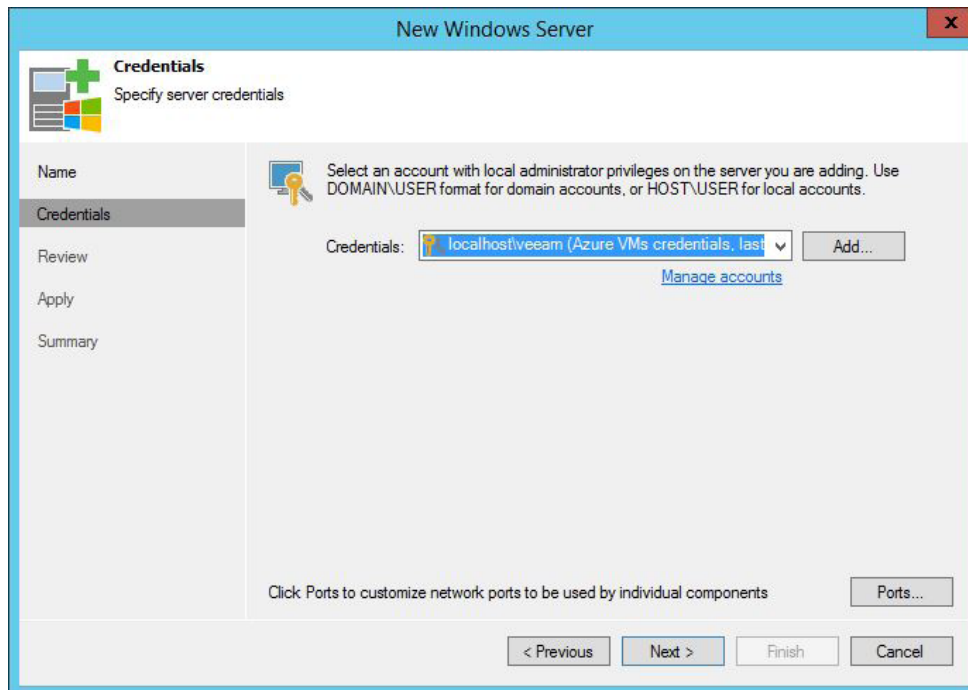


Figure 115: Select or enter credentials

The VM is being registered as a new Windows server in the Veeam infrastructure.

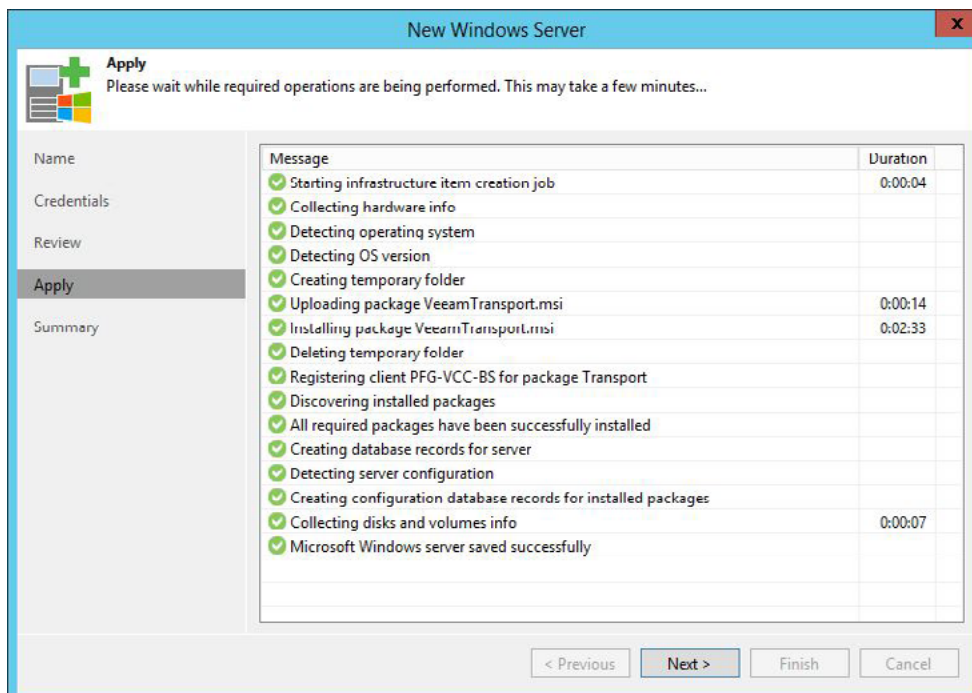


Figure 116: New Windows server being registered

Review the summary and click **Finish**.

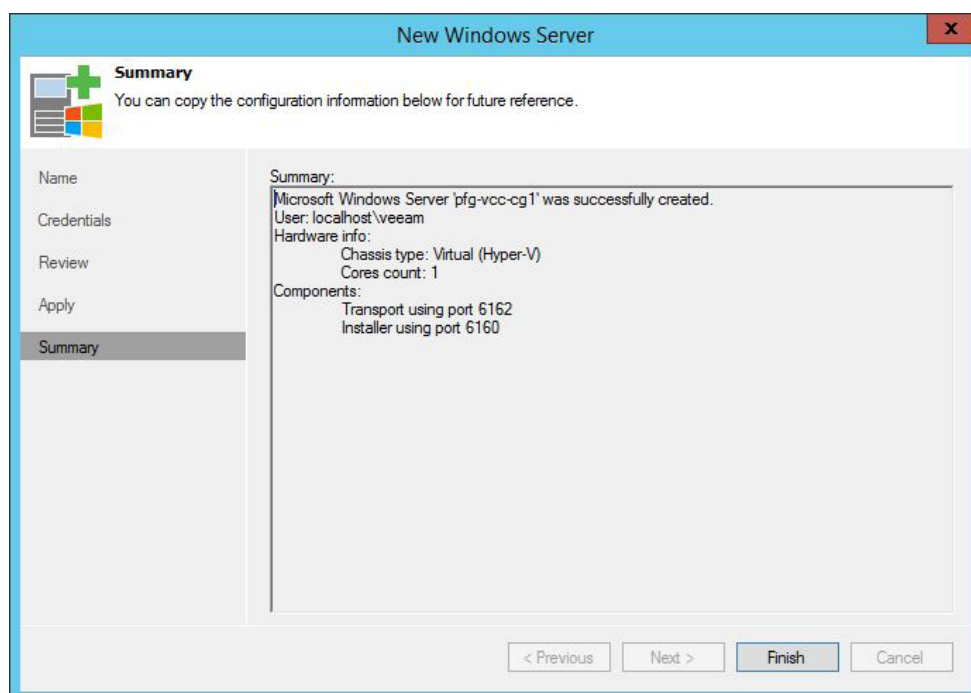


Figure 117: New Windows Server wizard summary

Back to the **Add Cloud Gateway** wizard: Enter an optional description and choose the external port for Veeam Cloud Connect communications that tenants will need to configure on their side. TCP port **6180** is used by default. Click **Next**.

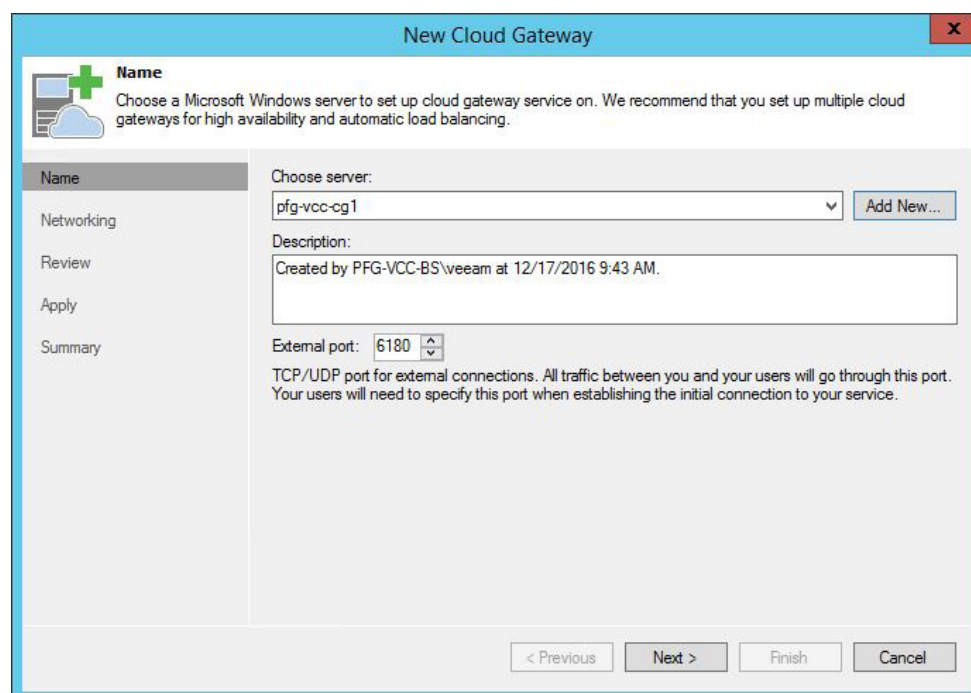
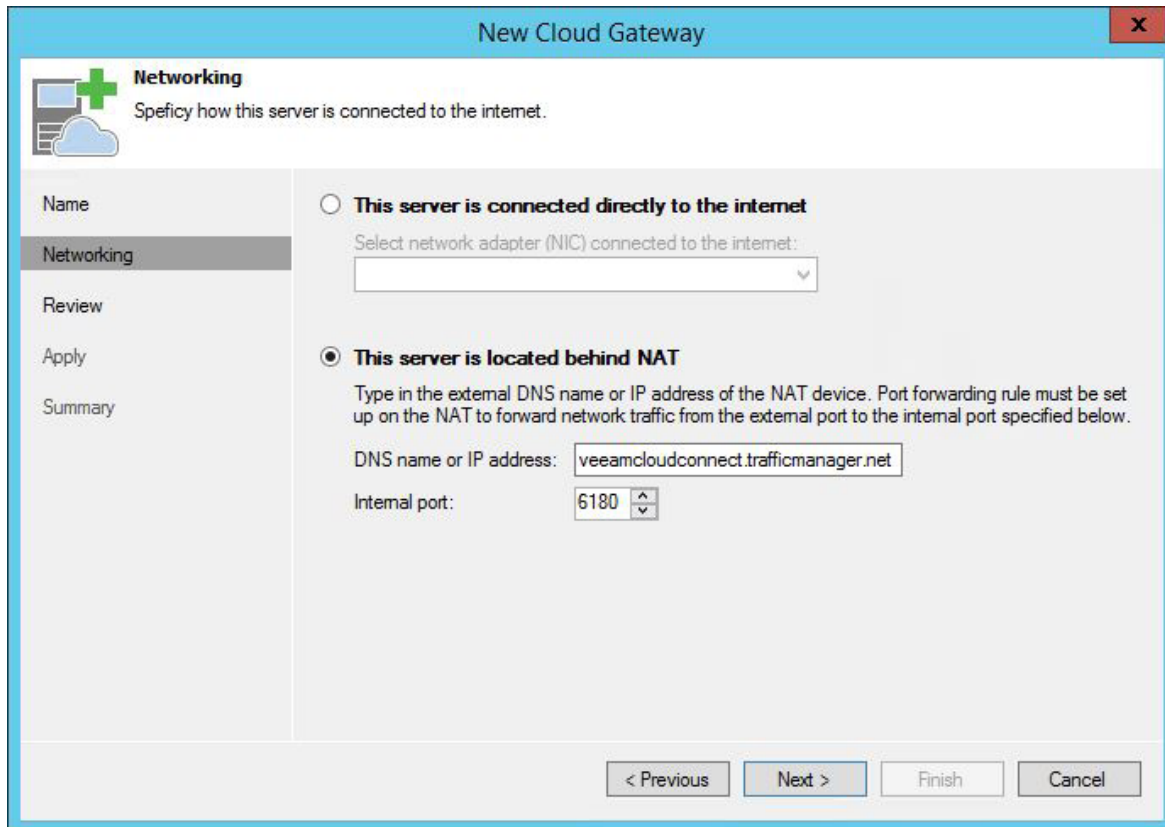


Figure 118: Name, description and external port for the new cloud gateway

The next step is one of the most important ones for a Veeam cloud gateway. We need to specify how the server is connected to the internet, directly or through a NAT. In Azure, public IP addresses are not configured at the guest OS level so we have to choose the option **The server is located behind NAT**. And because we have configured a Traffic Manager profile to balance connections between our different cloud gateway servers, we enter its name in the **DNS name or IP address field**.

**Note:** You can also choose the internal port for Veeam Cloud Connect communications, but you need to make sure that internal port and external port are the same, otherwise it would require to configure port forwarding rules. However, the only way to achieve this in the Azure Resource Manager model is through an Azure Load Balancer, but a load balancer with shared IP address cannot be used to publish multiple cloud gateways.



**New Cloud Gateway**

**Networking**  
Specify how this server is connected to the internet.

**Name**

**Networking**

**Review**

**Apply**

**Summary**

☐ **This server is connected directly to the internet**  
Select network adapter (NIC) connected to the internet:  
[Dropdown menu]

☒ **This server is located behind NAT**  
Type in the external DNS name or IP address of the NAT device. Port forwarding rule must be set up on the NAT to forward network traffic from the external port to the internal port specified below.  
DNS name or IP address: [veeamcloudconnect.trafficmanager.net]  
Internal port: [6180]

< Previous   Next >   Finish   Cancel

Figure 119: New Cloud Gateway networking configuration

Review the settings, you'll see that the **Cloud Gateway** component will be deployed on the virtual machine.

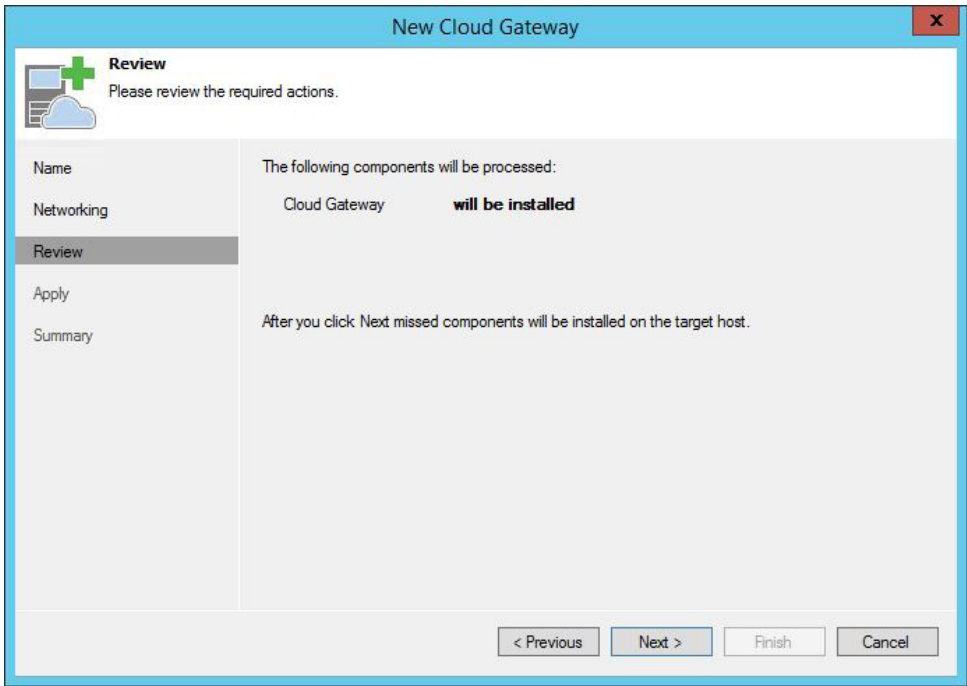


Figure 120: Review screen

On the next page of the wizard, follow the actions taken to turn the VM into an actual Veeam cloud gateway. Once complete, click **Next** and then **Finish**.

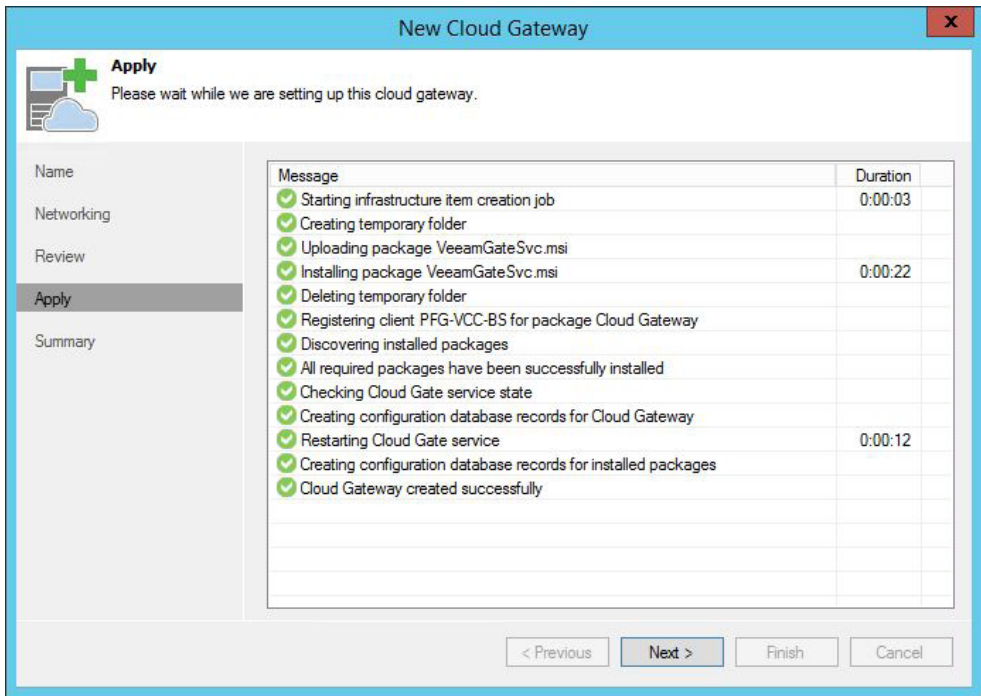


Figure 121: Apply settings

Repeat the steps above to add other cloud gateways to your environment. In this case, we end up with three cloud gateways, two of which we want to keep. Select the cloud gateway that has been automatically deployed and configured on the Veeam Backup & Replication server, look for the one with **This server** in the **Host** column. Click on the **Remove Gateway** button in the ribbon and confirm by clicking **Yes**.

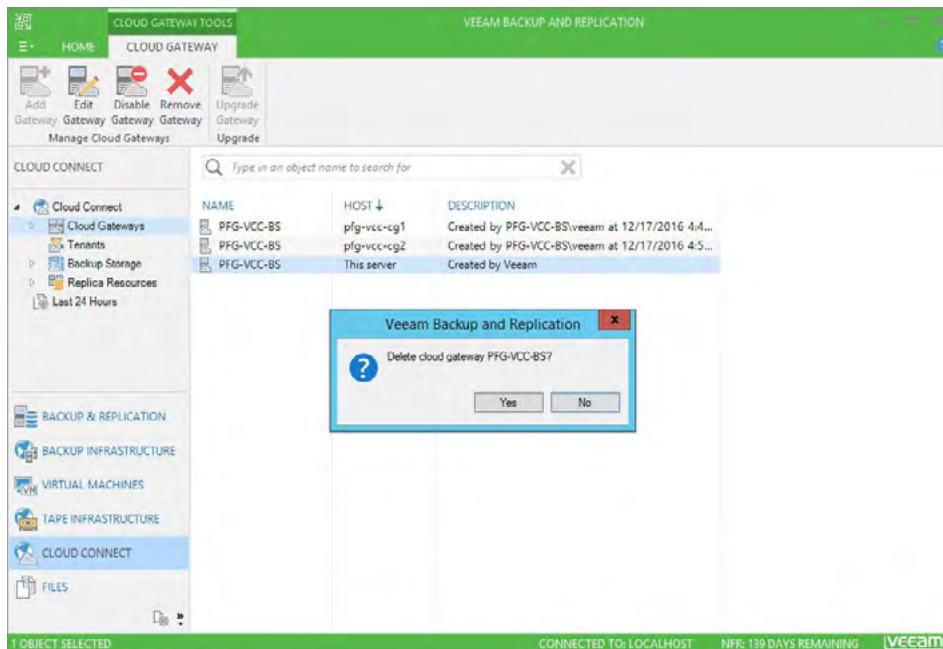


Figure 122: Delete a cloud gateway

Finally, create a new tenant to which we will allocate a quota on the Scale-out Backup Repository. Click on **Tenants**, then **Add Tenant** in the ribbon. Name your new tenant, set a password, an optional description and under **Assigned resources**, select **Backup storage** and click **Next**.

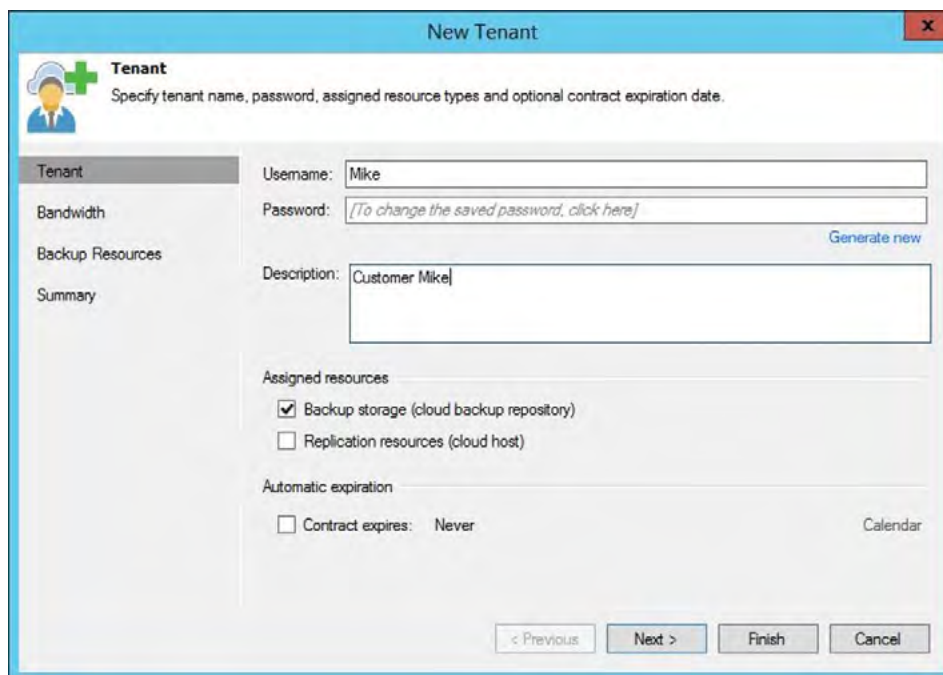
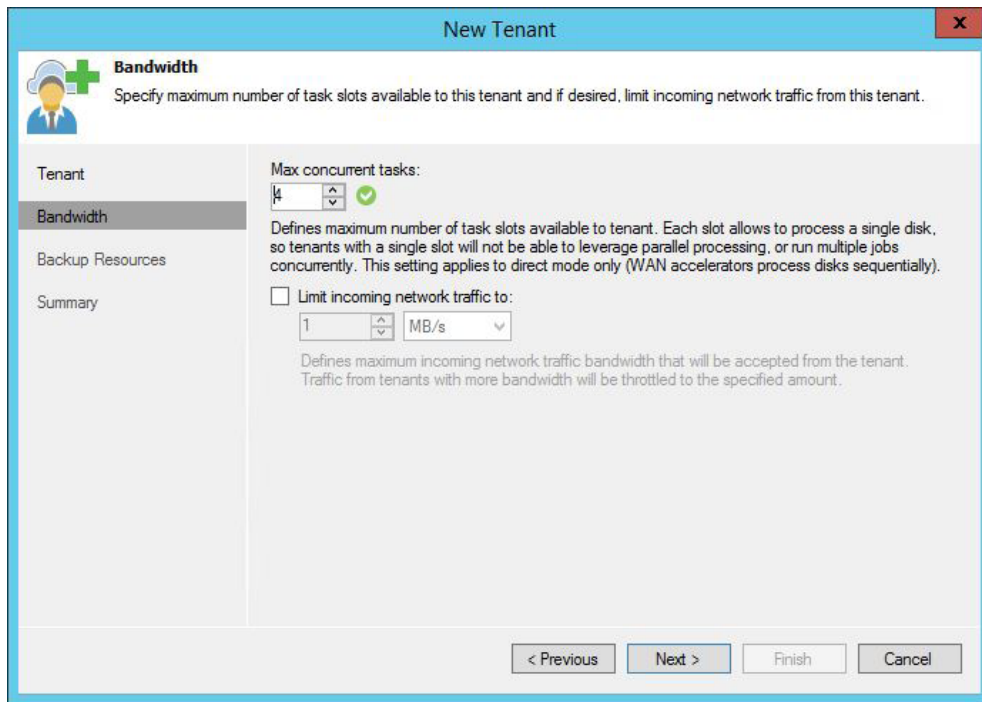


Figure 123: Configure a new tenant



On the **Bandwidth** page, select a maximum number of concurrent tasks for this tenant and a maximum bandwidth if you want to set a limit. Click **Next**.



**New Tenant**

**Bandwidth**  
Specify maximum number of task slots available to this tenant and if desired, limit incoming network traffic from this tenant.

Tenant

**Bandwidth**

Backup Resources

Summary

Max concurrent tasks:  
4

Defines maximum number of task slots available to tenant. Each slot allows to process a single disk, so tenants with a single slot will not be able to leverage parallel processing, or run multiple jobs concurrently. This setting applies to direct mode only (WAN accelerators process disks sequentially).

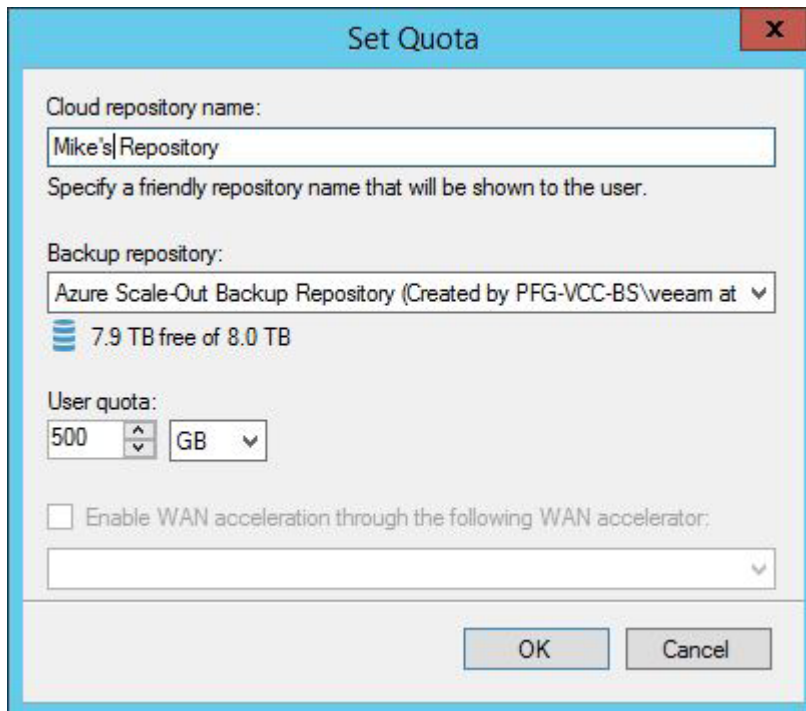
☐ Limit incoming network traffic to:  
1 MB/s

Defines maximum incoming network traffic bandwidth that will be accepted from the tenant. Traffic from tenants with more bandwidth will be throttled to the specified amount.

< Previous   Next >   Finish   Cancel

Figure 124: Configure bandwidth utilization

On the **Backup resources** page, click **Add**. Specify a name, choose the target repository and a quota. Here we choose the Scale-out Backup Repository with a quota of **500 GB**. Click **OK**, then **Next**.



**Set Quota**

Cloud repository name:  
Mike's Repository

Specify a friendly repository name that will be shown to the user.

Backup repository:  
Azure Scale-Out Backup Repository (Created by PFG-VCC-BS\veeam at 7.9 TB free of 8.0 TB)

User quota:  
500 GB

☐ Enable WAN acceleration through the following WAN accelerator:

OK   Cancel

Figure 125: Configure tenant quota

Review the summary and click **Finish**. We're now ready for [The on-premises side!](#)

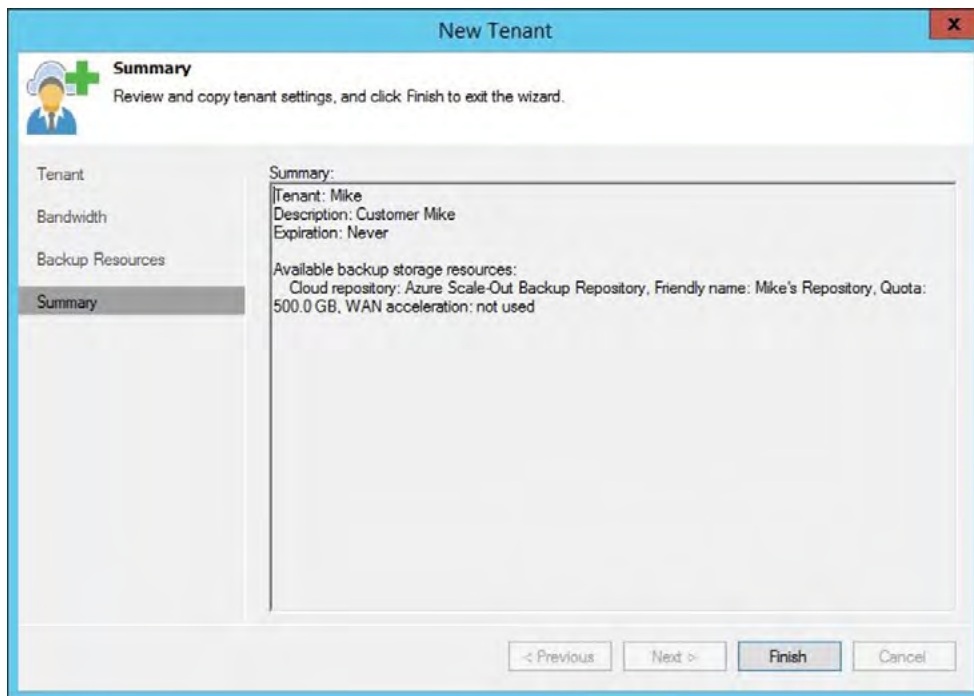


Figure 126: New tenant summary.

## The on-premises side!

You have successfully configured your cloud infrastructure. Now it is time to connect your first tenant. In the scenario for this white paper, the tenant is already protecting their VMs on-premises.

The tenant needs to take the next steps. But first, you need to provide the tenant with some information:

- Cloud gateway: DNS name or IP address. Depending on whether you are using a single-VM deployment or a distributed one, you'll have to give information to connect to the backup server itself, or to the pool of cloud gateways. In this example, we'll point the tenant to the DNS name of the Traffic Manager profile
- Username: The username of the tenant
- Password: The password of the tenant
- Port of the cloud gateway (e.g., **6180**)
- In case of a self-signed certificate<sup>6</sup>, provide the fingerprint of the certificate (e.g., **32E709CD6F0FF598A1A46FBF5A3BB940E0931EF3**)

The tenant will take the following steps:

- Connect to a **service provider**, as labeled in Veeam Backup & Replication's GUI
- Create a Backup Copy job

<sup>6</sup>In case you are using a self-signed certificate, consider using the **Copy to clipboard** button on the last page of the certificate wizard. This will allow you to copy and send the information related to the certificate to the tenant.

## Connect to Veeam Cloud Connect in Azure

In Veeam Backup & Replication, there is an option in the Backup Infrastructure UI called **Service Providers**. From there, the tenant can choose **Add Service Provider** from the ribbon.

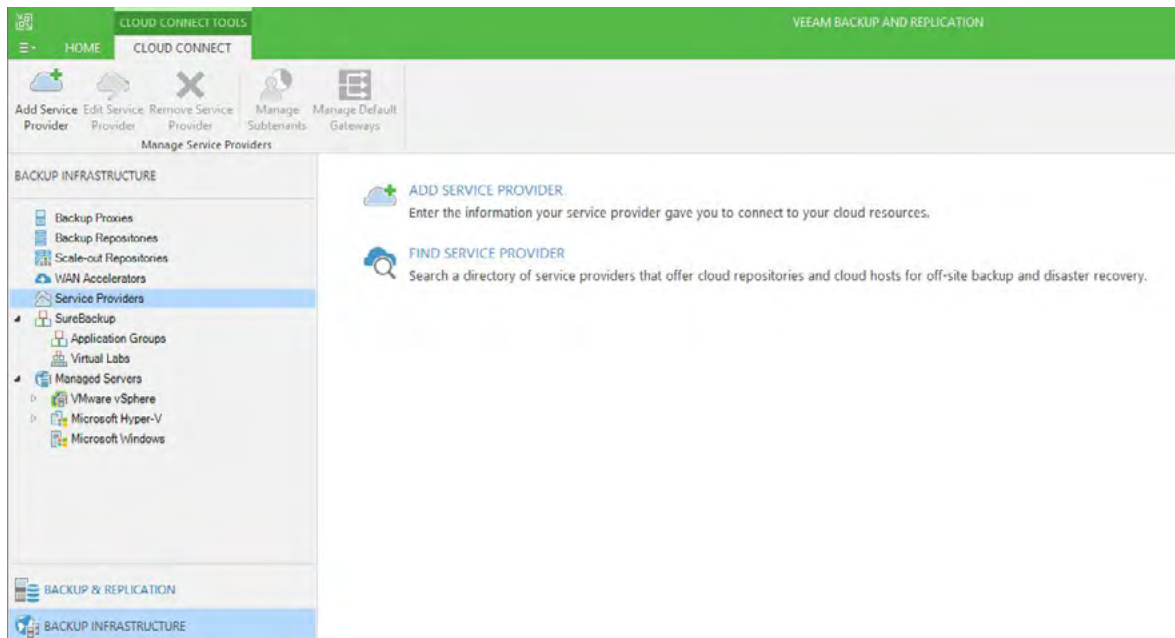


Figure 127: Service provider section

Type in the DNS name or IP address of the cloud gateway or Azure Traffic Manager profile and see if the port matches the port that is configured for the cloud gateway.

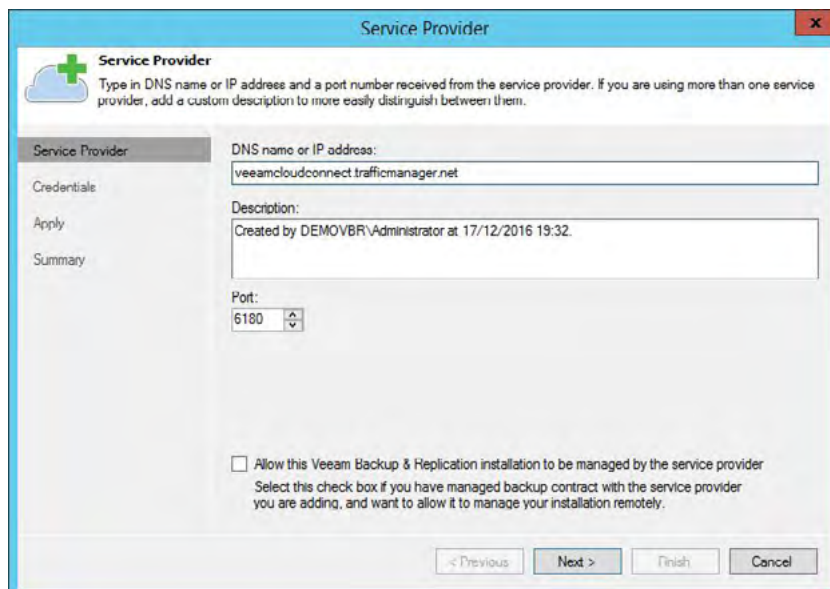


Figure 128: Add a service provider

The wizard will connect to the cloud gateway and request the certificate. The tenant can paste the fingerprint information into the verification box and verify the certificate. The tenant will also add the username and password that IT administrators provide and connect to the cloud infrastructure.

Figure 129: Verification of the certificate and username / password

On the resources page, the tenant will be able to see the available cloud repositories and capacity and whether or not the WAN acceleration is enabled.

Repository	Capacity	WAN Acceleration
Mike's Repository	500.0 GB	Disabled

Figure 130: Review the resources

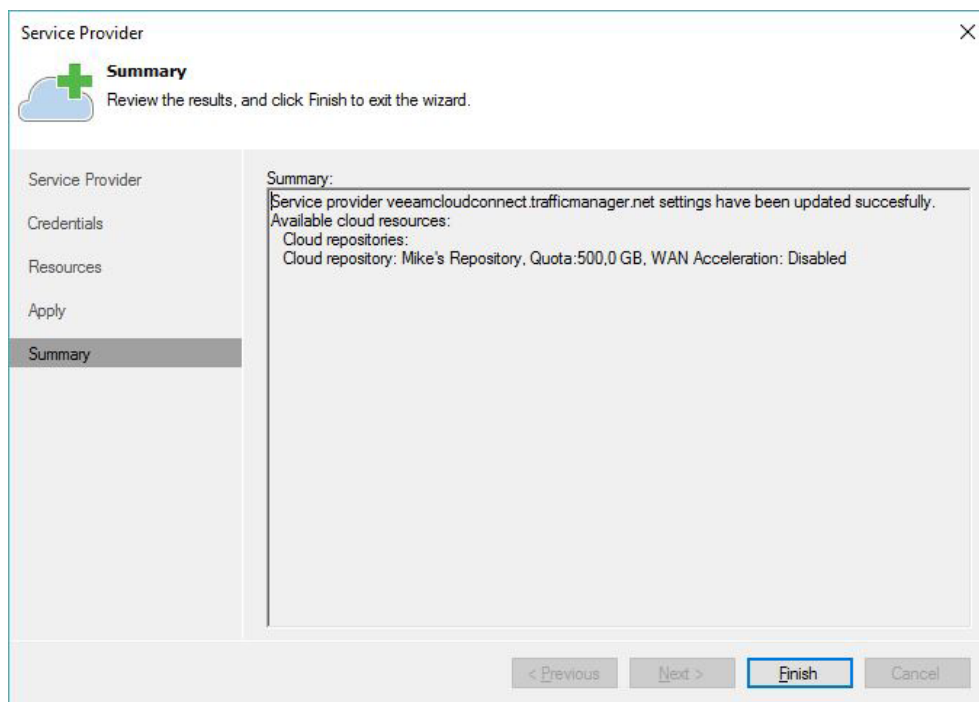


Figure 131: The service provider is added

The tenant is now ready to use the cloud repository as well as view it in their infrastructure.

Local Backup Repository	Windows	demovbr	D:\Backups	200,0 GB	97,9 GB	Created by DEMOVBR
Mike's Repository	Cloud	veeamcloudconnect.trafficmanager.net	\Mike	500,0 GB	461,8 GB	Cloud repository

Figure 132: Cloud repository seen in the infrastructure

## Create a Backup Copy job

The tenant can now configure a Backup Copy job to the service provider to store their data off site. Please note that a Backup Copy job is not the only option. The tenant can use backup, Backup Copy and file copy jobs. However, it is a best practice to maintain the 3-2-1 Rule: three copies of your data on two different types of media and one off site. This rule provides the tenant with quick on-premises restore points for fast recoveries while leveraging Azure as their off-premises archive store.

In this scenario, the tenant is regularly protecting a few virtual machines and the tenant wants to store them in your cloud infrastructure.

On the tenant side, go to **Backup & Replication** and choose the **Backup Copy** button from the ribbon.

Type in a name for the job, a description and how frequently you want to run the Backup Copy job.

**New Backup Copy Job**

**Job**

Backup copy job efficiently creates local and remote copies of your backups, making it easy to maintain multiple copies of your data. Type in a name and description for the job, and specify backup copy interval.

**Job**

Virtual Machines

Target

Data Transfer

Schedule

Summary

Name: Hyper-V Backup Copy to Azure

Description: Created by localhost\administrator at 29/12/2016 15:59.

Copy every: 1 Day starting at 00:00

Controls how often backup copies are created. Backup Copy job creates a new backup file for each copy interval, and starts copying the most recent restore point of each processed VM into this backup file immediately, or as soon as the new restore point appears in the source backup repository.

< Previous Next > Finish Cancel

Figure 133: New Backup Copy job

On the **Virtual Machines** page, add the VMs from the backup and select the required VMs.

**Edit Backup Copy Job [Hyper-V Backup Copy to Azure]**

**Virtual Machines**

Add virtual machines to the job. Consider using containers (such as backup jobs, or infrastructure folders) for dynamic selection scope. No matter how you choose to select VMs, the job will always get VM data from the existing backups files.

**Job**

Virtual Machines

Target

Data Transfer

Schedule

Summary

Objects to process:

Name	Type	Size
demohvcentos	VM	2.3 GB
demohvmail01	VM	39.6 GB
demohvdc	VM	32.0 GB

Add... Remove Exclusions... Source... Up Down Recalculate

Total size: 73.9 GB

< Previous Next > Finish Cancel

Figure 134: Add VM(s) to the job



On the **Target** page, select the cloud repository and the number of restore points to keep. Alternatively, you can create an archive schedule. If necessary, configure advanced settings such as encryption, notifications and deduplication.

**New Backup Copy Job**

**Target**  
Specify the target backup repository, amount of most recent restore points to keep, and retention policy for full backups. You can use map backup functionality to seed the backup files.

**Job**  
Virtual Machines  
**Target**  
Data Transfer  
Schedule  
Summary

**Backup repository:**  
Mike's Repository (Cloud repository)  
461.8 GB free of 500.0 GB [Map backup](#)

**Restore points to keep:** 5

☐ Keep the following restore points as full backups for archival purposes

Weekly backup: 4 Sunday 22:00 [Schedule...](#)  
Monthly backup: 0 First Sunday of the month  
Quarterly backup: 0 First Sunday of the quarter  
Yearly backup: 0 First Sunday of the year

☐ Read the entire restore point from source backup instead of synthesizing it from increments

Advanced settings include health check and compact schedule, notifications settings, and automated post-job activity options. [Advanced](#)

< Previous **Next >** Finish Cancel

Figure 135: Choose the cloud repository

On the **Data Transfer** page, select **Direct**, or choose one of the tenant's WAN accelerators and cloud WAN accelerators when it is enabled by the service provider. See [Appendix A: Using WAN acceleration](#) for more information.

**New Backup Copy Job**

**Data Transfer**  
Choose how VM data should be transferred from source to target backup repository.

**Job**  
Virtual Machines  
Target  
**Data Transfer**  
Schedule  
Summary

☒ **Direct**  
VM data will be sent directly from source to target repository. This mode is recommended for copying backups on-site, and off-site over a fast connection.

☐ **Through built-in WAN accelerators**  
VM data will be sent to target repository through WAN accelerators that must be deployed in both source and target sites. This mode provides for significant bandwidth savings.

Source WAN accelerator:  
Target WAN accelerator:  
Service Provider's WAN Accelerator (Unavailable)

< Previous **Next >** Finish Cancel

Figure 136: Choose direct or a WAN accelerator

On the **Schedule** page, select **Any time** or choose specific time periods to upload data.

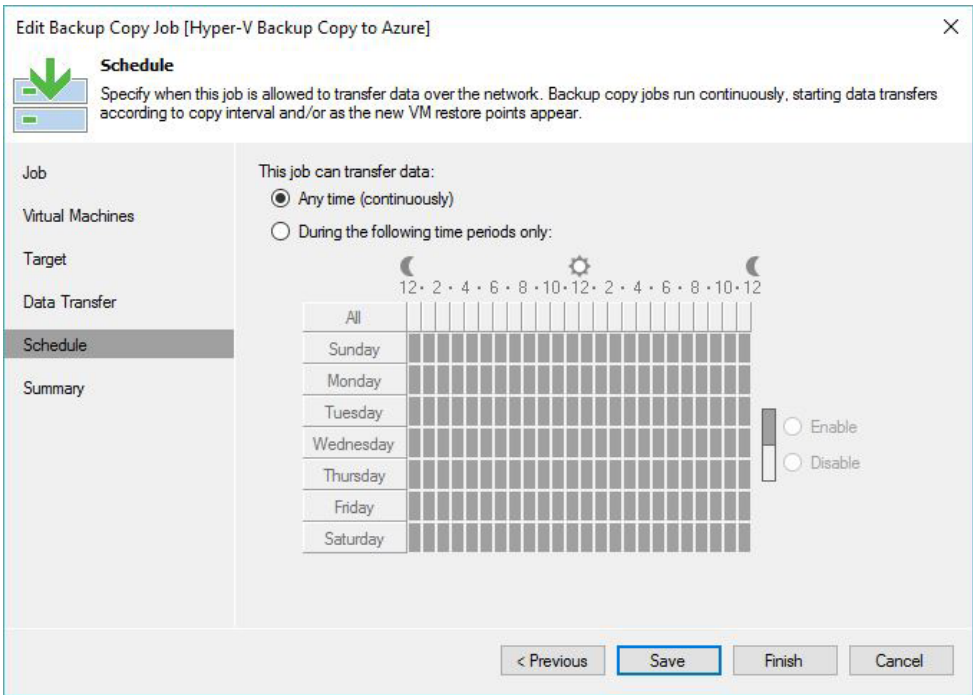


Figure 137: Time to allow upload of data

Review the summary and press **Finish**.

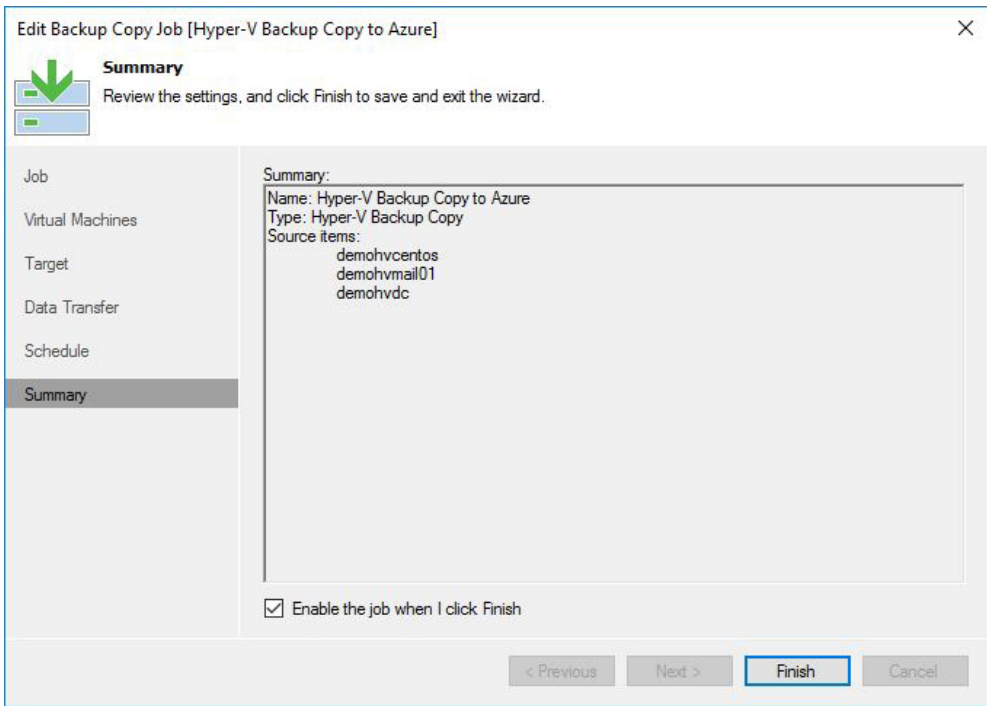


Figure 138: Review the job settings and apply

At the scheduled time, the Backup Copy job will connect to your cloud infrastructure and start the job.

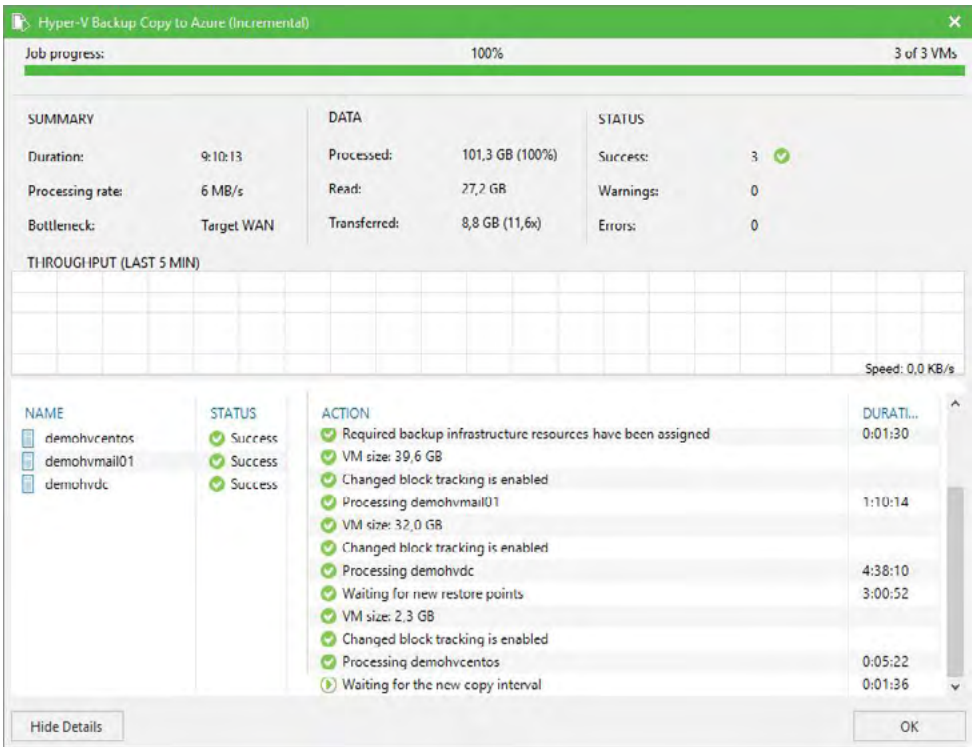


Figure 139: Finished job

When the job has finished, IT administrators can see the used space in the cloud infrastructure.

NAME ↓	QUOTA	FREE SPACE	USED SPACE	USED SPACE %	DESCRIPTION	PATH
Maria	100,0 GB	100,0 GB	0,0 MB	0%	Customer Maria	\\Maria
Mike	150,0 GB	118,7 GB	31,3 GB	21%	Customer Mike	\\Mike

Figure 140: View of used space on the cloud repository

# Restore

Backing up to the cloud is one thing – restoring the data is another thing (and it's the most important). To fully test the capabilities, the tenant, in this example, will do a few restore tests.

The tenant will perform four different tests:

- Recover an Active Directory item (since it is a Domain Controller)
- Recover a file out of the VM
- Restore the VHDX file
- Restore an entire VM

When you perform file recovery or application-item recovery, you don't need to download the entire VM first. You only need to download the requested file or application item.

## Test 1: Recovering an Active Directory item

The first test in this scenario is to recover an Active Directory item. In the scenario, we have accidentally deleted the administrator and want to recover that as soon as possible.

Go to **Backup & Replication > Backups > Cloud**. Select the created Backup Copy job and the VM in that job and press **Application Items > Microsoft Active Directory** from the ribbon.

Choose the restore point. In this example, we only have one restore point, but if you have multiple restore points stored within the cloud repository, you will be able to choose from different restore points.

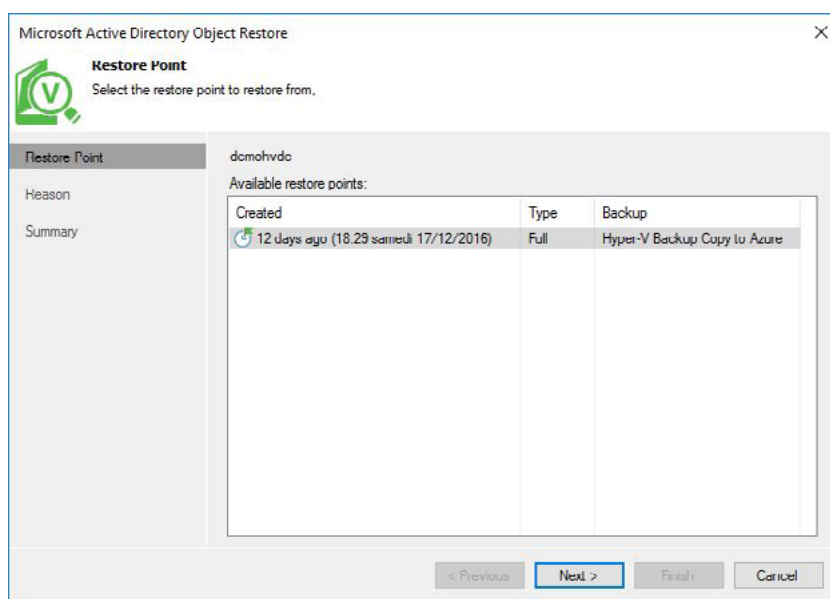


Figure 141: Choose your restore point

Give a reason for restoring.

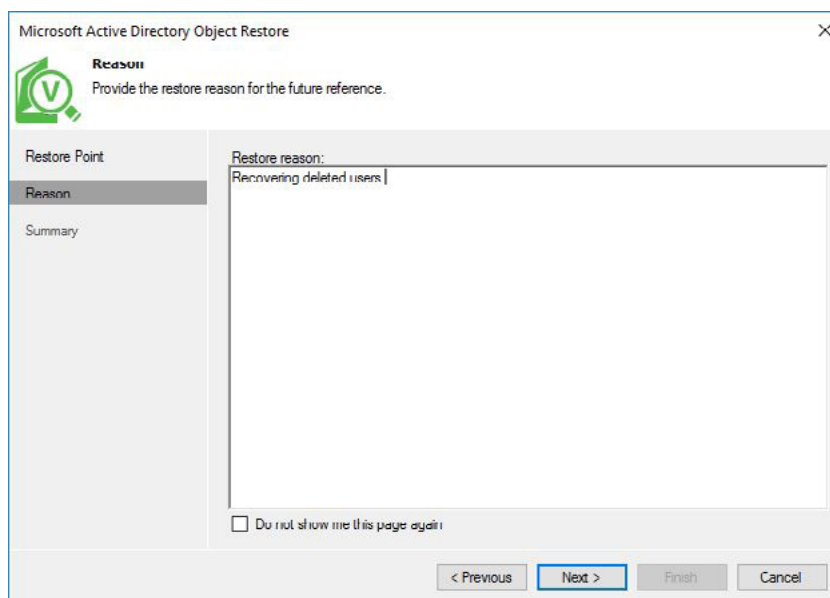


Figure 142: Add the restore reason

Press Finish on the **Completing the Restore Wizard** page.

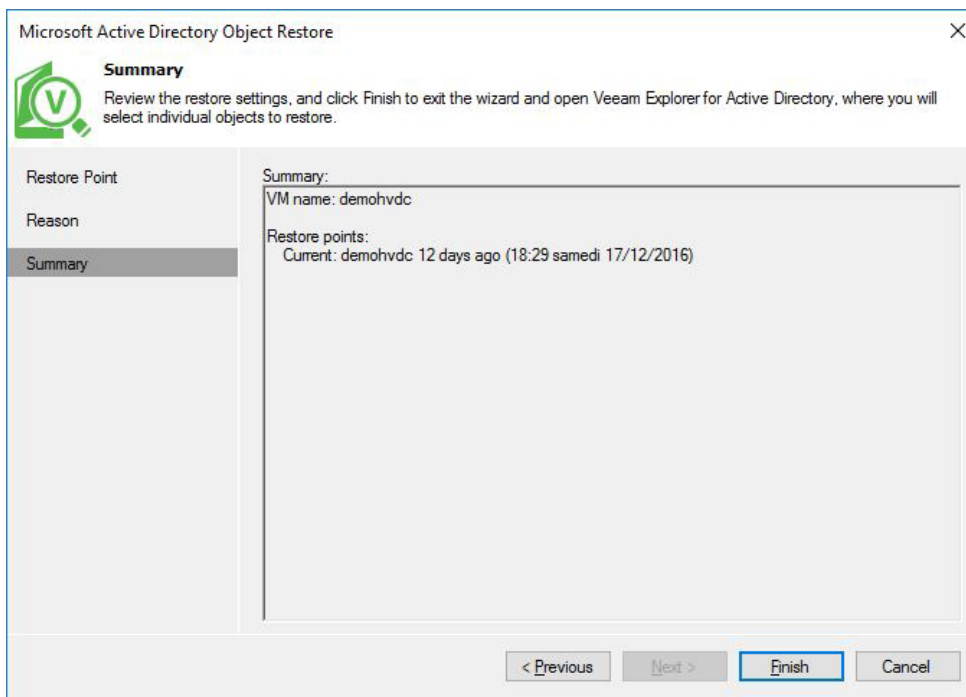


Figure 143: Finish

The system will now connect to the backup on Azure side and open Veeam Explorer™ for Microsoft Active Directory. It will take some time to enumerate the information and load the Active Directory tree in Explorer. The time required depends on the line speed.

When the tree is loaded, you can browse through it and select the object you want to restore.

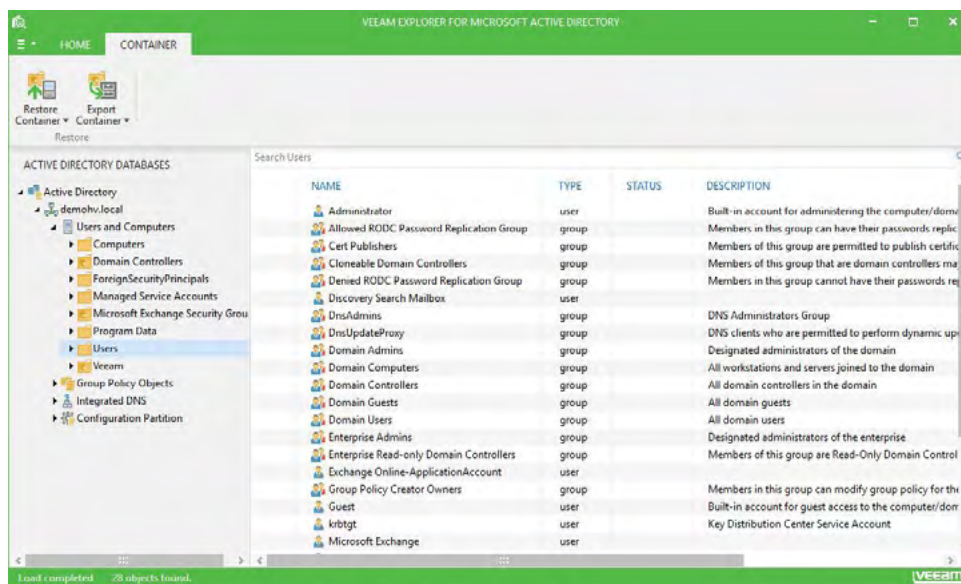


Figure 144: Search or select your object

Right-click on the object and choose your preference for restoring. In this scenario, we decided to export the object to the desktop. It will become an LDF file that can be imported later.




Figure 145: Export options



Figure 146: Finished exporting object(s)

And this is the actual LDF file you can import.

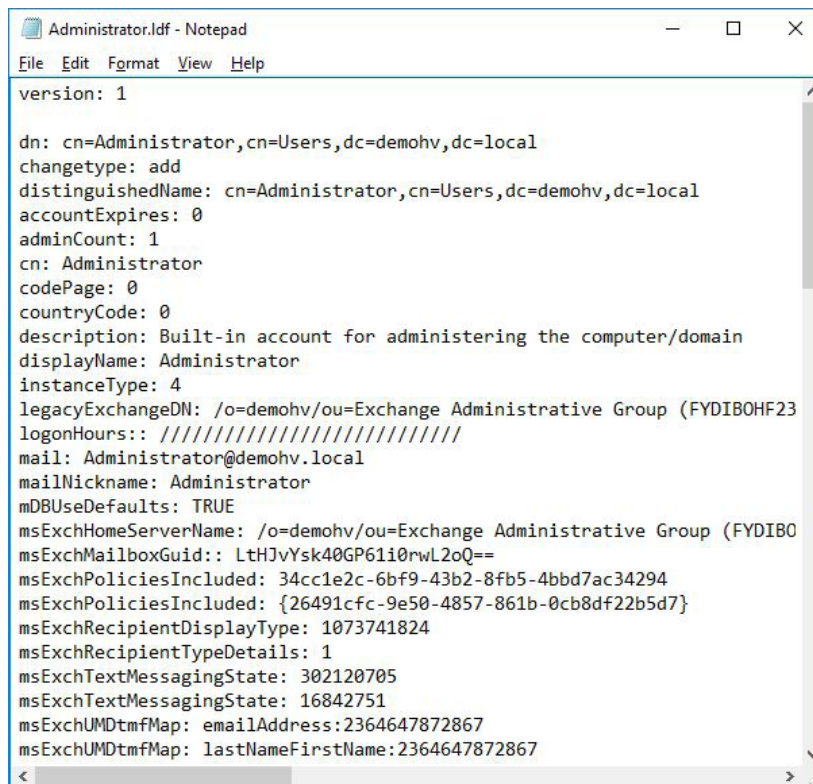


Figure 147: LDF file as a result of the export



## Test 2: Recovering a file

The second test is similar to the first, but now you will restore a file out of the VM. And instead of saving it somewhere else, you will recover it directly to the running VM.

Go to **Backup & Replication > Backups > Cloud**. Select the created Backup Copy job and the VM in that job and press **Guest Files > Microsoft Windows** from the ribbon.

Choose the restore point. In this example, we only have one restore point.

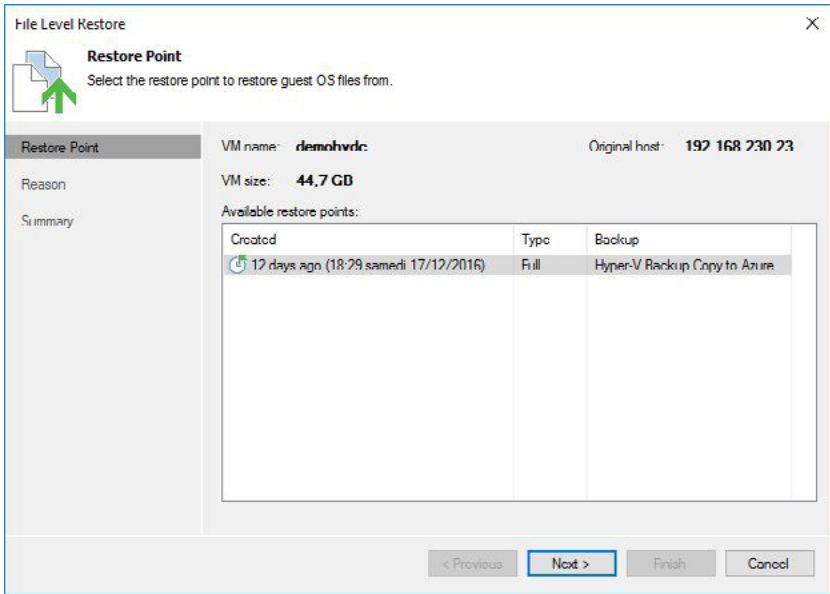


Figure 148: Choose your restore point

Give a reason for restoring.

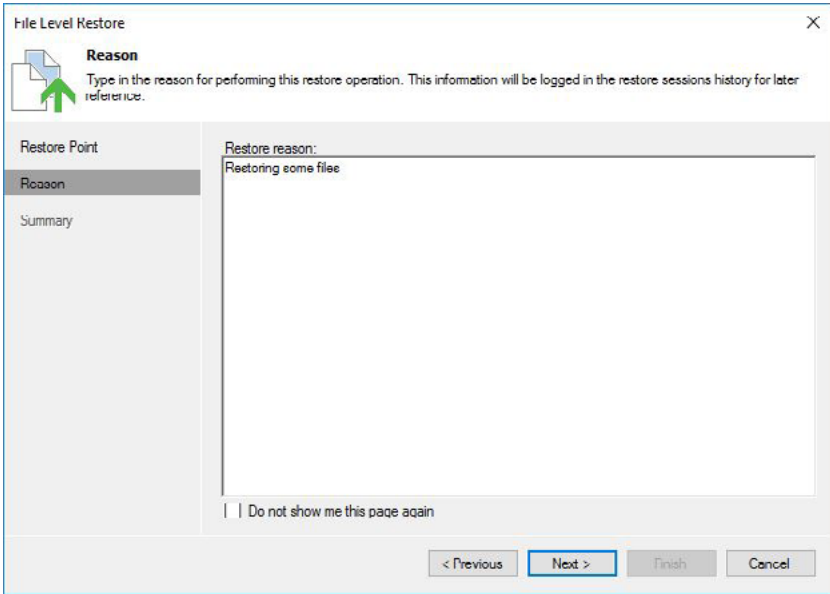


Figure 149: Give a reason for restoring

Press **Finish** to start the backup browser.

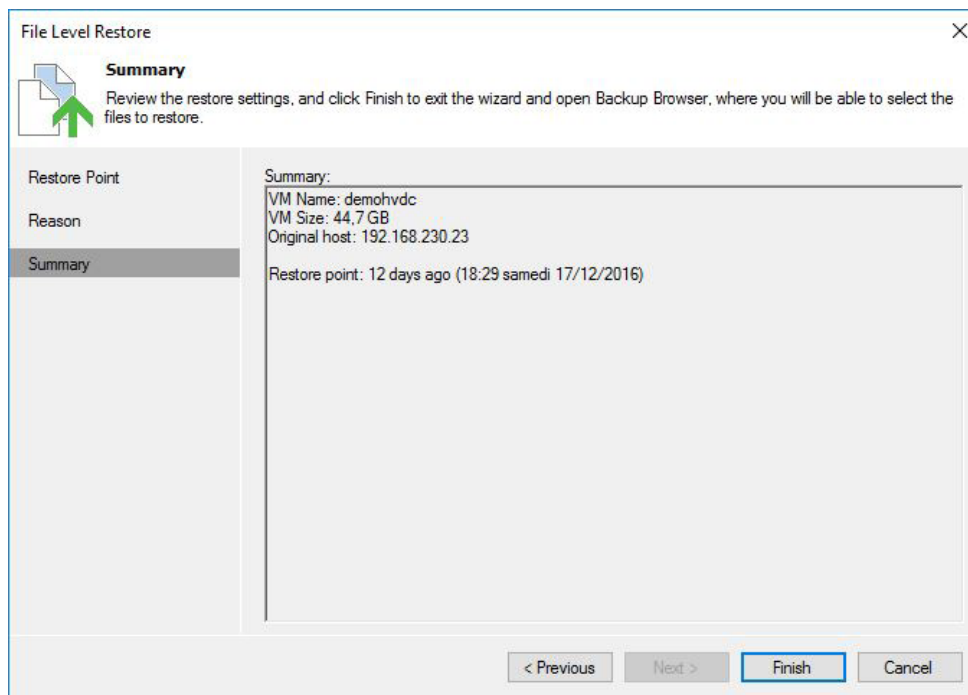


Figure 150: Finish

Next, the system displays a browser that shows the entire file system tree. You can browse in that tree and select the file(s) necessary to restore. As with Veeam Explorer for Microsoft Active Directory, this will take a few minutes. After selecting a file or files, press the **Restore** button.

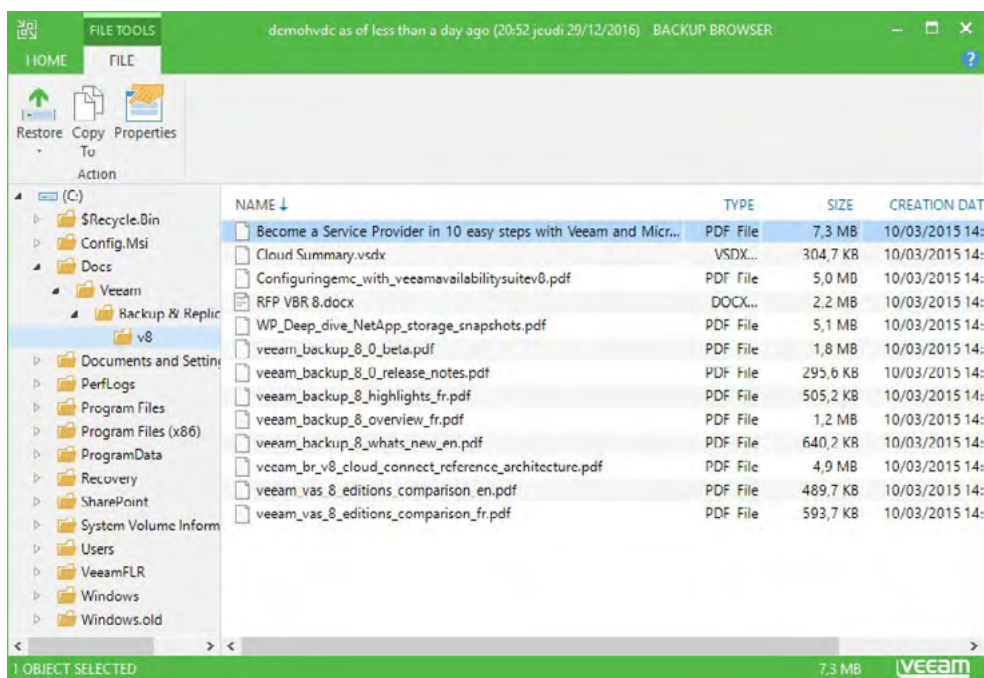


Figure 151: Browse through the file system and select your file

The system will now restore the file.




Figure 152: Restore successful

## Test 3: Recovering a virtual disk

The third test is a bit different. In this example, we want to restore the entire VHDX and import it in a different environment.

Go to **Backup & Replication > Backups > Cloud**. Select the created Backup Copy job and the VM in that job and press **VM Files > VM Files** from the ribbon.

Choose the restore point. In this case, we only have one restore point.




Figure 153: Choose your restore point

Now you can select the VM files that you want to restore. In this case, we're only interested in the VHDX file, so we chose that specific file. Also, select the server location and file path to where you want to restore.

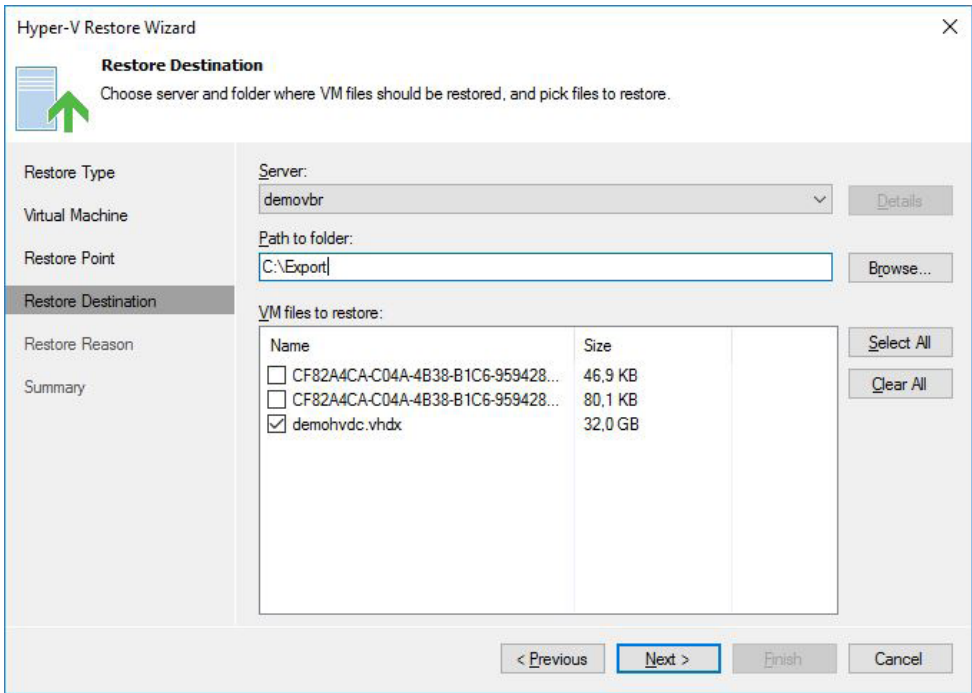


Figure 154: Select the file(s) to restore

Give a restore reason.

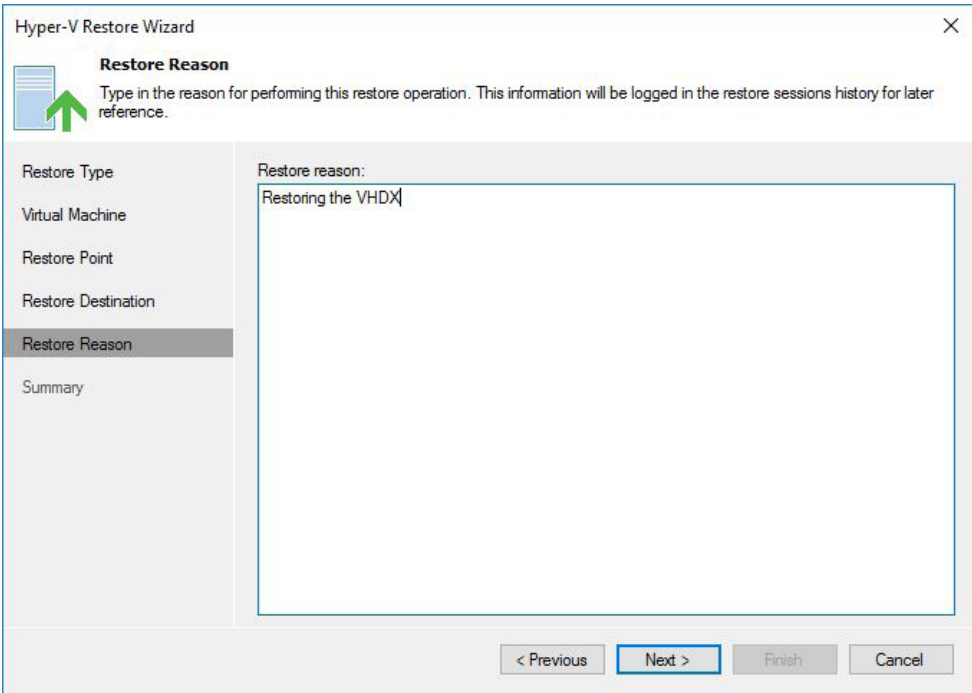


Figure 155: Type in a restore reason

Review the summary and press **Finish**.

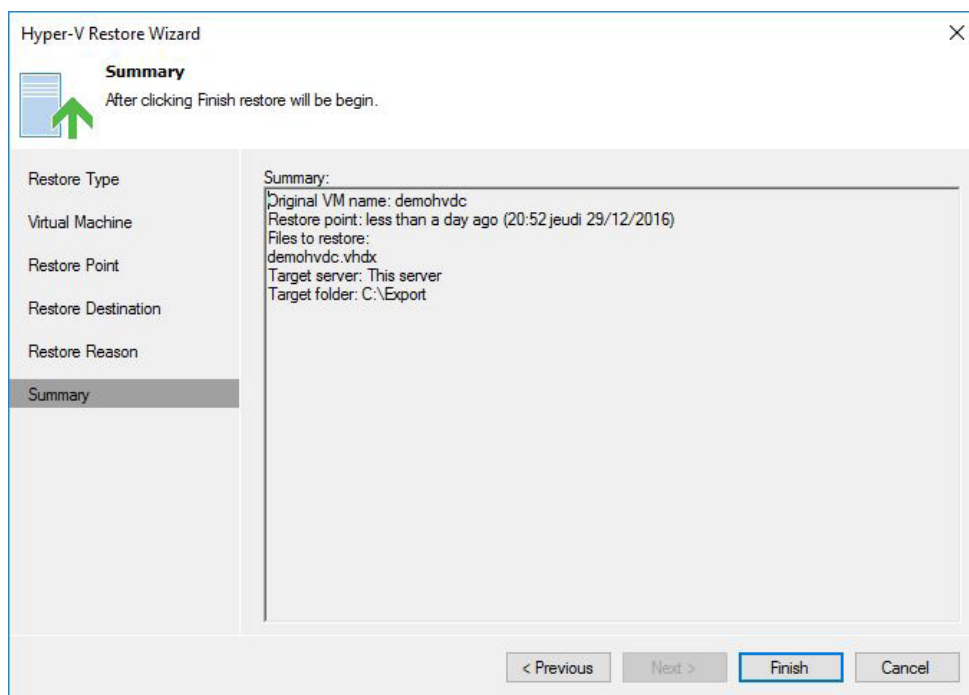


Figure 156: Finish and let the job run

The system is now restoring the VHDX file to the server and location of your choice.




Figure 157: Job window – Log




Figure 158: Job window – Progress




Figure 159: Job success



## Test 4: Recovering the entire virtual machine

The fourth and last test is to recover the entire VM. In this scenario, we don't want to recover the VM to the original host. Assume that we've lost that host and we want to restore it to a different host with settings that differ from the original one.

Go to **Backup & Replication > Backups > Cloud**. Select the created Backup Copy job and the VM in that job and press **Entire VM** from the ribbon.

Select the VM and the restore point (if you have multiple) and press **Next**.

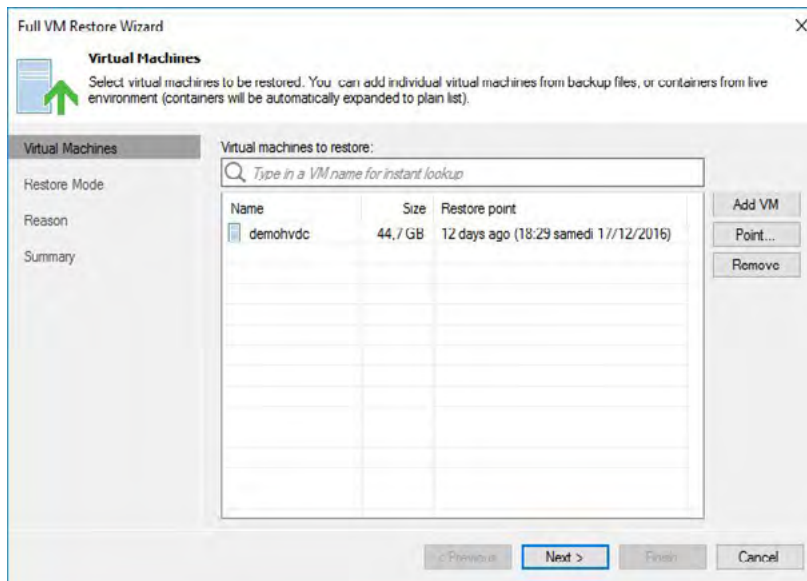


Figure 160: Choose your restore point

Select **Restore to a new location, or with different settings**.

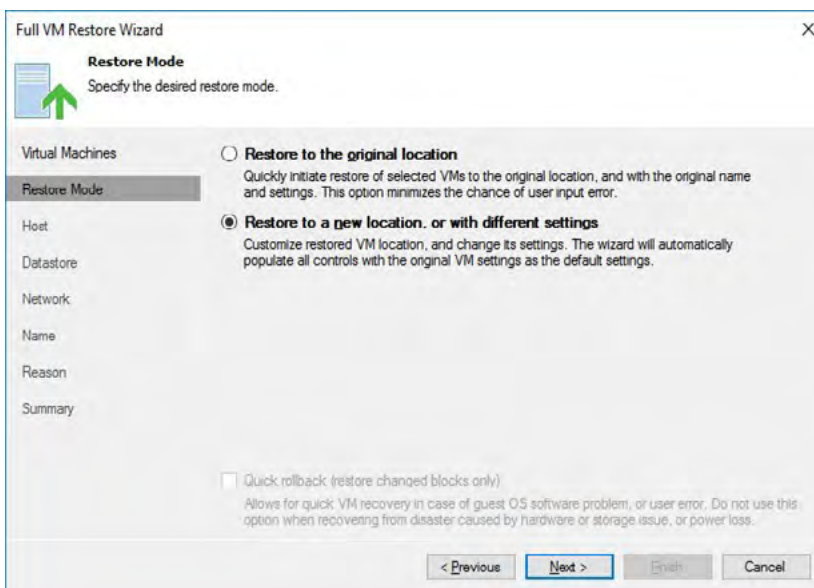


Figure 161: Choose to restore to new location, or with different settings

On the **Host** page, select a different host to restore to and press **Next**.

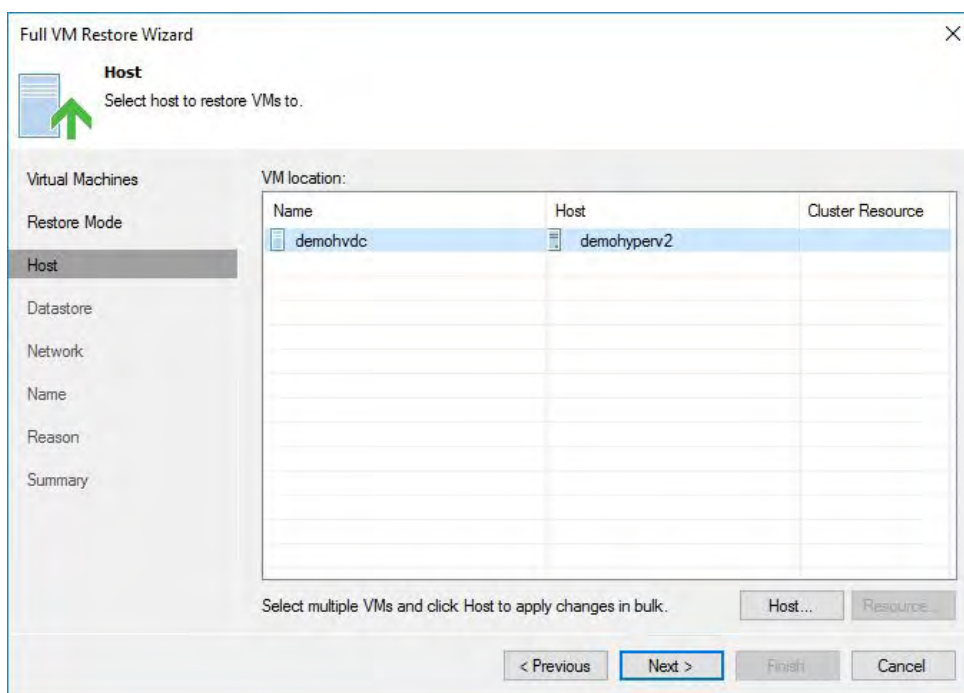


Figure 162: Choose the host to restore to

On the **Datastore** page, change the path location to store this VM.

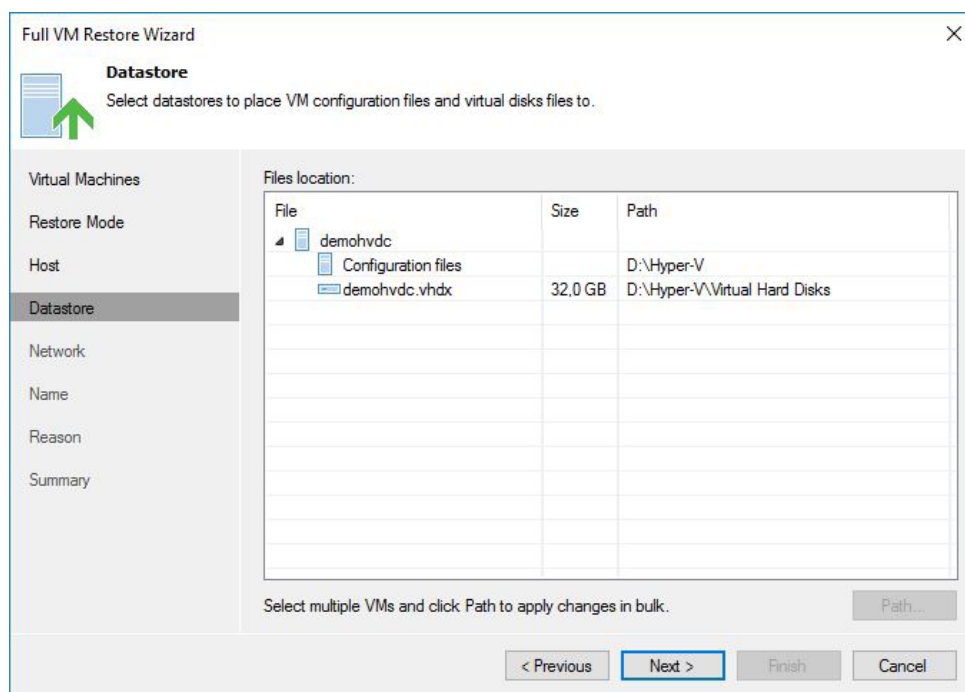


Figure 163: Choose the path location

On the **Network** page, select the network from that host where you want to connect your VM to.

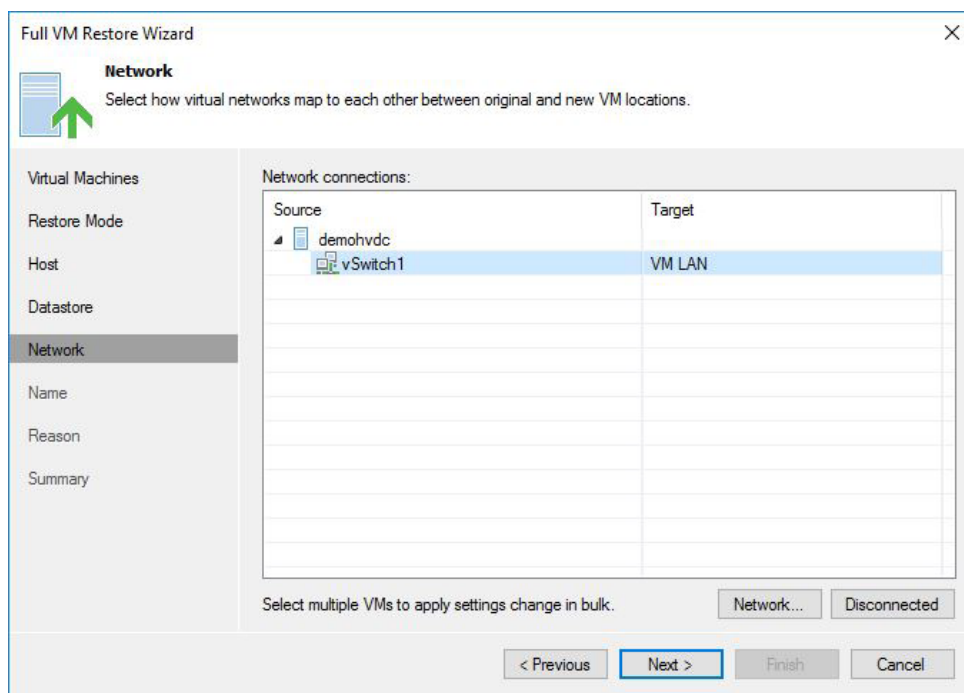


Figure 164: Choose the virtual network to connect to

Finally, choose a new name if necessary and decide whether you want to preserve the VM UUID or not.

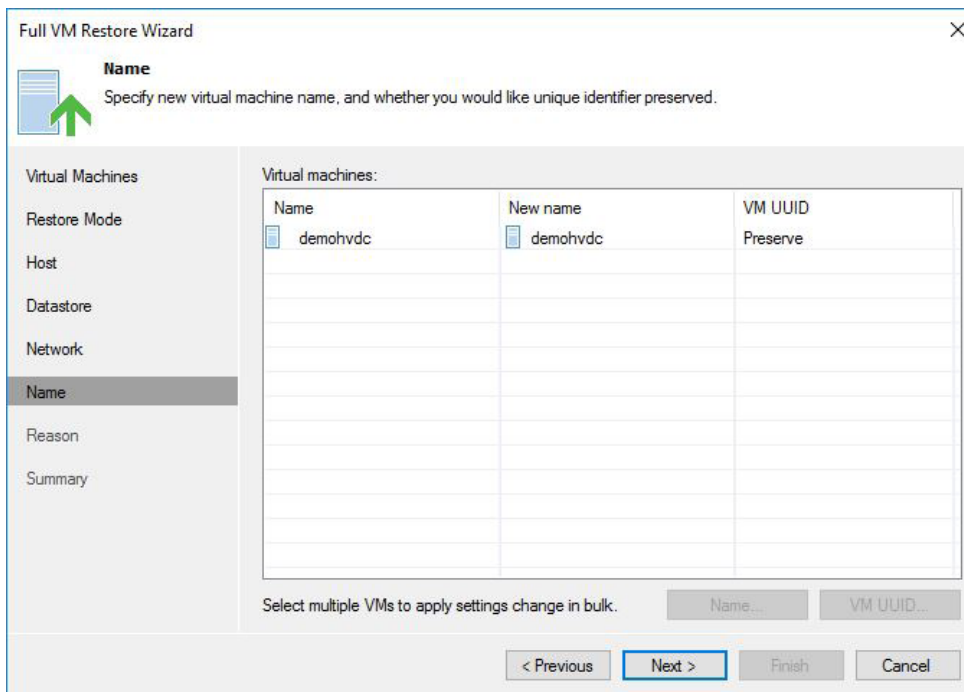
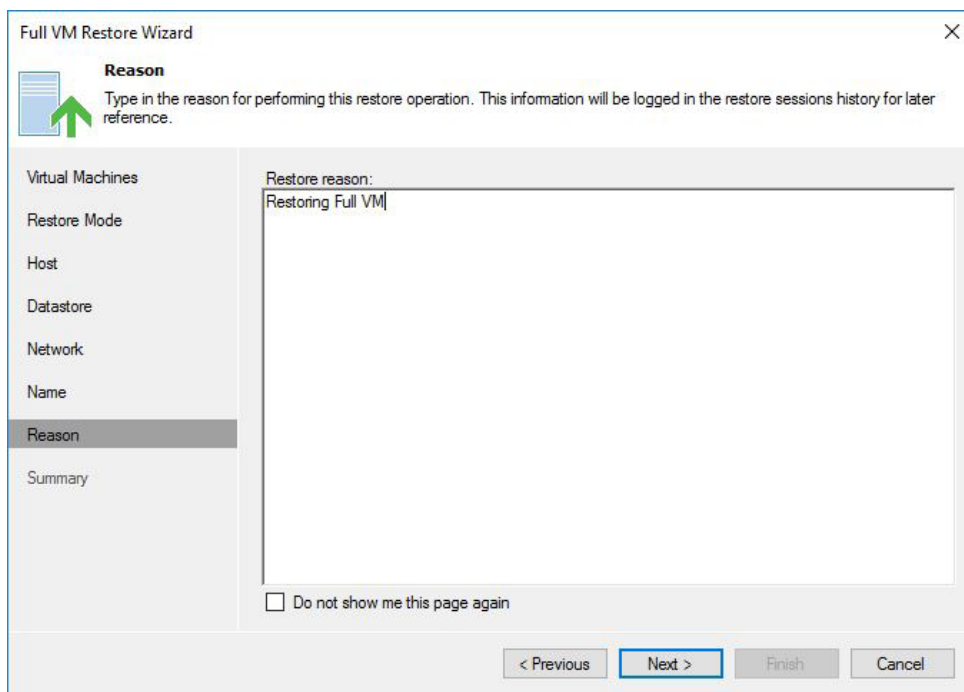


Figure 165: Change the VM name and UUID if necessary

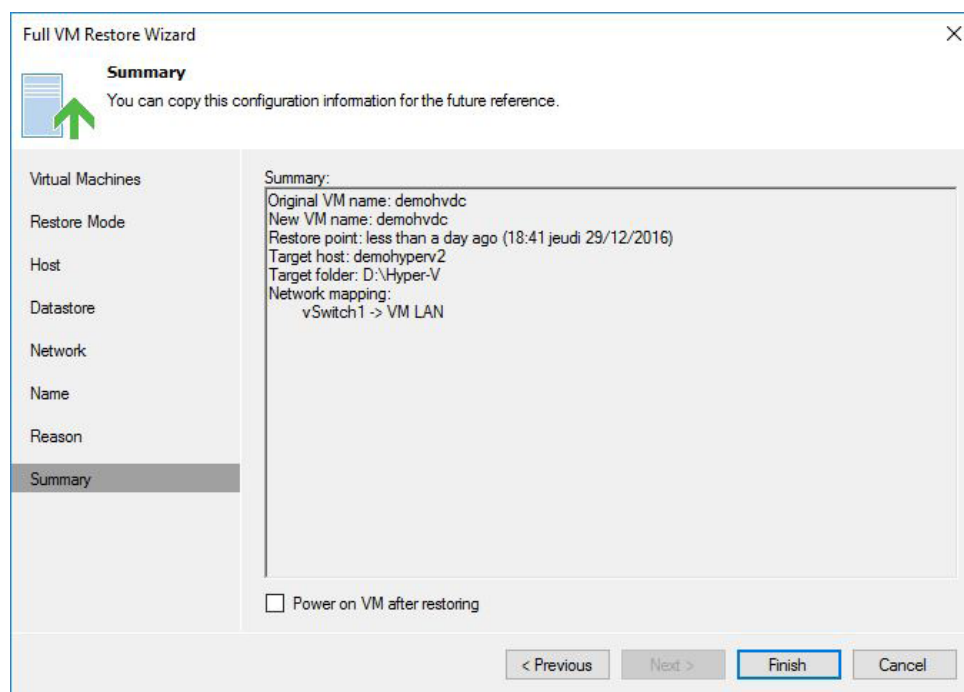
Give a reason for the restore.



The screenshot shows the 'Full VM Restore Wizard' window, specifically the 'Reason' step. The left sidebar contains a list of steps: Virtual Machines, Restore Mode, Host, Datastore, Network, Name, Reason (highlighted), and Summary. The main area has a title 'Reason' with a green arrow icon and a text box for 'Restore reason:'. The text 'Restoring Full VM' is entered in the text box. Below the text box is a checkbox labeled 'Do not show me this page again'. At the bottom, there are four buttons: '< Previous', 'Next >', 'Finish', and 'Cancel'.

Figure 166: Type in the restore reason

Review your settings and decide if you want to automatically start the VM after the restore or not.



The screenshot shows the 'Full VM Restore Wizard' window, specifically the 'Summary' step. The left sidebar contains a list of steps: Virtual Machines, Restore Mode, Host, Datastore, Network, Name, Reason, and Summary (highlighted). The main area has a title 'Summary' with a green arrow icon and a text box for 'Summary:'. The text box contains the following information: Original VM name: demohvdc, New VM name: demohvdc, Restore point: less than a day ago (18:41 jeudi 29/12/2016), Target host: demohyperv2, Target folder: D:\Hyper-V, Network mapping: vSwitch1 -> VM LAN. Below the text box is a checkbox labeled 'Power on VM after restoring'. At the bottom, there are four buttons: '< Previous', 'Next >', 'Finish', and 'Cancel'.

Figure 167: Review and start the restore job

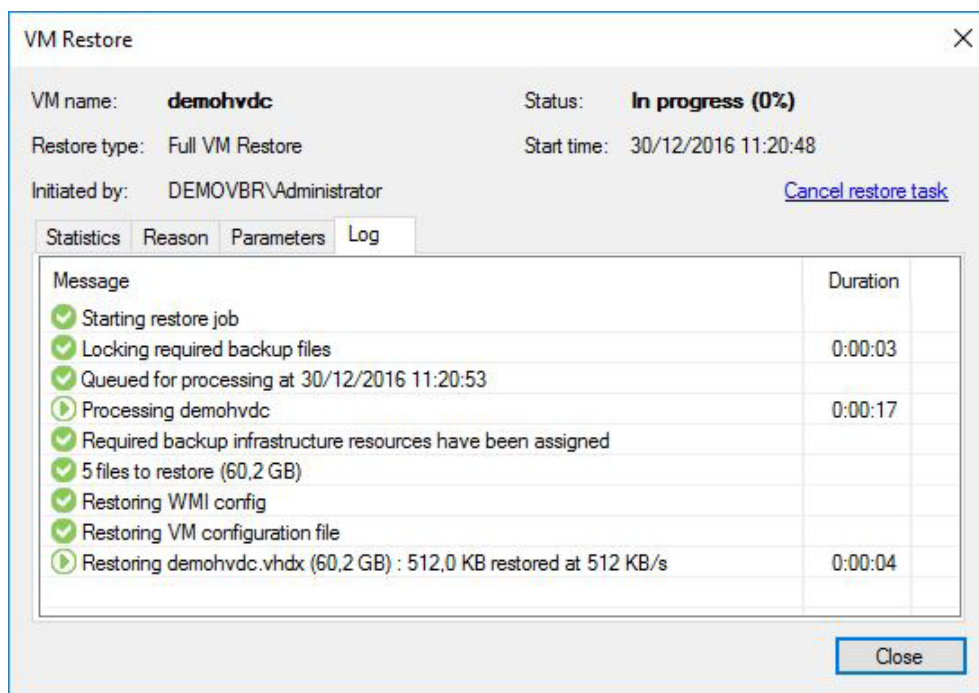


Figure 168: Job log window




Figure 169: Finished and successful restore

## Conclusion

Becoming your own internal service provider with minimum effort is easy with Veeam Cloud Connect *for the Enterprise* and Microsoft Azure. You can set up your cloud infrastructure quickly and easily in Microsoft Azure's IaaS offering. You can complete the entire setup and configuration in a few hours.

There is not much effort required on the tenant side to connect to your cloud infrastructure. And handling backups, Backup Copy jobs, file copy and restores works with the same easy-to-use interface that the tenant already knows from Veeam Backup & Replication.

By following this guide, you can quickly set up a test or POC environment and try it out for yourself. Building your internal Backup as a Service (BaaS) offering with Veeam Backup & Replication and Microsoft Azure only takes eight steps:

1. Sign up for an Azure subscription
2. Create the Veeam Cloud Connect *for the Enterprise* from Azure Marketplace
3. Create and assign a data disk to the VM
4. Complete the initial configuration
5. Configure the repository and optional WAN acceleration
6. Configure Veeam Cloud Connect
7. Connect to a service provider as a tenant
8. Create a Backup Copy job as a tenant



## Appendix A: Using WAN acceleration

If you want to allow your tenants to use WAN acceleration as part of their agreement, you need to deploy one or more WAN accelerators in your Veeam Cloud Connect infrastructure. After you have done that, you can assign a WAN accelerator to the tenant's cloud repository (see Create tenant).

To create a WAN accelerator:

Go to **Backup Infrastructure > WAN Accelerators** and press the **Add WAN Accelerator** button in the ribbon.

Choose the server that will host the role, enter a description and choose the port and maximum number of streams.

**New WAN Accelerator**

**Server**  
Choose a server to install WAN accelerator components on. You can only select between 64-bit Microsoft Windows servers added to the managed servers tree in the console.

Server  
Cache  
Review  
Apply  
Summary

Choose server:  
pfg-vcc-wan1 Add New...

Description:  
Created by PFG-VCC-BS\veeam at 12/29/2016 4:28 PM.

Traffic port : 6165  
TCP/IP port to use for data transfer. Ensure this port is open in any firewall between sites.

Streams: 5  
Using multiple upload streams helps to fully saturate WAN links.

< Previous Next > Finish Cancel

Figure 170: WAN Accelerator wizard

On the next page, select the folder and the cache size that you want to dedicate, press **Next**, review the parameters and press **Finish**. Note that for the purposes of this test, the default OS disk of the Azure virtual machine is selected for storing the WAN accelerator's cache, but it's recommended to have a dedicated disk for it.

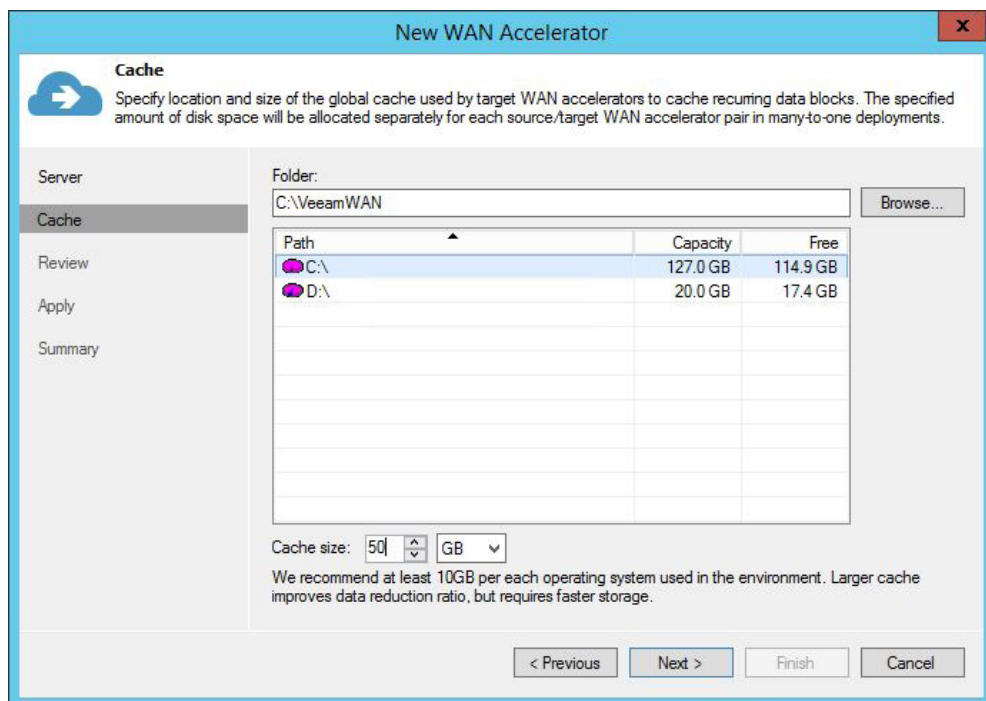


Figure 171: Cache size and location

Note that in order for this to work, both the tenant and the cloud infrastructures need to have a WAN accelerator configured. You don't need to open additional ports as the traffic will be tunneled through the cloud gateway.

## About the Authors



**Mike Ressler** is a product strategy specialist for Veeam. Mike is focused on technologies around Hyper-V and System Center. With years of experience in the field, he frequently presents at large events such as MMS, TechEd and TechDays. Mike has been awarded the MVP for System Center Cloud and Datacenter Management since 2010 and received the Hyper-V MVP since 2014. His major hobby is discussing and developing solid disaster recovery scenarios. Additionally, he has enterprise-class experience in private cloud architecture, deployment with marked focus on protection from the bottom to the top. He holds certifications in many Microsoft technologies such as MCITP.

Follow Mike on Twitter [@MikeRessler](#) or [@Veeam](#) and on [Google+](#).



**Pierre-François Guglielmi** is an Alliance systems engineer for Veeam, in charge of Microsoft and Hewlett Packard Enterprise (HPE) in EMEA. With almost 15 years of experience in IT and a strong virtualization background, he now focuses on Backup as a Service, Disaster Recovery as a Service in hybrid and public clouds, specializing in Microsoft Azure. Pierre-François is also a speaker at large events such as Microsoft TechDays and Microsoft Experience. He's also been certified on different Microsoft technologies since 2004.

Follow Pierre-François on Twitter [@pfguglielmi](#)

## About Veeam Software

[Veeam](#)® recognizes the new challenges companies across the globe face in enabling the Always-On Business™, a business that must operate 24.7.365. To address this, Veeam has pioneered a new market of Availability for the Always-On Enterprise™ by helping organizations meet recovery time and point objectives (RTPO™) of < 15 minutes for all applications and data, through a fundamentally new kind of solution that delivers high-speed recovery, data loss avoidance, verified protection, leveraged data and complete visibility. [Veeam Availability Suite](#)™, which includes [Veeam Backup & Replication](#)™, leverages virtualization, storage, and cloud technologies that enable the modern data center to help organizations save time, mitigate risks, and dramatically reduce capital and operational costs.

Founded in 2006, Veeam currently has 47,000 ProPartners and more than 242,000 customers worldwide. Veeam's global headquarters are located in Baar, Switzerland, and the company has offices throughout the world. To learn more, visit <http://www.veeam.com>.

AVAILABILITY for the Always-On Enterprise™

VEEAM

A nighttime photograph of a city skyline, featuring several illuminated skyscrapers. The foreground shows a body of water reflecting the city lights. The text is overlaid on the image.

# Veeam makes the Fortune 500 Available. 24.7.365

To enable its **Digital Transformation**, 70% of the Fortune 500 rely on Veeam to ensure Availability of all data and applications. **24.7.365**