Mastering Azure Virtual Desktop

The ultimate guide to the implementation and management of Azure Virtual Desktop

Ryan Mangan

Foreword by Jim Moyle, Senior Program Manager, Azure Virtual Desktop, Microsoft

Mastering Azure Virtual Desktop

The ultimate guide to the implementation and management of Azure Virtual Desktop

Ryan Mangan

Packt

BIRMINGHAM—MUMBAI

Mastering Azure Virtual Desktop

Group Product Manager: Rahul Nair

Publishing Product Manager: Preet Ahuja

Senior Editor: Athikho Sapuni Rishana

Content Development Editor: Nihar Kapadia

Technical Editor: Shruthi Shetty

Copy Editor: Safis Editing

Project Coordinator: Shagun Saini

Proofreader: Safis Editing

Indexer: Tejal Daruwale Soni

Production Designer: Prashant Ghare

Marketing Coordinator: Nimisha Dua

First published: March 2022

Production reference: 2290722

Published by Packt Publishing Ltd.

Livery Place

35 Livery Street

Birmingham

B3 2PB, UK.

ISBN 978-1-80107-502-2

www.packt.com

Technology made large populations possible; large populations now make technology indispensable.

— Joseph Wood Krutch

Just because something doesn't do what you planned it to do doesn't mean it's useless.

— Thomas Edison

Be passionate and bold. Always keep learning. You stop doing useful things if you don't learn.

— Satya Nadella

Contributors

About the author

Ryan Mangan is an end user computing specialist. He is a speaker, presenter, and author who has helped customers and technical communities with end user computing solutions, ranging from small to global, 30,000-user enterprise deployments in various fields. Ryan is the owner and author of `ryanmangansitblog`, and has over 3 million visitors and over 200+ articles. Some of Ryan's community and technical awards include Microsoft **Most Valuable Professional** (**MVP**), VMware vExpert 2014, 2015, 2016, 2017, 2018, 2019, 2020, & 2021, VMware vExpert EUC 2021, VMware vExpert Desktop Hypervisor 2021, **Very Important Parallels professional program** (**VIPP**) 2019, 20, & 21, and LoginVSI Technology Advocate 19, and 20.

Writing a book does require lots of time, energy, and dedication, especially in the midst of a pandemic where the customer demand for technology and services increased significantly. I'd like to thank my wife, Alexandra, for supporting me and providing continued motivation as well as the private time to get the book finished. Also, my daughter, Sienna, who continues to this day to ask "what are you doing on the computer, Daddy?"

About the reviewers

Marcel Meurer is responsible for the professional IT services business unit at sepago GmbH in Cologne and is the founder of the development company ITProCloud GmbH. In this role, he leads a team of consultants who provide their expertise in Microsoft and Citrix technologies for customers and partners. His technical focuses are Microsoft Azure platform services, and he has been a Microsoft Azure MVP since 2016.

He loves working in the community. Besides his blog, he publishes tools that simplify working with the Azure cloud – especially in the context of Azure Virtual Desktop. His well-known tools include WVDAdmin and Hydra for Azure Virtual Desktop.

Marcel Meurer graduated as an engineer in electrical engineering from the University of Applied Science, Aachen.

Marco Moioli is a cloud solution architect working for Microsoft's Italian subsidiary.

His goal is to enable Microsoft partners in understanding and to propose solutions based on the Azure cloud and Microsoft 365.

He spent the first part of his career as a consultant/presales engineer at Microsoft specializing in Windows deployment and security.

In 2019, he joined the Microsoft Partner division in West Europe with the role of cloud solution architect, dedicated to Azure Virtual Desktop.

In 2021, he took care of the infrastructure, identity, security, and compliance streams for the Microsoft Partner division in Italy.

He's also the author of the free ebook *Azure Virtual Desktop (Succinctly)*, which will be published in 2022 by Synchfusion.

I'd like to thank Michel Roth and Christiaan Brinkhoff for helping me to become an Azure Virtual Desktop expert.

Neil McLoughlin is based in Manchester in the UK. He has worked in the IT industry for over 20 years, working across many different sectors and roles. He spent around 10 years providing Citrix consultancy for large enterprise customers. Around 5 years ago, Neil discovered the cloud and DaaS and since then has specialized in cloud-based desktop solutions, mainly Azure Virtual Desktop and M365.

Neil is very passionate about community work and runs the UK Azure Virtual Desktop User Group and the Virtual Desktops Community, which is a worldwide community of people interested in Azure Virtual Desktop.

He is currently employed as the UK Field CTO for Nerdio but has previously worked for New Signature, Computacenter, and Cap Gemini as a senior consultant and architect specializing in end user computing.

You can find Neil on Twitter `@virtualmanc`.

Toby Skerritt is an experienced end user architect and engineer. He currently works as technology director for Foundation IT. He has been with Foundation IT for over 10 years, working mainly in the professional service and presales functions. Toby helped the organization to achieve multiple Microsoft accreditations and competencies, including Microsoft Gold Competency status for Cloud Platforms. Toby has been working in the technology space for the past 20 years, working predominantly with Windows OS deployment and virtual desktop technologies. He holds both Azure Administrator Associate and Azure Virtual Desktop Specialty accreditations and has written a number of blogs and opinion pieces on the cloud, Windows Desktop, and cloud desktop solutions.

5
Implementing and Managing Storage for Azure Virtual Desktop

In this chapter, we'll learn how to implement and manage storage for AVD. We'll create a storage account and configure Azure Files for FSLogix Profile Containers.

The following topics will be covered in this chapter:

- Configuring storage for FSLogix components
- Configuring storage accounts
- Creating file shares
- Configuring disks

Configuring storage for FSLogix components

This chapter looks at the storage options that are available for FSLogix Profile Containers when preparing and configuring AVD. We will focus on Azure Files as the storage option of choice as this is the most commonly used storage option for AVD.

FSLogix Profile container storage options

There are three common storage options available for **Azure Virtual Desktop** (**AVD**). This section provides a comparison of the options available to you.

> **Important Note**
>
> Microsoft recommends storing FSLogix Profile Containers in Azure Files unless there is a specific requirement not to. However, this may not meet all organization's requirements.

FSLogix is a profile solution that was acquired by Microsoft to provide Azure Virtual Desktop with roaming profiles by dynamically attaching a virtual hard disk at sign-in. The user profile that's stored on the virtual disk becomes immediately available and appears in the system like a typical user profile.

> **Important Note**
>
> You can use the FSLogix Profile solution outside of AVD.

The following table provides a comparison of the different storage options and features:

Features	Azure Files	Azure NetApp Files	Storage Spaces Direct
Use case	General-purpose.	Ultra performance or migration from NetApp on-premises.	Cross-platform.
Platform service	Yes; Azure-native solution.	Yes; Azure-native solution.	No, self-managed.
Regional availability	All regions.	Select regions.	All regions.
Redundancy	Locally redundant/ zone-redundant/ geo-redundant/ geo-zone-redundant.	Locally redundant/ cross-region replication.	Locally redundant/zone-redundant/geo-redundant.
Tiers and performance	Standard (transaction optimized). Premium. Up to a maximum of 100K IOPS per share with 10 GBps per share at about 3 ms latency.	Standard. Premium. Ultra. Up to 4.5GBps per volume at about 1 ms latency. For IOPS and performance details, see the Azure NetApp Files performance considerations and the FAQ.	Standard HDD: Up to 500 IOPS per-disk limits. Standard SSD: Up to 4K IOPS per disk limits. Premium SSD: Up to 20K IOPS per-disk limits. We recommend Premium disks for Storage Spaces Direct.
Max capacity	100 TiB per share, up to 5 PiB per general-purpose account.	100 TiB per volume, up to 12.5 PiB per subscription.	Maximum 32 TiB per disk.
Required infrastructure	Minimum share size 1 GiB.	Minimum capacity pool 4 TiB, minimum volume size 100 GiB.	Two VMs on Azure IaaS (plus Cloud Witness) or at least three VMs without and costs for disks.
Protocols	SMB 3.0/2.1, NFSv4.1 (preview), REST.	NFSv3, NFSv4.1 (preview), SMB 3.x/2.x.	NFSv3, NFSv4.1, SMB 3.1.

This table was taken from the following site: `https://docs.microsoft.com/en-us/azure/virtual-desktop/store-fslogix-profile?WT.mc_id=modinfra-17152-thmaure#azure-platform-details`.

As shown in the preceding table, Azure Files is the likely candidate for AVD deployments, while Azure NetApp Files offers high performance. There is also an Azure **Virtual Machine** (**VM**) option for using Storage Spaces Direct.

The following table details the features available for Azure Files, Azure NetApp Files, and Storage Spaces Direct:

Features	**Azure Files**	**Azure NetApp Files**	**Storage Spaces Direct**
Access	Cloud, on-premises, and hybrid (Azure File Sync)	Cloud, on-premises (via ExpressRoute)	Cloud, on-premises
Backup	Azure backup snapshot integration	Azure NetApp Files snapshots	Azure backup snapshot integration
Security and compliance	All Azure supported certificates	ISO completed	All Azure supported certificates
Azure Active Directory integration	Native Active Directory and Azure Active Directory Domain Services	Azure Active Directory Domain Services and Native Active Directory	Native Active Directory or Azure Active Directory Domain Services support only

This table was taken from the following site: `https://docs.microsoft.com/en-us/azure/virtual-desktop/store-fslogix-profile?WT.mc_id=modinfra-17152-thmaure#azure-management-details`.

The preceding table shows the features that are available for each service, including Azure Files, Azure NetApp Files, and Storage Spaces Direct.

This section looked at the three storage options that are available when you're planning to configure FSLogix Profile Containers. In the next section, we will look at the two different Azure Files tiers.

The different Azure Files tiers

Azure Files has two different tier types of file storage: standard and premium. The key difference between the two is performance, as premium uses **solid-state drives** (**SSDs**) and are deployed in the file storage account type. Premium file share types are helpful in larger organizations where the requirement for higher performance and low latency is required due to the number of users accessing the file share storage.

Standard file shares use **hard disk drives** (**HDDs**) and are deployed as **general-purpose version 2** (**GPv2**) storage account types. Therefore, you should expect to use standard file shares in small environments or organizations with low I/O needs.

> **Important Note**
> Standard file shares are only available in pay-as-you-go billing models. This means that billing is based on the total storage used, whereas when you're using premium file shares storage, you pay for the configured capacity.

The following table provides examples of when you should use standard file shares versus premium file shares:

Deployment Type	Recommended Storage Tier
Fewer than 200 users	Standard file shares
Greater than 200 users	Premium file shares or standard with multiple file shares
Medium	Premium file shares
Heavy	Premium file shares
Power	Premium file shares

This section explored the two different Azure files storage tiers and when to use each type. The next section looks at Azure Files integration with Azure Active Directory Domain Service.

Best practices for Azure Files with AVD

The following are some of the best practices associated with Azure Files when you're configuring it for use with AVD:

- It is advised that you create your storage accounts in the same region as the session host VMs. This is to ensure that latency is kept to a minimum. This also applies to optimal performance when you're using FSLogix Profile Containers.
- It is recommended that you should be using Active Directory integrated file shares for security and that the following permissions should be set:

User Account	**Folder**	**Permissions**
Users	This Folder Only	Modify
Creator/Owner	Subfolders and Files Only	Modify
Administrator (optional)	This Folder, Subfolders, and Files	Full Control

- When you're storing images in Azure Files, it is advised that you store the master image in the same region as where the VMs are being provisioned.

This section looked at the different storage options available to you, including Azure Files, Azure NetApp Files, and Storage Spaces Direct. We also looked at the different storage tiers for Azure Files, Active Directory Domain Services integration, and storage best practices when configuring FSLogix Profile Containers.

Now, let's learn how to configure a storage account.

Configure storage accounts

This section will look at creating a storage account and configuring data protection.

To create a storage account, you need to follow the stepwise procedure detailed in the following subsections.

Step 1 – create a new storage account

From the left menu within the Azure portal, select **Storage accounts** to display a list of your storage accounts. You can also search for `storage accounts` in the top search bar. This is shown in the following screenshot:

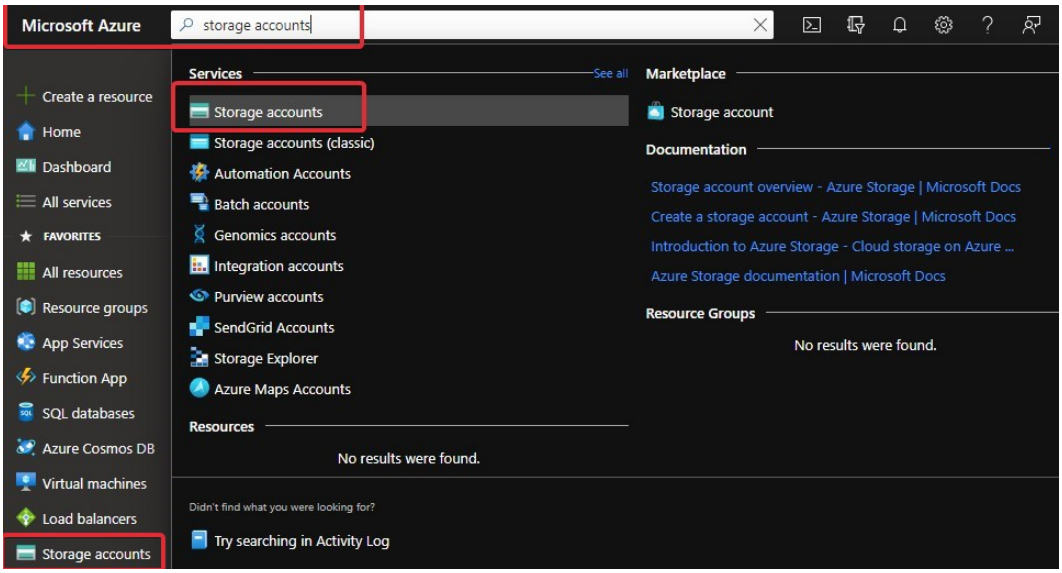

Figure 5.1 – Using the search bar to show the Storage accounts service in the Azure portal

Once on the **Storage accounts** page, you will see all the storage accounts and an icon to create one in the page's navigation bar. To create a new storage account, click **Create**:

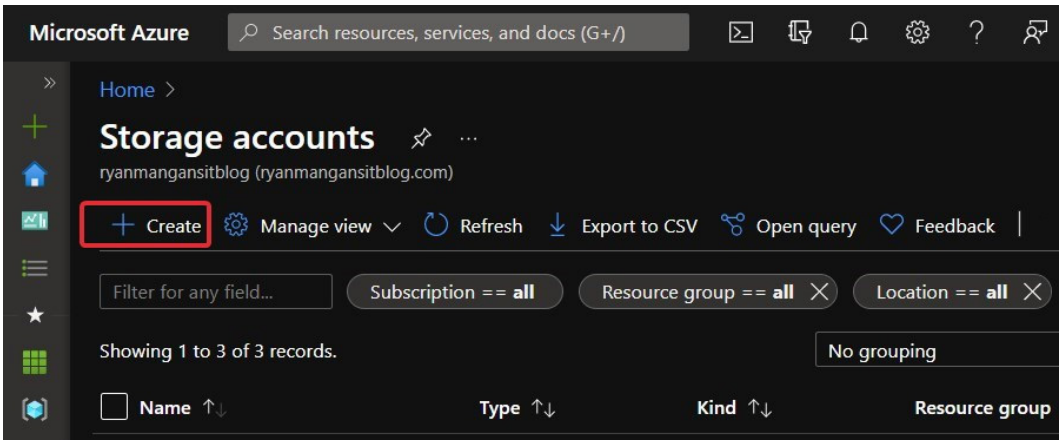

Figure 5.2 – Storage accounts

Once you have clicked **Create**, we can move on to the next section.

Step 2 – configure the basics

Once you have selected **Storage accounts** and clicked **Create**, you will see the basic **Create a storage account** page:

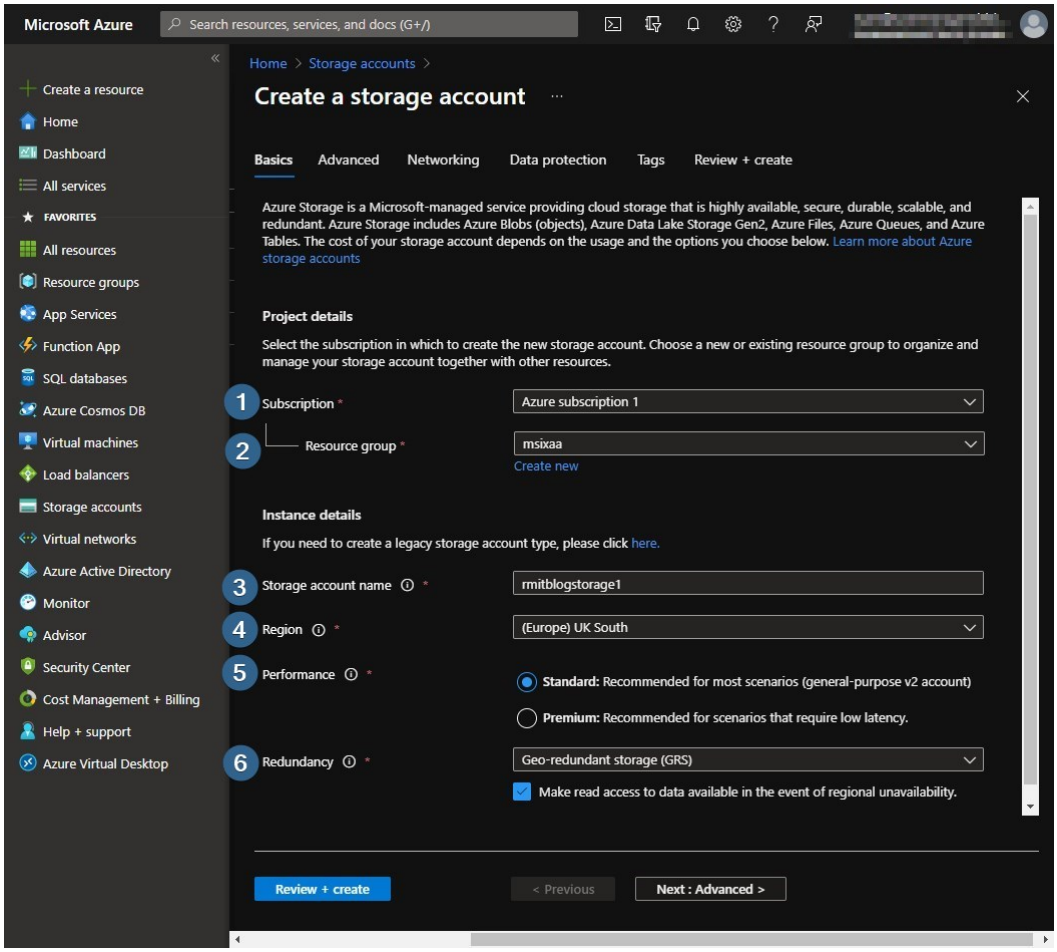

Figure 5.3 – Creating a storage account

The following table details the steps shown in the preceding screenshot. You are required to complete these steps before progressing to the **Advanced** tab:

SN	Name	Required or Optional	Description
1	Subscription	Required	Select the subscription that's required for the storage account.
2	Resource group	Required	Select an existing resource group for this storage account, or create a new one.
3	Storage account name	Required	Choose a name for your storage account; this must be unique. Storage account name must be between 3 and 24 characters in length and contain numbers and lowercase letters only.
4	Region	Required	Select the region for your storage account.
5	Performance	Required	Select Standard performance for general-purpose v2 storage accounts; this is the default. Microsoft recommends this type of account for most scenarios. Use Premium storage for low latency.
6	Redundancy	Required	Select the required redundancy configuration. Remember, not all redundancy options are available in all regions. By selecting geo-redundant configuration (GRS or GZRS), your data is replicated to a data center in a different region. For read access to data in the secondary region, ensure you select Make read access to data available in the event of regional unavailability.

Once you have configured the **Basics** section of creating a new storage account, we can look at configuring advanced settings.

> **Important Note**
>
> Not all regions are supported for all types of storage accounts or redundancy configurations. The choice of region can also have a billing impact.

Step 3 – configure advanced settings

Once you're in the **Advanced** tab, you will see several security and storage configuration options. You can leave these as-is or customize them as required:

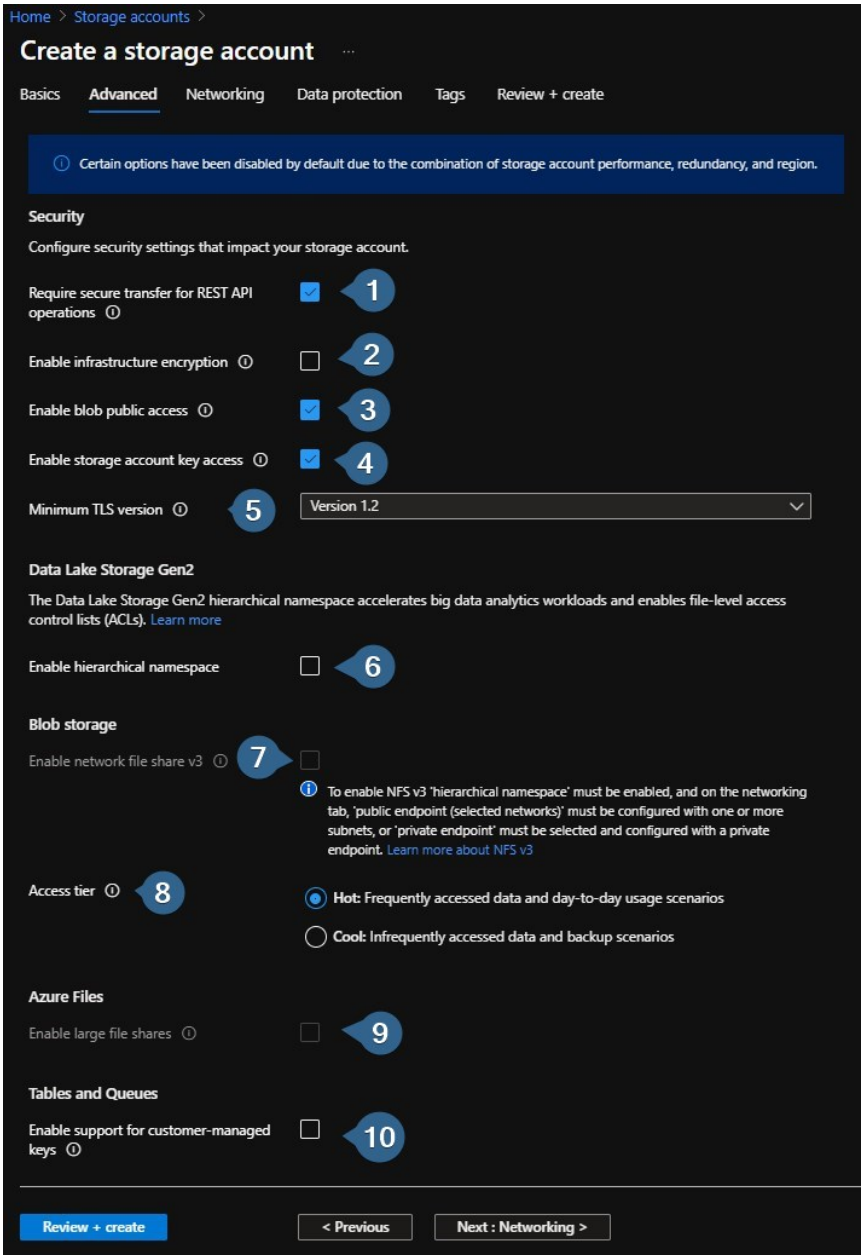

Figure 5.4 – Advanced tab – Create a storage account

The following table details the 10 configuration options. Eight are optional, while two are mandatory. These configuration settings are cross-referenced in the preceding screenshot:

SN	Name	Required or Optional	Description
1	Enable secure transfer	Optional	Enabling secure transfer requires that incoming requests to this storage account are made only via HTTPS (default). This is recommended for optimal security.
2	Enable infrastructure encryption	Optional	Infrastructure encryption is not enabled by default. You can enable infrastructure encryption to encrypt your data at both the service level and the infrastructure level.
3	Enable blob public access	Optional	Users with the appropriate permissions can enable anonymous public access to a container in the storage account when enabling this setting. If you disable this setting, it prevents all anonymous public access to the storage account, making it private.
4	Enable storage account key access	Optional	When enabled, this setting allows clients to authorize requests to the storage account using either the account access keys or an **Azure Active Directory** (**Azure AD**) account. Disabling this setting prevents authorization with the account access keys.

SN	Name	Required or Optional	Description
5	Minimum TLS version	Required	Select the minimum version of **Transport Layer Security** (**TLS**) for incoming requests to the storage account. The default value is TLS version 1.2. When set to the default value, incoming requests made using TLS 1.0 or TLS 1.1 are rejected.
6	Enable hierarchical namespace	Optional	You need to configure a hierarchical namespace to use this storage account for Azure Data Lake Storage Gen2 workloads.
7	Enable **network file system** (**NFS**) v3 (preview)	Optional	NFS v3 provides Linux filesystem compatibility at object storage scale and enables Linux clients to mount a container in Blob storage from an Azure VM or a computer on-premises.
8	Access tier	Required	Blob access tiers enable you to store blob data cost-effectively, based on usage.
9	Enable large file shares	Optional	This is only available for premium storage accounts for file shares.
10	Enable support for customer-managed keys	Optional	To enable support for customer-managed keys for tables and queues, it is advised that you ensure this setting is enabled while creating the storage account.

Once you have chosen the required advanced settings, you can start configuring the **Networking** section.

Step 4 – configure networking

This step is where you configure specific network connectivity requirements, including public and private endpoints. You can also specify the required routing option:

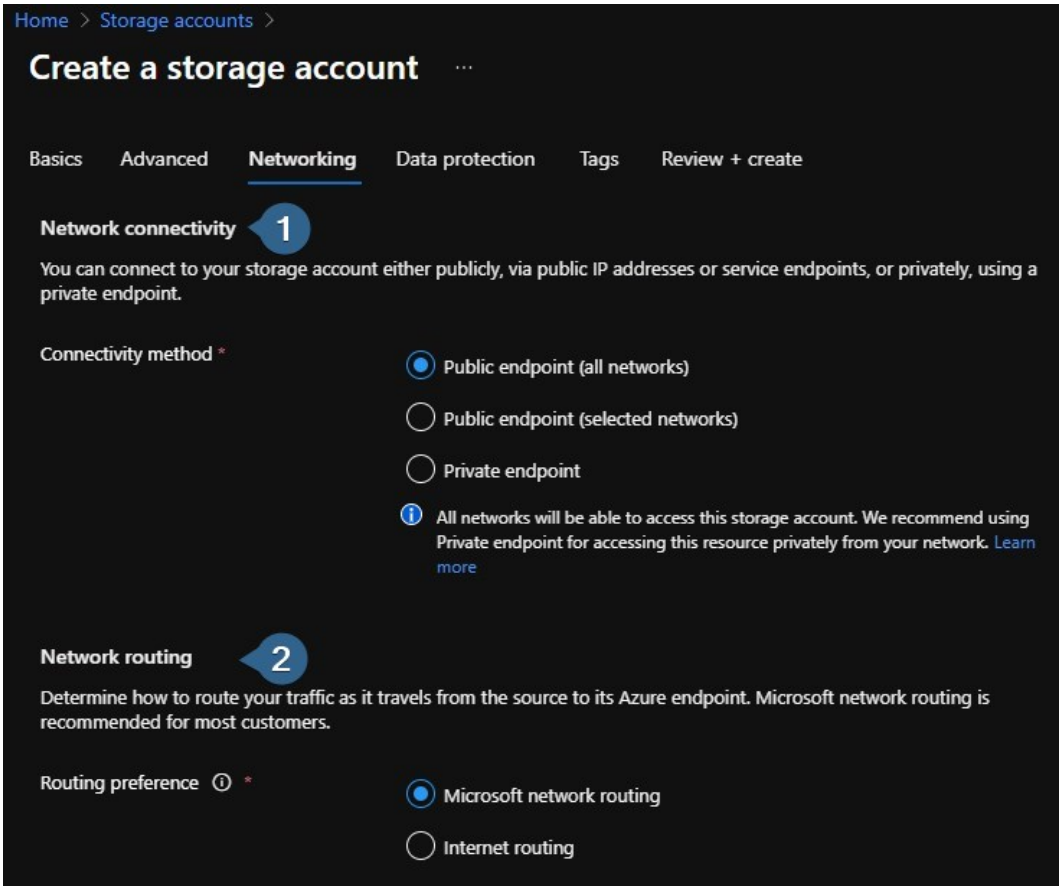

Figure 5.5 – The Networking tab

The preceding screenshot is numbered to reference the **Connectivity method** and **Routing preference** areas shown in the following table:

SN	Name	Required or Optional	Description
1	Connectivity method	Required	Incoming network traffic is routed to the public endpoint for your storage account by default. You can specify that traffic must be routed to the public endpoint through an Azure virtual network. You can also configure private endpoints for private network communication to the storage account.
2	Routing preference	Required	This setting specifies how network traffic is routed to the public endpoint of your storage account from clients over the internet. A new storage account uses Microsoft network routing by default. However, you can also configure how network traffic is routed through the **point of presence** (**POP**) closest to the storage account, which may lower networking costs.

Now that you have configured the networking section of the *Creating a storage account*, we can move on to step five, where we will configure the data protection settings for the storage account.

Step 5 – configure data protection

Within this tab, you can configure the various recovery and tracking options for your storage account. The following screenshot, whose numbers are referenced in the following table, shows several options that are available to you:

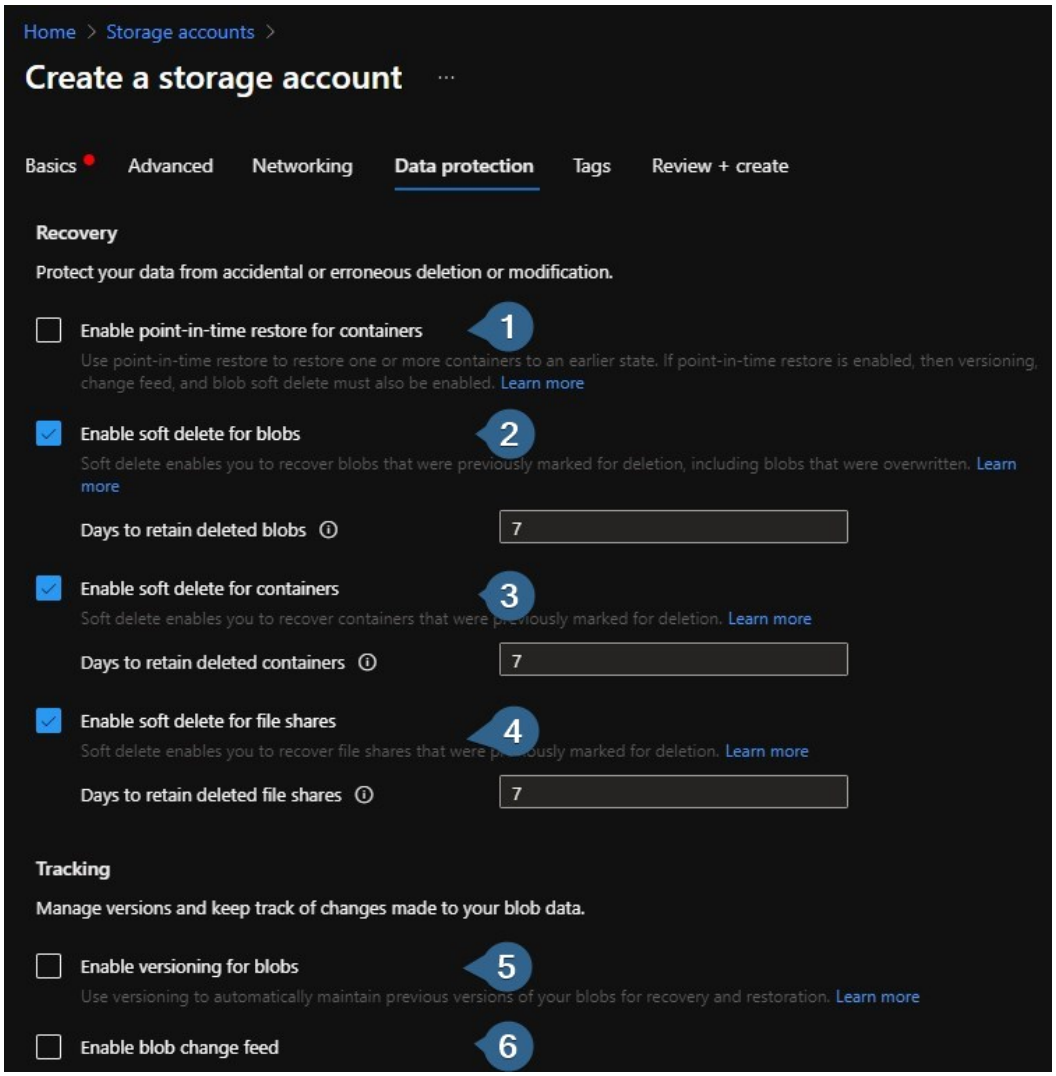

Figure 5.6 – The Data protection tab

The preceding screenshot is annotated with numbers one to six; this correlates with the following table, which shows the options for configuring data protection for the new storage account:

SN	Name	Required or Optional	Description
1	Enable point-in-time restore for containers	Optional	Point-in-time restore protects against accidental deletion or corruption by enabling you to restore block blob data to an earlier state. Enabling point-in-time restore enables blob versioning, blob soft delete, and blob change feed. These prerequisite features have a cost impact, so be sure to check their pricing first.
2	Enable soft delete for blobs	Optional	Blob soft delete protects an individual blob, snapshot, or version from accidental deletes or overwrites by maintaining the deleted data in the system for a specified retention period. During the retention period, you can restore a soft-deleted object to its state when it was deleted.
3	Enable soft delete for containers	Optional	Container soft delete protects a container and its contents from accidental deletes by maintaining the deleted data in the system for a specified retention period. During the retention period, you can restore a soft-deleted container to its state at the time it was deleted.
4	Enable soft delete for file shares	Optional	Soft delete for file shares protects a file share and its contents from accidental deletion by maintaining the deleted data in the system for a specified retention period. During the retention period, you can restore a soft-deleted file share to its state at the time it was deleted.

SN	Name	Required or Optional	Description
5	Enable versioning for blobs	Optional	Blob versioning automatically saves the state of a blob in a previous version when the blob is overwritten. It is recommended that you enable blob versioning for optimal data protection for the storage account.
6	Enable blob change feed	Optional	The blob change feed provides transaction logs of all the changes that have been made to all the blobs in your storage account, as well as their metadata.

Once you have selected the required options for data protection, you can set **Tags** or proceed to the **Review + create** tab.

Within the **Review + create** tab, check if all the settings are as you require, then proceed to create the storage account:

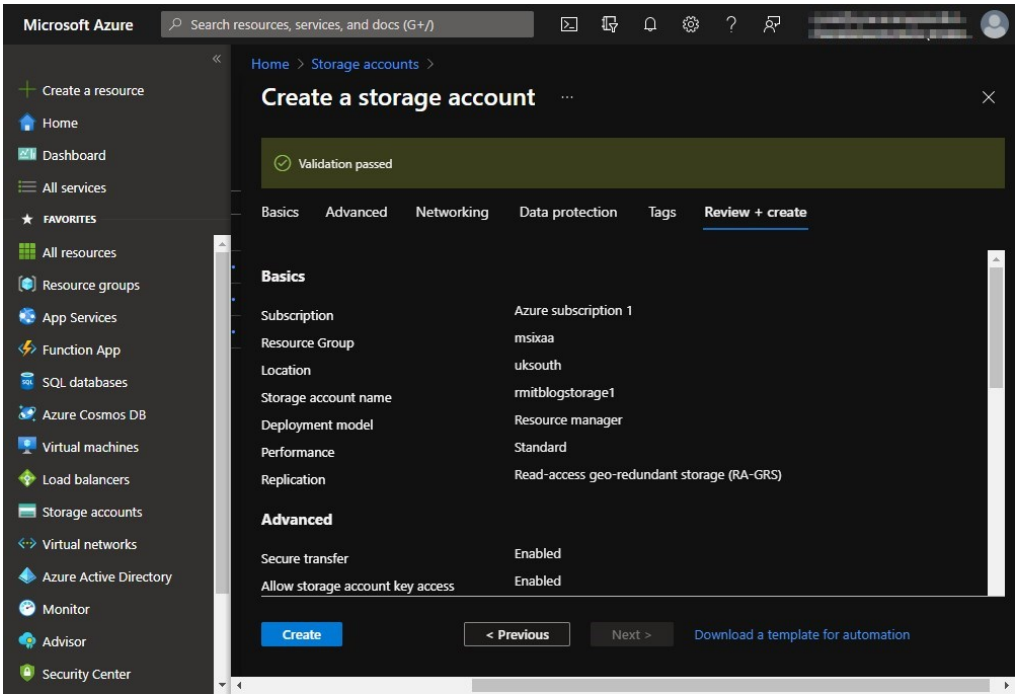

Figure 5.7 – The Review + create tab

Once the storage account has been created, you will see it appear on the **Storage accounts** page. In the next section, we will look at configuring an Azure file share.

Configuring file shares

Once you have created your storage account, you need to create a file share for FSLogix Profile Containers. This section will look at configuring a file share in a storage account ready for use with FSLogix.

Before we get started with Azure file shares, let's have a look at the different tiers that are available per share:

- **Premium** file shares use SSDs, which provide higher constant performance and lower latency than standard storage. This file share tier type is beneficial for larger shares or high I/O workload requirements.
- **Transaction** optimized file shares, similar to standard storage, use HDDs. This is suitable for heavy workloads but does not provide the required latency that premium file shares offer.
- **Hot** file shares provide storage optimized for general-purpose file sharing for items such as department shares. Hot files use HDDs.
- **Cool** file shares provide cost-effective storage for archive storage requirements. This type of storage tier uses HDDs.

> **Important Note**
>
> For larger organizations and high I/O workloads, it is recommended that you use the premium storage tier for Azure file shares.

Creating an azure file share is quite simple. You need to make sure you have created a storage account before proceeding. Within the storage account, you need to navigate to the **File Shares** icon within the table of contents for the storage account:

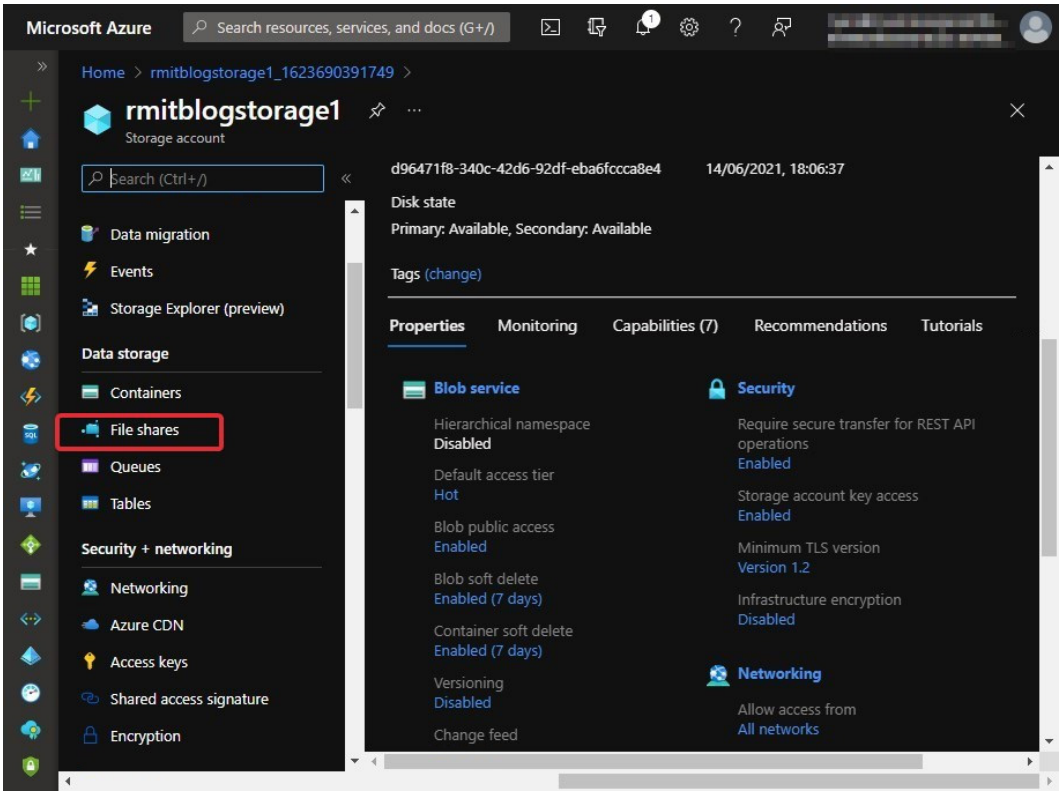

Figure 5.8 – The File shares link within the storage account

On the **File shares** page, click the **File share** button, as shown in the following screenshot:

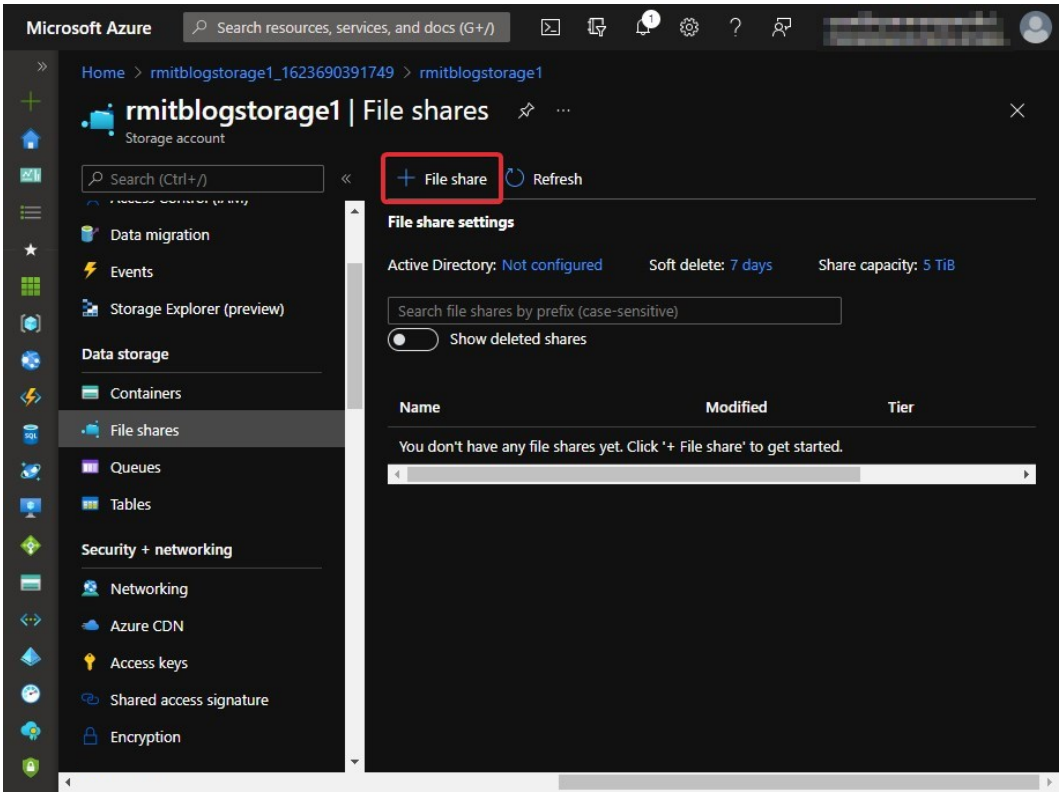

Figure 5.9 – The File share button

Once you have clicked the **File share** button, you will see the **New file share** blade appear. Fill in the following fields in this blade to create a new file share:

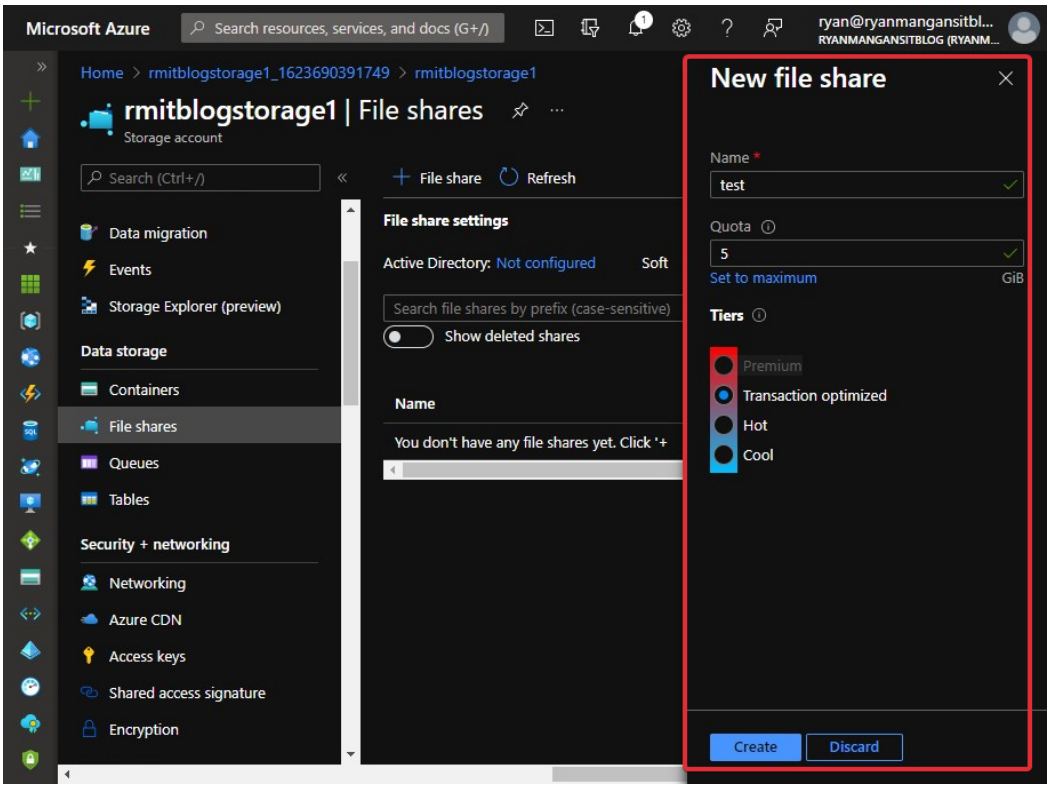

Figure 5.10 – The New file share blade

You will need to enter a **Name** for the share, a **Quota** size, and choose the tier you would like.

Once you have entered the required details, click **Create** to finish creating the new share:

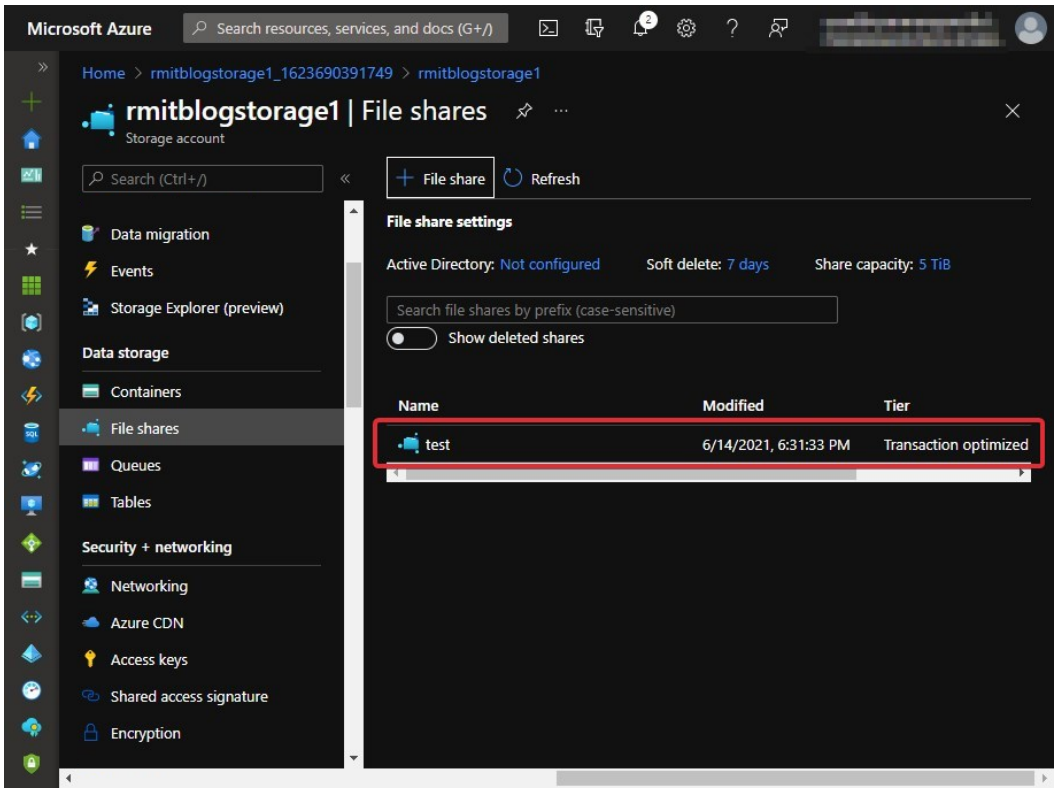

Figure 5.11 – The newly created file share

Note that the experience within a storage account using premium storage for file storage has a slightly different UI experience, as shown in the following screenshot:

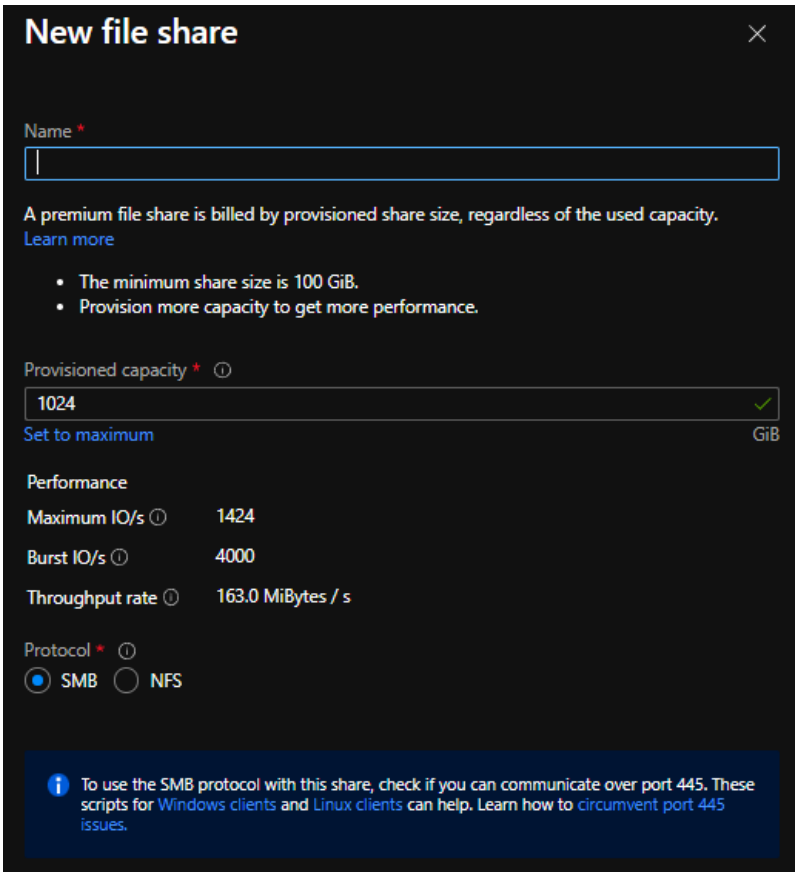

Figure 5.12 – Configuring premium file shares

This section summarized the different Azure file share storage tier options and how to create a new Azure file share. In the next section, we will look at Azure Managed Disks, ephemeral OS disks, and learn how to prepare a custom image.

Configuring disks

This section will look at Azure Managed Disks, the different available options, and how to prepare a custom VHD image.

An Azure managed disk is essentially a virtual disk (block-level storage volume) in conjunction with Azure VMs. Managed disks are designed to provide an availability of 99.999%. This is achieved by providing three replica copies of your data, which provides high durability.

The following table details the different types of managed disks that are available:

Detail	**Ultra Disk**	**Premium SSD**	**Standard SSD**	**Standard HDD**
Disk type	SSD	SSD	SSD	HDD
Scenario	I/O-intensive workloads such as SAP HANA, top tier databases (for example, SQL, Oracle), and other transaction-heavy workloads	Production and performance-sensitive workloads	Web servers, lightly used enterprise applications, and dev/test	Backup, non-critical, and infrequent access
Max disk size	65,536 gibibytes (GiB)	32,767 GiB	32,767 GiB	32,767 GiB
Max throughput	2,000 MB/s	900 MB/s	750 MB/s	500 MB/s
Max IOPS	160,000	20,000	6,000	2,000

This table was taken from the following site: `https://docs.microsoft.com/en-us/azure/virtual-machines/disks-types`.

As shown in the preceding table, each type of disk has a specific use case. For AVD multi-session deployments, it is recommended that you use premium SSDs to avoid any IOPs bottlenecks. You can use standard SSDs for personal desktop deployments. It is not recommended to use standard HDD disks for AVD deployments as performance could be degraded:

> **Important Note**
> It is recommended that premium SSDs be used for session hosts.

Premium SSD Sizes	P10	P15	P20	P30	P40	P50	P60	P70	P80
Disk size in GiB	128	256	512	1,024	2,048	4,096	8,192	16,384	32,767
Provisioned IOPS per disk	500	1,100	2,300	5,000	7,500	7,500	16,000	18,000	20,000
Provisioned Throughput per disk	100 MB/sec	125 MB/sec	150 MB/sec	200 MB/sec	250 MB/sec	250 MB/sec	500 MB/sec	750 MB/sec	900 MB/sec
Max burst IOPS per disk	3,500	3,500	3,500	30,000*	30,000*	30,000*	30,000*	30,000*	30,000*
Max burst throughput per disk	170 MB/sec	170 MB/sec	170 MB/sec	1,000 MB/sec*	1,000 MB/sec*	1,000 MB/sec*	1,000 MB/sec*	1,000 MB/sec*	1,000 MB/sec*
Max burst duration	30 min	30 min	30 min	Unlimited*	Unlimited*	Unlimited*	Unlimited*	Unlimited*	Unlimited*
Eligible for reservation	No	No	No	Yes, up to 1 year	Yes, up to 1 year	Yes, up to 1 year	Yes, up to 1 year	Yes, up to 1 year	Yes, up to 1 year

This table was taken from Microsoft's documentation site: `https://docs.microsoft.com/en-us/azure/virtual-machines/disks-types#premium-ssd-size`.

Typically, Azure Managed Disks are **locally redundant storage** (**LRS**). This means that the storage is replicated three times within a single data center in the region where you deployed the VM.

You can also configure **zone-redundant storage** (**ZRS**) for managed disks. ZRS replicates Azure Managed Disks synchronously across three Azure availability zones within a selected Azure region. Each zone is a separate physical location with independent networking, cooling, and power.

There is no difference in latency or performance; the only improvement when using ZRS is the improved data protection.

Ephemeral OS disks

Ephemeral OS disks, also known as stateless disk storage, are created on the Azure Hypervisor's local storage as part of the VM cache. One benefit of using ephemeral disks over Azure Managed Disks is that ephemeral disks are free. This allows the stateless disk storage to provide lower latency and faster reads and writes.

The following table details the differences between Azure Managed Disks and ephemeral disks:

	Azure Managed Disks	**Ephemeral OS Disks**
Size limit for OS disk	2 TiB.	Cache size for the VM size or 2 TiB, whichever is smaller. For the cache size in GiB, see DS, ES, M, FS, and GS.
VM sizes supported	All.	VM sizes that support (cache disk) premium storage such as DSv1, DSv2, DSv3, Esv3, Fs, FsV2, GS, and M.
Disk type support	Managed and unmanaged OS disk.	Managed OS disk only.
Region support	All regions.	All regions.
Data persistence	OS disk data that's written to OS disks is stored in Azure Storage.	Data written to OS disk is stored in the local Hypervisor storage and is not persisted to Azure Storage.
Stop-deallocated state	VMs and scale set instances can be stop-deallocated and restarted from the stop-deallocated state.	VMs and scale set instances cannot be stop-deallocated.
Specialized OS disk support	Yes.	No.
OS disk resize	Supported during VM creation and after the VM is stop-deallocated	Supported during VM creation only.
Resizing to a new VM size	OS disk data is preserved.	Data on the OS disk is deleted, OS is re-provisioned.
Page file placement	For Windows, the page file is stored on the resource disk.	For Windows, the page file is stored on the OS disk.

This table was taken from the following Microsoft site: `https://docs.microsoft.com/en-us/azure/virtual-machines/ephemeral-os-disks`.

Note that you cannot start and stop/deallocate an Azure VM that's been configured with an ephemeral OS disk (OS cache). The only options that are available to you are to restart or reimage.

> **Important Note**
>
> If you want to use ephemeral disks, you need to use a custom ARM template or third-party tooling and PowerShell.

In this section, we looked at what ephemeral disks are, the pros and cons, and the differences between Azure Managed Disks and ephemeral disks. In the next section, we will create a custom master VHD image.

Creating a VHD image

In this section, you will learn how to prepare a master **virtual hard disk** (**VHD**) image for Azure. Note that Microsoft recommends that you use an image from the Azure image gallery. However, this section covers both options, giving you the ability to customize an image offline and upload it to Azure when you're finished. You can also use Microsoft Deployment Toolkit and SCCM to create images for AVD. To upload these images, you can use the following tools:

- **Azure portal**: Use the upload feature within the storage account.
- **Azure Storage Explorer**: `https://azure.microsoft.com/features/storage-explorer/`
- **Az copy**: `https://docs.microsoft.com/azure/storage/common/storage-ref-azcopy`

> **Important Note**
>
> Ensure your image does not have the AVD agent installed on the VM. The agent can cause issues, including blocking registration and preventing user session connections.

Creating a VM

There are two options for creating a VM. First, you can provision the VM in Azure, and then customize and install the required software. Alternatively, you can create an image locally using Hyper-V and customize it to your requirements.

First, let's look at deploying a VM in Azure:

1. Within the Azure search bar, type `virtual`; the **Virtual machines** page link will be shown. Click on **Virtual machines**:

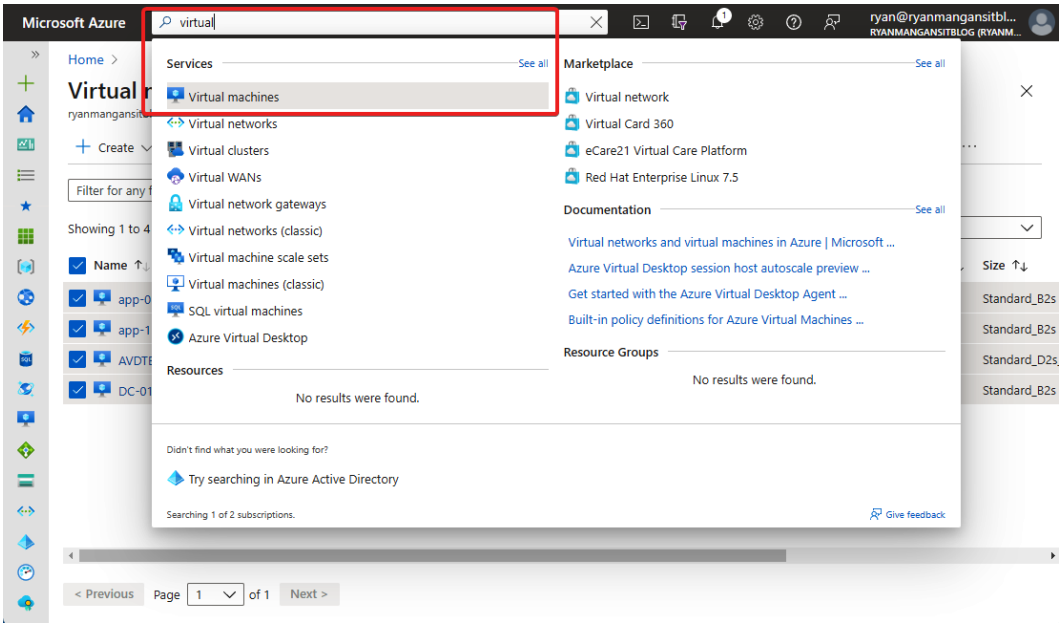

Figure 5.13 – Search bar displaying the Virtual machines page link in the Azure portal

2. Within the **Virtual machines** page, click **Create** and select **Virtual Machine**. This will open the **Create a virtual machine** page:

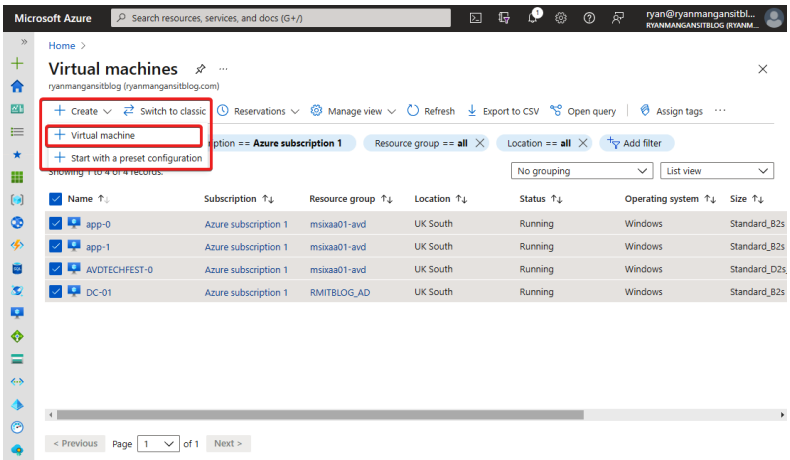

Figure 5.14 – Creating a VM within the Virtual machines page in the Azure portal

3. Within the **Create a virtual machine** page, you will need to fill in all the required fields:

   - Under the **Subscription** section, select the required subscription and select an existing **Resource group** or create a new one.
   - Under the **Instance details** section, provide a **Virtual machine name**, select a **Region**, select an **Image**, and specify a **Size**. This will be a VM skew:

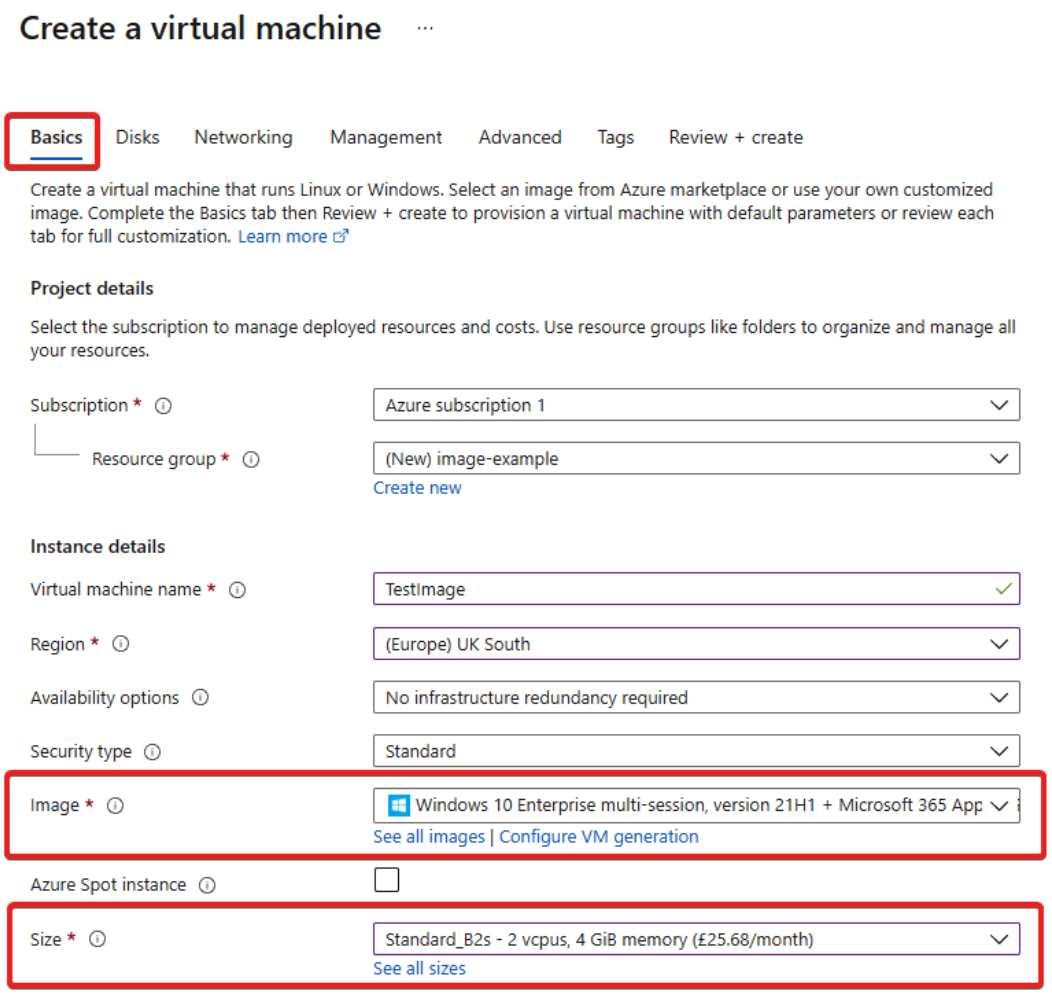

Figure 5.15 – The Basics tab within the Create a virtual machine page

4. Within the same tab, provide an administrator username and password.
5. Set the inbound port rules if required.
6. Check the **Licensing** check box to confirm that you have the correct licensing rights:

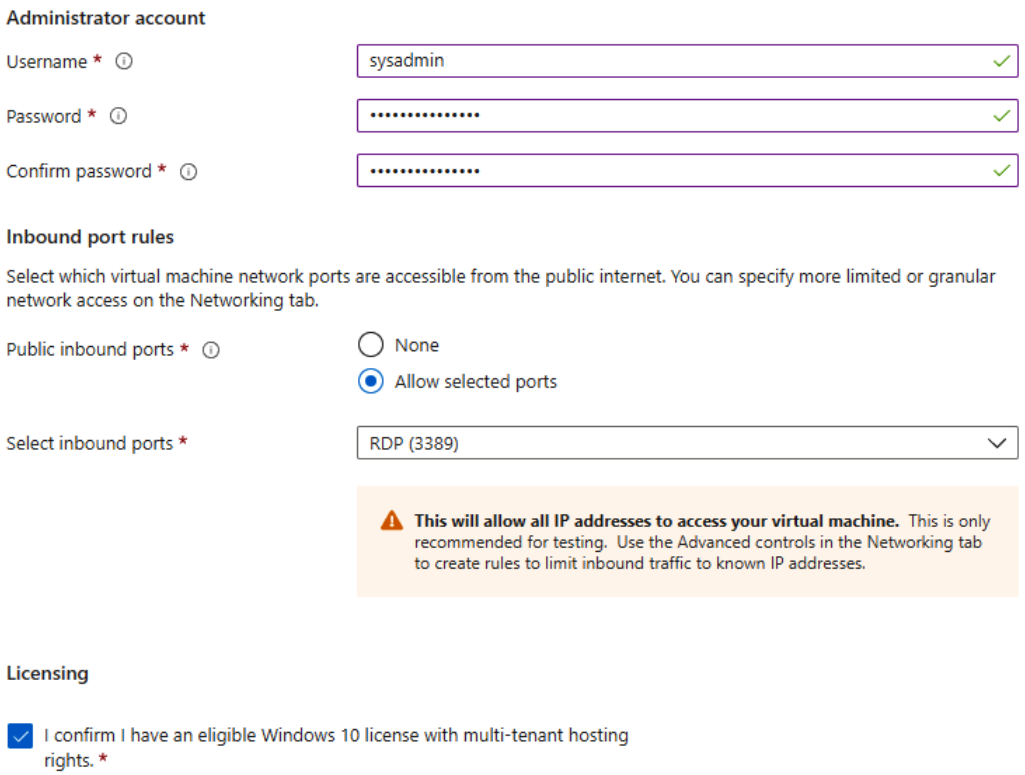

Figure 5.16 – The Administrator account section within the Basics tab of the Create a virtual machine page within the Azure portal

7. On the **Disks** tab, select the required disk. As we mentioned previously, a premium SSD is recommended:

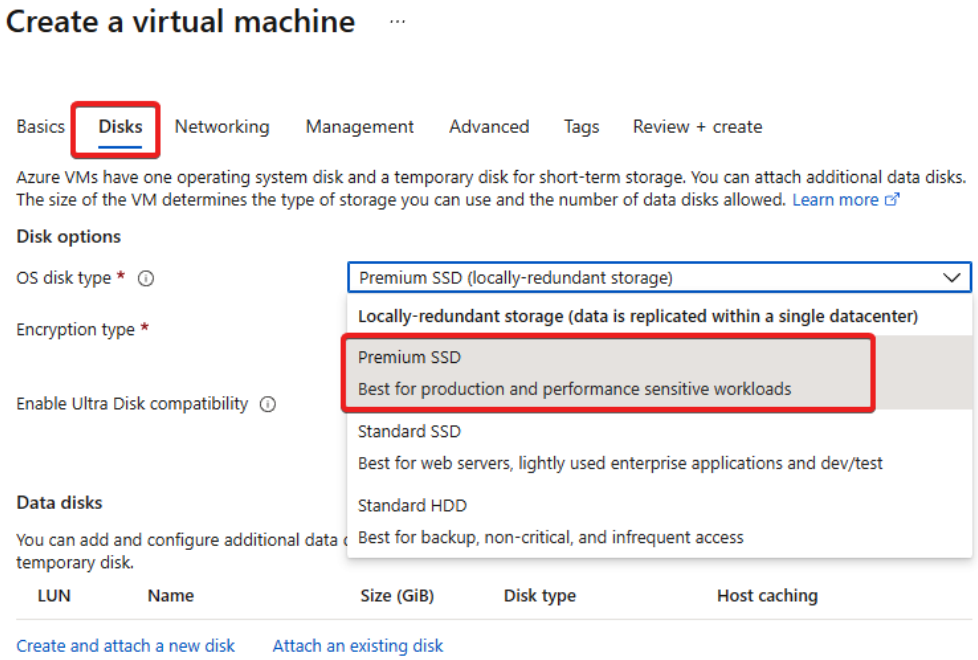

Figure 5.17 – The Disks tab within the Create a virtual machine page of the Azure portal

8. Once you have finished choosing the required disk and settings within the **Disks** tab, click the **Networking** tab and configure the required networking.

9. Under the **Networking** tab, configure the following:

   - Select the required **Virtual network**.
   - Select the required **Subnet**.
   - Set a public VM, if required.
   - Set the network security groups, if required:

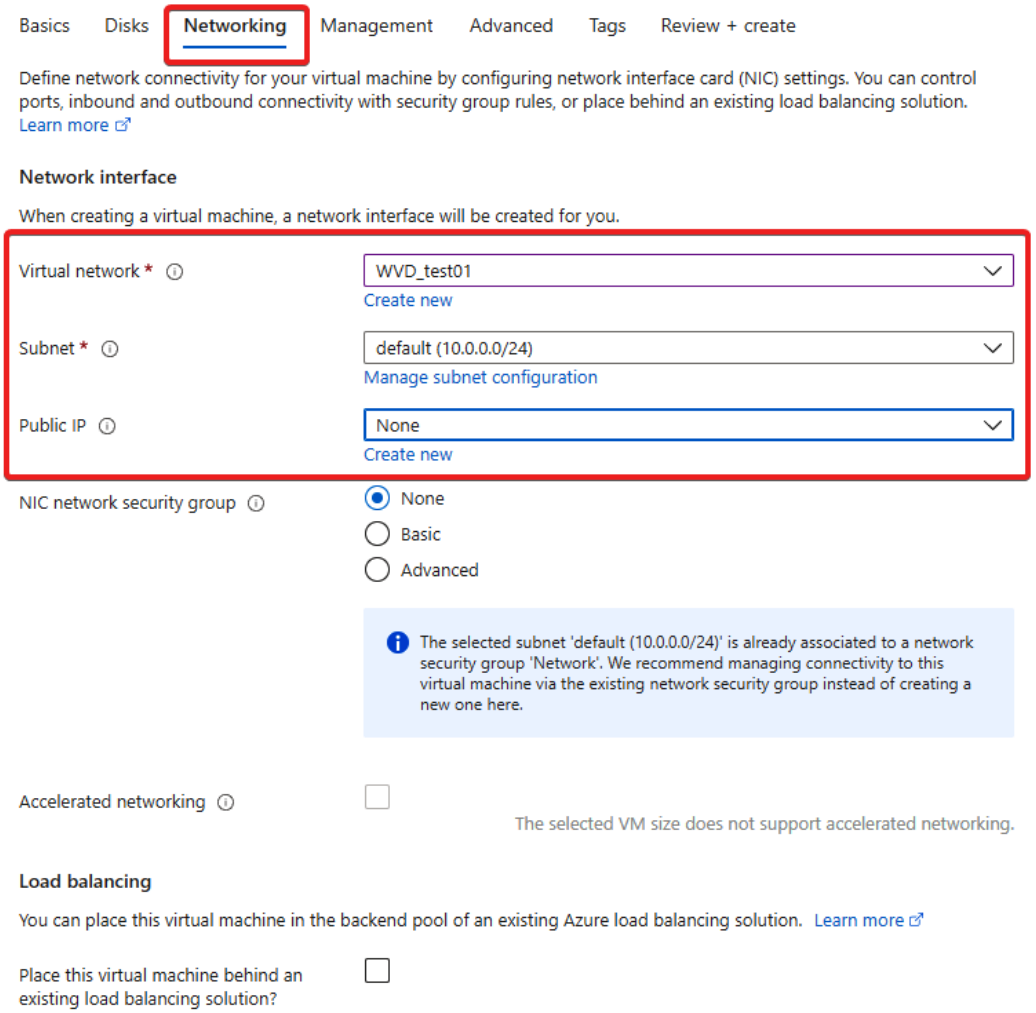

Figure 5.18 – The Networking tab within the Create a virtual machine page of the Azure portal

10. If you require specific settings under the **Management**, **Advanced**, and **Tags** tabs, complete the required settings and progress to the **Review + create** tab. If you do not require specific settings under these tabs, skip to the **Review + create** tab:

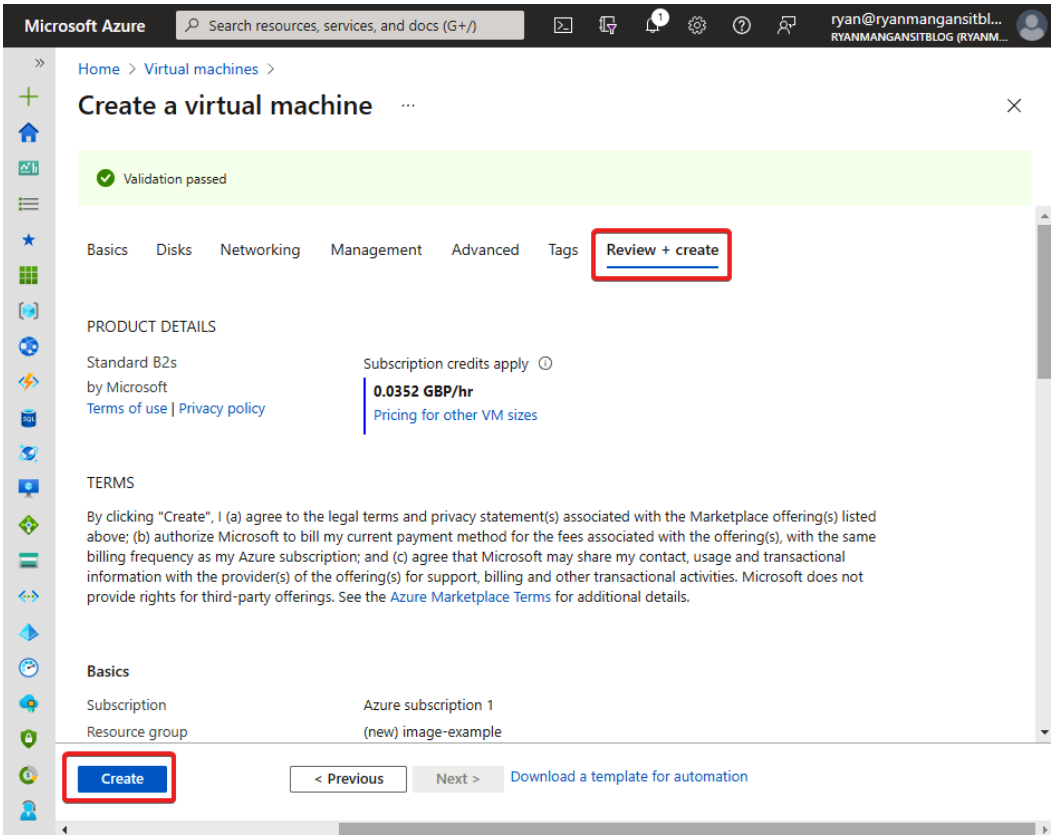

Figure 5.19 – The Review + create tab of the Create a virtual machine page within the Azure portal

This section showed you how to deploy a VM image template for AVD in the Azure portal. Next, we will learn how to create a local image on Hyper-V.

Creating a local image

First, you will need to download the required OS image. Then, using Hyper-V, you must create a VM using the downloaded VHD. You need to ensure that you complete the following steps:

1. Specify the generation as **Generation 1**:

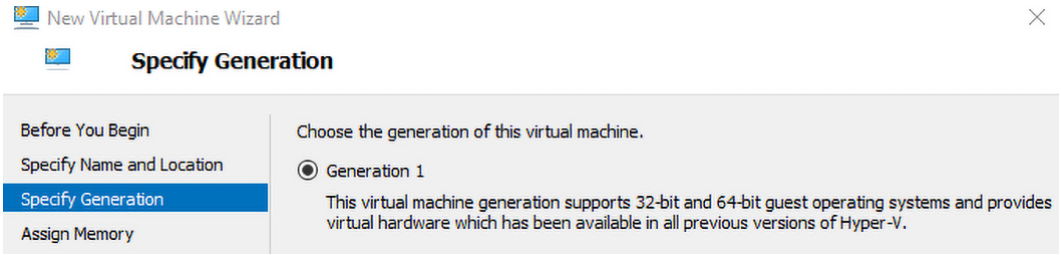

Figure 5.20 – Choosing Generation 1 in Hyper-V

2. Disable the checkpoints for the VM:

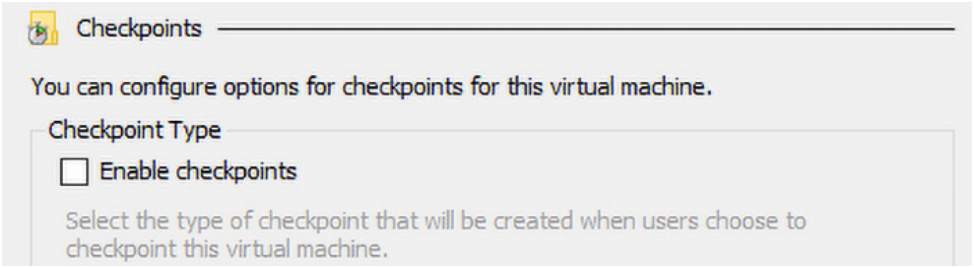

Figure 5.21 – Disabling the Enable checkpoints box

The following PowerShell cmdlet allows you to disable checkpoints:

```
Set-VM -Name <VMNAME> -CheckpointType Disabled
```

Now, let's look at the difference between dynamic and fixed disks since Azure only supports the fixed disk format.

Dynamic disks versus fixed disks

When creating a VM from an existing VHD, it creates a dynamic disk by default. However, you can change this by selecting the **Edit Disk…** option within Hyper-V.

You can also use PowerShell to change a dynamic disk to a fixed disk, as follows:

```
Convert-VHD –Path c:\test\MY-VM.vhdx –DestinationPath c:\test\
MY-NEW-VM.vhd -VHDType Fixed
```

This section detailed the options available to you when creating an image. We also covered some of the requirements for if you decide to customize an image outside of AVD using Hyper-V.

Summary

In this chapter, we looked at implementing and managing storage for AVD. First, we explored the requirements for storing FSLogix Profile Containers, storage account tiers, Azure Files storage tiers, and Azure Files integration with Active Directory Domain Services. Next, we looked at creating a new storage account and configuring Azure File Shares. Then, we reviewed the differences between Azure Managed Disks and ephemeral Operating System disks and finished by looking at the options available for creating a VM with Azure.

In the next chapter, we will look at creating and configuring host pools and session hosts.

Questions

Here are a few questions to test your understanding of this chapter:

1. What is the recommended storage solution for FSLogix Profile Containers?
2. Do all regions support all types of storage accounts and redundancy configurations?
3. When it comes to storage accounts for larger organizations and high I/O workloads, what is the recommended storage tier?
4. What is the recommended disk type for session hosts?
5. What disk format does a virtual hard disk need to be to upload and function correctly within Azure?