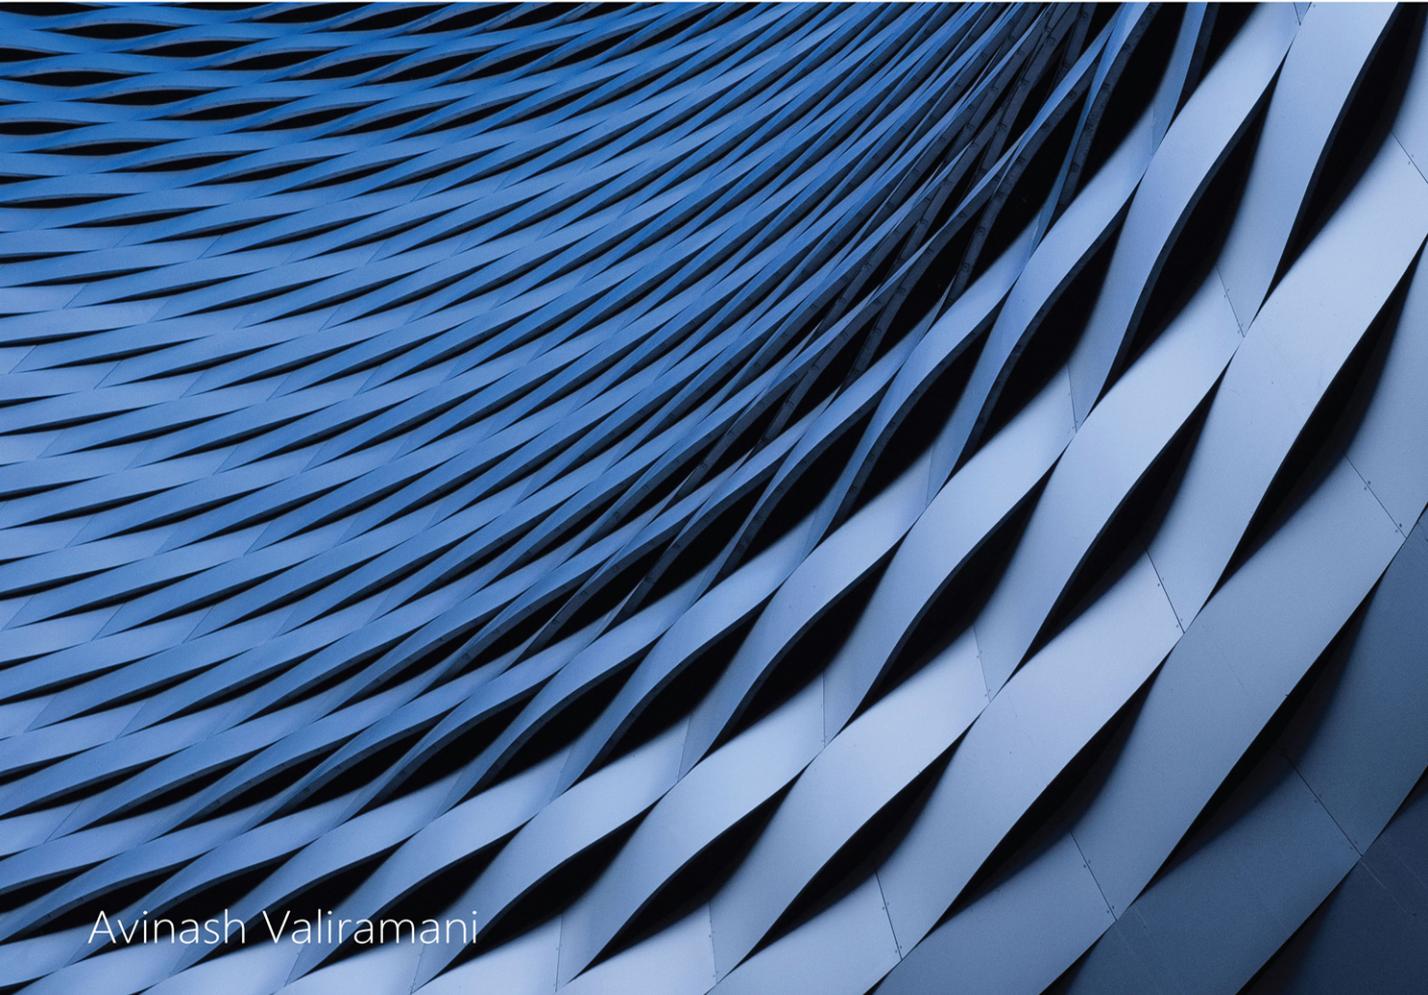




Microsoft Azure Storage

The Definitive Guide



Avinash Valiramani

FREE SAMPLE CHAPTER |



Microsoft Azure Storage: The Definitive Guide

Avinash Valiramani

Microsoft Azure Storage: The Definitive Guide

Published with the authorization of Microsoft Corporation by:
Pearson Education, Inc.

Copyright © 2024 by Pearson Education, Inc.

All rights reserved. This publication is protected by copyright, and permission must be obtained from the publisher prior to any prohibited reproduction, storage in a retrieval system, or transmission in any form or by any means, electronic, mechanical, photocopying, recording, or likewise. For information regarding permissions, request forms, and the appropriate contacts within the Pearson Education Global Rights & Permissions Department, please visit www.pearson.com/permissions.

No patent liability is assumed with respect to the use of the information contained herein. Although every precaution has been taken in the preparation of this book, the publisher and author assume no responsibility for errors or omissions. Nor is any liability assumed for damages resulting from the use of the information contained herein.

ISBN-13: 978-0-13-759318-7

ISBN-10: 0-13-759318-X

Library of Congress Control Number: 2023938511

\$PrintCode

TRADEMARKS

Microsoft and the trademarks listed at <http://www.microsoft.com> on the "Trademarks" webpage are trademarks of the Microsoft group of companies. All other marks are property of their respective owners.

WARNING AND DISCLAIMER

Every effort has been made to make this book as complete and as accurate as possible, but no warranty or fitness is implied. The information provided is on an "as is" basis. The author, the publisher, and Microsoft Corporation shall have neither liability nor responsibility to any person or entity with respect to any loss or damages arising from the information contained in this book or from the use of the programs accompanying it.

SPECIAL SALES

For information about buying this title in bulk quantities, or for special sales opportunities (which may include electronic versions; custom cover designs; and content particular to your business, training goals, marketing focus, or branding interests), please contact our corporate sales department at corpsales@pearsoned.com or (800) 382-3419.

For government sales inquiries, please contact governmentsales@pearsoned.com.

For questions about sales outside the U.S., please contact intlcs@pearson.com.

CREDITS

EDITOR-IN-CHIEF
Brett Bartow

EXECUTIVE EDITOR
Loretta Yates

DEVELOPMENT EDITOR
Kate Shoup

MANAGING EDITOR
Sandra Schroeder

SENIOR PROJECT EDITOR
Tracey Croom

COPY EDITOR
Sarah Kearns

INDEXER
Ken Johnson

PROOFREADER
Donna E. Mulder

TECHNICAL EDITOR
Thomas Palathra

EDITORIAL ASSISTANT
Cindy Teeters

COVER DESIGNER
Twist Creative, Seattle

COVER ILLUSTRATION
O.C Ritz / www.shutterstock.com

COMPOSITOR
codeMantra

GRAPHICS
codeMantra

Pearson's Commitment to Diversity, Equity, and Inclusion

Pearson is dedicated to creating bias-free content that reflects the diversity of all learners. We embrace the many dimensions of diversity, including but not limited to race, ethnicity, gender, socioeconomic status, ability, age, sexual orientation, and religious or political beliefs.

Education is a powerful force for equity and change in our world. It has the potential to deliver opportunities that improve lives and enable economic mobility. As we work with authors to create content for every product and service, we acknowledge our responsibility to demonstrate inclusivity and incorporate diverse scholarship so that everyone can achieve their potential through learning. As the world's leading learning company, we have a duty to help drive change and live up to our purpose to help more people create a better life for themselves and to create a better world.

Our ambition is to purposefully contribute to a world where:

- Everyone has an equitable and lifelong opportunity to succeed through learning.
- Our educational products and services are inclusive and represent the rich diversity of learners.
- Our educational content accurately reflects the histories and experiences of the learners we serve.
- Our educational content prompts deeper discussions with learners and motivates them to expand their own learning (and worldview).

While we work hard to present unbiased content, we want to hear from you about any concerns or needs with this Pearson product so that we can investigate and address them.

- Please contact us with concerns about any potential bias at <https://www.pearson.com/report-bias.html>.

Contents at a Glance

	<i>Acknowledgments</i>	<i>xii</i>
	<i>About the author</i>	<i>xiii</i>
	<i>Introduction to Microsoft Azure Storage</i>	<i>xiv</i>
Chapter 1	Azure Blob Storage	1
Chapter 2	Azure Files	79
Chapter 3	Azure Managed Disks	139
Chapter 4	Azure Queue Storage	193
Chapter 5	Azure Data Box	231
Chapter 6	Azure Data Share	251
	<i>Index</i>	<i>273</i>

Contents

<i>Acknowledgments</i>	<i>xii</i>
<i>About the author</i>	<i>xiii</i>
<i>Introduction to Microsoft Azure Storage</i>	<i>xiv</i>

Chapter 1	Azure Blob Storage	1
	Overview	1
	Key concepts	1
	Storage components	1
	Storage tiers	5
	Storage redundancy types	5
	Storage endpoints	10
	Storage encryption for at-rest data	10
	Storage data integrity	12
	Storage account walkthrough	12
	Data access authorization	24
	Azure Active Directory (Azure AD)	25
	Shared Key	26
	Shared access signature (SAS)	29
	Networking	33
	Network routing	33
	Network File System (NFS) 3.0 protocol	35
	SSH File Transfer (SFTP) protocol	36
	Storage account firewall and virtual networks	36
	Networking endpoints	41
	Storage access tiers	49
	Early deletion fees	51
	Default access tier configuration	51
	Blob lifecycle management	54
	Storage reservations	58

Static website hosting	58
Data protection	59
Soft delete for containers and blobs	59
Blob versioning	59
Blob change feed	60
Point-in-time restore	61
Data protection walkthrough	61
Azure Backup integration	65
Blob snapshots	70
Disaster recovery	73
Storage account failover	74
Last Sync Time	74
Best practices	75
Chapter 2 Azure Files	79
Overview	79
Key features	79
Key concepts	81
Deployment models	81
Storage accounts	82
File shares	90
Storage tiers for file shares	99
Networking considerations	101
Network protocols	101
Networking endpoints	103
Network routing	110
Encryption in transit	111
Storage account firewall	112
SMB Multichannel	116
Identity and access considerations	118
Identity and access considerations walkthrough	120
Data redundancy	122
Data protection	123

	Encryption for at-rest data	124
	Soft delete	124
	Share snapshots	124
	Azure Backup integration	127
	Best practices	137
Chapter 3	Azure Managed Disks	139
	Overview	139
	Key features.....	139
	Key concepts.....	140
	Disk roles	141
	Disk types	141
	Managed disk creation walkthrough	143
	Private Link integration	151
	Encryption	158
	Managed disk snapshots	158
	Managed images	167
	Performance tiering	172
	Disk redundancy	176
	Shared disks	177
	Managed disk bursting	180
	Managed disk backup	181
	Best practices	192
Chapter 4	Azure Queue Storage	193
	Overview	193
	Key features.....	194
	Key concepts.....	195
	Azure Storage account	195
	Queues and messages	202
	Networking considerations	208
	Storage firewall and virtual networks	209
	Private endpoints	213

Requiring secure transfers	220
Enforcing TLS versions	220
Identity and access considerations	220
Data redundancy	223
Disaster recovery	225
Storage account failover	225
Last Sync Time	226
Data encryption	227
Infrastructure encryption	227
Service-level encryption	228
Client-side encryption	228
Best practices	228

Chapter 5 Azure Data Box 231

Overview	231
Key features	231
Key concepts	232
Data Box components	232
Import/export workflow	233
Data security	234
Data-transfer speeds	235
Supported Azure services	236
Supported client operating systems	236
Availability	236
Data resiliency	236
Partner integrations	237
Preserving ACLs, file attributes, and timestamps	237
Limitations	238
Azure Data Box walkthrough	238
Data Box use cases	248
Best practices	249

Chapter 6	Azure Data Share	251
	Overview	251
	Key features.....	251
	Key concepts.....	252
	Data types	252
	Data provider	252
	Data consumer	252
	Sharing models	253
	Data stores	254
	Sharing caveats	254
	Managed identities	256
	Share and receive data with Azure Data Share	257
	Set up an Azure Data Share resource to share data walkthrough	257
	Set up an Azure Data Share resource to receive data walkthrough	265
	Best practices	271
	<i>Index</i>	273

Acknowledgments

At the outset, I want to express my deepest gratitude to Loretta Yates for bestowing upon me this tremendous responsibility. Only because of your unwavering trust and belief in my abilities, these books have come to fruition. I am forever grateful for the opportunity you have given me.

To my amazing mom, I am incredibly grateful for your unwavering support throughout the past two years as I wrote these books. Your love and understanding have meant the world to me. Thank you for being my rock.

To Celine, my sincere gratitude for being a constant source of guidance and assistance, whenever I needed you, throughout the journey of these last three books. Celine, thank you for your constant presence and encouragement. It has made this journey all the more meaningful.

To my beloved family, I am forever grateful for your understanding and patience during the countless hours I spent engrossed in writing these books.

To my extended family, thank you for tolerating my absence for over two years as I immersed myself in this writing endeavor. Hope to catch up with you all soon.

A heartfelt thank you goes to Kate Shoup for her exceptional editing and review work throughout all four books in the series. Your keen eye for detail and guidance throughout these books have been immeasurably valuable. Collaborating with you has been an enriching experience, and I am grateful for your exceptional skills as an editor.

I would also like to express my appreciation to Thomas Palathra, Sarah Kearns, and Tracey Croom for their meticulous contributions that brought this book to its completion. This endeavor has been a collective labor of love, and I am elated and grateful for our collaborative efforts.

Lastly, I extend my thanks to the entire Microsoft Press/Pearson team for their support and guidance throughout this project. Your expertise and guidance have been instrumental in shaping this book, and I am grateful for the opportunity to work alongside such a dedicated team.

Thank you all for being a part of this incredible journey. Your contributions and support have made these books a reality, and I am humbled and grateful for each and every one of you.

About the author

Avinash Valiramani is a highly experienced IT Infrastructure and Cloud Architect, specializing in Microsoft Technologies such as Microsoft Azure, Microsoft 365, Windows Server, Active Directory, Microsoft Exchange, SCCM, Intune, and Hyper-V. With over 17 years of expertise, he has worked with large and mid-size enterprises globally, designing their Cloud Architecture, devising migration strategies, and executing complex implementations. Avinash holds multiple certifications in Azure Infrastructure, Azure Artificial Intelligence, Azure Security, and Microsoft 365.

As part of the Microsoft Azure Best Practices series, Avinash is currently publishing four books, including this one, that draw from extensive real-world experiences. These books provide a comprehensive and concise resource for aspiring technologists and professionals alike. In addition to his Microsoft expertise, Avinash is also certified in Amazon AWS, Barracuda, Citrix, VMware, and other IT/Security industry domains, which further complements his skill set.

Avinash's contributions extend beyond writing books. He has authored an Azure Virtual Desktop course for O'Reilly Media and has plans for creating additional courses in the near future. You can stay updated with Avinash's insights and updates by following him on Twitter at @avaliramani. Furthermore, he will be sharing frequent blogs on his websites www.avinashvaliramani.com and www.cloudconsulting.services.

With his wealth of experience, industry certifications, and passion for advancing cloud technologies, Avinash Valiramani is a trusted advisor and sought-after resource in the realm of Microsoft Azure and Microsoft Office365. His expertise and dedication make him an invaluable asset for anyone seeking to leverage the full potential of the cloud.

Introduction to Microsoft Azure Storage

Welcome to *Microsoft Azure Storage: The Definitive Guide*. This book includes in-depth information about the various Azure services that provide storage capabilities and shares best practices based on real-life experiences with these services in different environments.

This book focuses primarily on Azure storage services generally available during 2022, encompassing development work done on these services over the years. A few storage features and functionalities were under preview at the time of this writing and could change before they are widely available; thus, we will cover the most notable ones in subsequent iterations of this book as they go live globally.

Overview

Over the years, Microsoft has introduced services related to the Azure storage stack to address various types of application and infrastructure requirements. Microsoft has released regular updates to these services, introducing additional features and functionality, enhancing each service's support matrix, and making these services easier to deploy and manage with each iteration.

Following is a brief timeline of the announcement of each of these services in public preview or general availability:

- **Azure Blob Storage** February 2010
- **Azure Queue Storage** February 2010
- **Azure Files** September 2015
- **Azure Managed Disks** February 2017
- **Azure Data Box** September 2017
- **Azure Data Share** July 2019

Each service provides customers with different options and features to address their storage requirements. This book dives into each of these services to highlight important considerations in deploying and managing them and to share associated best practices.

Each chapter focuses first on the features provided by a service. The chapter then explores in-depth the concepts behind that service and the components that comprise it so you will understand how that service can deliver value in your Azure deployment. Finally, each chapter focuses on deployment considerations and strategies where necessary, with step-by-step walkthroughs to illustrate deployment and management methods, followed by some best practices.

Cloud service categories

As in earlier books in this series, let's start by first discussing the different types of cloud service categories. Currently, cloud services are broken down into four main categories: infrastructure as a service (IaaS), platform as a service (PaaS), function as a service (FaaS), and software as a service (SaaS). SaaS is not relevant to the content covered in this Microsoft Azure book series; thus, we will focus on better understanding the first three categories:

- **Infrastructure as a service (IaaS)** Using virtual machines (VMs) with storage and networking is generally referred to as infrastructure as a service (IaaS). This is a traditional approach to using cloud services in line with on-premises workloads. Most on-premises environments use virtualization technologies such as Hyper-V to virtualize Windows and Linux workloads. Migrating to IaaS from such an environment is much easier than migrating to PaaS or FaaS. Over time, as an organization's understanding of various other types of cloud services grows, it can migrate to PaaS or FaaS.
- **Platform as a service (PaaS)** One of the biggest benefits of using a cloud service is the capability to offload the management of back-end infrastructure to a service provider. This model is called platform as a service (PaaS). Examples of back-end infrastructure include different layers of the application, such as the compute layer, storage layer, networking layer, security layer, and monitoring layer. Organizations can use PaaS to free up their IT staff to focus on higher-level tasks and core organizational needs instead of on routine infrastructure monitoring, upgrade, and maintenance activities. Azure Storage Service and Azure Data Share are examples of Azure PaaS offerings.

- **Function as a service (FaaS)** Function as a service (FaaS) offerings go one step beyond PaaS to enable organizations to focus only on their application code, leaving the entire back-end infrastructure deployment and management to the cloud service provider. This provides developers with a great way to deploy their code without worrying about the back-end infrastructure deployment, scaling, and management. It also enables the use of microservices architectures for applications. An example of an Azure FaaS offering is Azure Functions. There are no such examples for storage services.

In the Azure storage stack, some services fall under the PaaS category, including the following:

- **Azure Queue Storage** This PaaS service enables you to store large numbers of messages in a queue that can be ingested and processed by various application workloads.
- **Azure File Share** This PaaS service allows you to configure and manage SMB/NFS file shares in the Azure cloud platform and access them from Azure or on-premises environments.

Each cloud-service category has various features and limitations. Limitations might relate to the application, technological know-how, costs for redevelopment, among others. As a result, most organizations use some combination of different types of these cloud services to maximize their cloud investments.

Each service provides a different level of control and ease of management. For example:

- IaaS provides maximum control and flexibility in migration and use.
- FaaS provides maximum automation for workload deployment, management, and use.
- PaaS provides a mix of both at varying levels, depending on the PaaS service used.

Each service also offers varying levels of scalability or redundancy. For example:

- IaaS might require the use of additional services to achieve true geographical redundancy—for example, using Azure Site Recovery services, a PaaS service, to replicate Azure VMs and the underlying Azure managed disks across multiple Azure regions for redundancy and disaster recovery.
- PaaS and FaaS services are generally designed with built-in scalability and load-balancing features—for example, Azure Blob Storage with GRS redundancy level automatically replicates data to another Azure region.

Cost-wise, each service provides varying levels of efficiency. For example:

- FaaS offerings charge for compute based only on the usage hours for compute services, making them extremely cost-effective.
- IaaS offerings charge for compute services regardless of usage once the compute service (for example, a VM) is online.
- PaaS offerings are a mixed bag depending on how the services are configured. Some PaaS products charge for storage resources regardless of usage, while others, if configured correctly, charge based on usage alone. For example:
 - Azure standard file shares are charged based on the storage used to store the data in the primary region and secondary region, if configured for GRS.
 - Azure premium file shares are charged based on the storage allocated to store the data in the primary region and secondary region, if configured for GRS, regardless of the storage used.

Migration factors and strategies

Along with these features and limitations, there are certain migration factors to consider when deciding which category of cloud storage service might be the best solution in an organization's cloud journey. (See Figure I-1.) Of course, organizations can always start with one type of storage service and migrate to another type of storage service over time as their understanding of the cloud matures.

Let's examine the flow chart shown in Figure I-1 in more detail:

- **Lift-and-shift migration strategy** In a lift-and-shift migration, the organization migrates its existing on-premises environment as-is to the cloud, without redeveloping or redesigning the application stack. A lift-and-shift migration strategy generally involves less effort because no code changes are necessary. Application components remain as-is and are migrated in their current state to the cloud. This is a preferred migration approach for organizations in which:
 - A hardware refresh or procurement is planned.
 - Scaling or security limitations require the organization to migrate to the cloud as quickly as possible, with the least amount of disruption.
 - The organization wants to use IaaS mainly to host its application and database workloads.

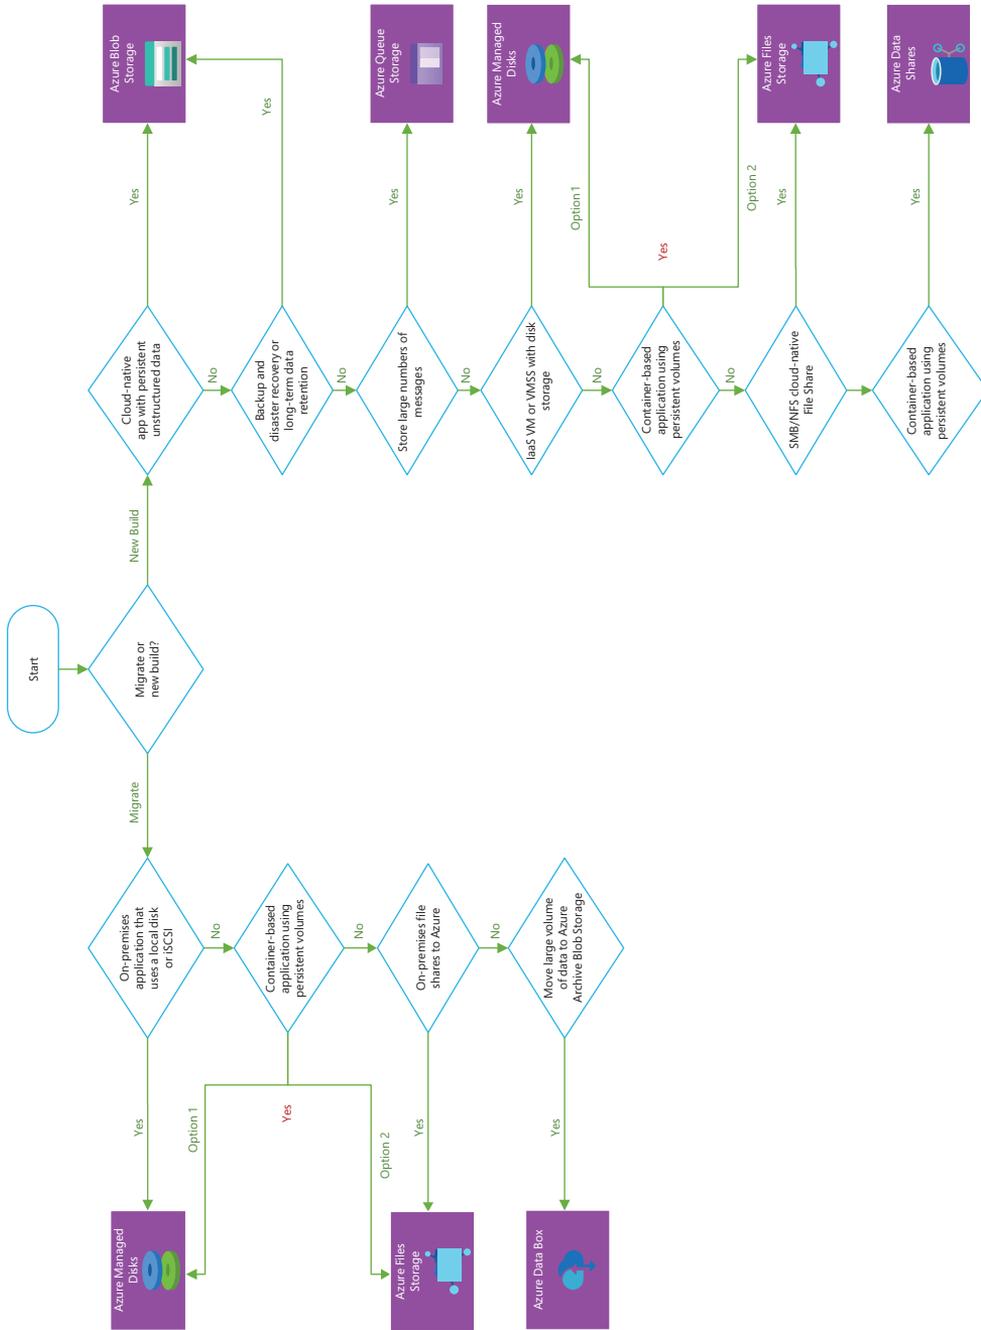


FIGURE I-1 Cloud-migration considerations.

- **Cloud-optimized strategy** With cloud-optimized migrations, the organization redesigns or recodes its application as necessary to use PaaS-based storage services. This enables the organization to use microservice architectures, allowing it to truly benefit from the scalability and cost benefits that a cloud service like Azure provides.

Organizations can use a lift-and-shift migration strategy, a cloud-optimized migration strategy, or a combination of the two. For example, an organization might use the flexibility provided by the Azure Managed Disks service to quickly migrate their existing on-premises VMs to Azure using a lift-and-shift approach to quickly benefit from the scaling and global availability of Azure. Then, over time, the organization could migrate to more cloud-optimized PaaS services, such as the Azure File Shares or Azure Blob Storage service, to meet those same needs.

Who is this book for?

Azure Storage: The Definitive Guide is for anyone interested in Azure infrastructure solutions—IT and cloud administrators, network professionals, security professionals, developers, and engineers. It is designed to be useful for the entire spectrum of Azure users. Whether you have basic experience using Azure or other on-premises or cloud virtualization technologies, or you are an expert, you will still derive value from this book. *Azure Storage: The Definitive Guide* provides introductory, intermediate, and advanced coverage of each widely used storage service.

The book especially targets those who are working in medium to large enterprise organizations; have at least basic experience in administering, deploying, and managing Azure infrastructure or other virtualization technologies such as Microsoft Hyper-V; and want to enhance their understanding of how to build resiliency and redundancy in their on-premises and cloud environments and to leverage the wide range of infrastructure services provided by Microsoft Azure.

How is this book organized?

This book is organized into six chapters:

- Chapter 1: Azure Blob Storage
- Chapter 2: Azure Files
- Chapter 3: Azure Managed Disks
- Chapter 4: Azure Queue Storage

- Chapter 5: Azure Data Box
- Chapter 6: Azure Data Share

Each chapter focuses on a specific Azure storage service, covering its inner workings in depth, with walkthroughs to guide you in building and testing the service and real-world best practices to help you maximize your Azure investments.

The approach adopted for the book is a unique mix of didactic, narrative, and experiential instruction:

- The didactic component covers the core introductions to the services.
- The narrative leverages what you already understand and acts as a bridge to introduce concepts.
- The experiential instruction takes into account real-world experiences and challenges in small and large environments and the factors to consider while designing and implementing workloads. Step-by-step walkthroughs on how to configure each Azure monitoring and management service and its related features and options enable you to take advantage of all the benefits each service has to offer.

System requirements

To get the most out of this book, your system must meet the following requirements:

- **An Azure subscription** Microsoft provides a 30-day USD200 trial subscription that can be used to explore most services covered in this book. Some services, such as dedicated hosts, cannot be created using the trial subscription, however. To test and validate these services, you will need a paid subscription. If you plan to deploy any of these restricted services, you will need to procure a paid subscription.
- **Windows 10/11** This should include the latest updates from Microsoft Update Service.
- **Azure PowerShell** For more information, see <https://docs.microsoft.com/en-us/powershell/azure/install-az-ps>.
- **Azure CLI** For more information, see <https://docs.microsoft.com/en-us/cli/azure/install-azure-cli>.
- **Display monitor** This must be capable of 1024 x 768 resolution.
- **Pointing device** You need a Microsoft mouse or compatible pointing device.

About the companion content

The companion content for this book can be downloaded from one of the following pages:

<https://MicrosoftPressStore.com/StorageTDG/downloads>

<https://github.com/avinashvaliramani/AzureStorageTDG>

The companion content includes the following:

- PowerShell code for each walkthrough in the book (where applicable)
- CLI code for each walkthrough in the book (where applicable)

Errata, updates, & book support

We've made every effort to ensure the accuracy of this book and its companion content. You can access updates to this book—in the form of a list of submitted errata and their related corrections—at:

MicrosoftPressStore.com/StorageTDG/errata

If you discover an error that is not already listed, please submit it to us at the same page.

For additional book support and information, please visit *MicrosoftPressStore.com/Support*.

Please note that product support for Microsoft software and hardware is not offered through the previous addresses. For help with Microsoft software or hardware, go to *<http://support.microsoft.com>*.

Stay in touch

Let's keep the conversation going! We're on Twitter: <http://twitter.com/MicrosoftPress>.

Azure Managed Disks

Overview

In February 2017, Microsoft announced the general availability for the Azure Managed Disks service, starting with the Standard and Premium disk types. Managed disks enable Azure customers to reduce overhead associated with managing and scaling storages account while creating or managing virtual machine (VM) disks. Microsoft also introduced numerous features that made managed disks a compelling solution for every Azure-hosted infrastructure as a service (IaaS) environment and for customers considering migrating to the cloud. Over time, the list of features and benefits associated with the Azure Managed Disks service has grown, and it has become the default disk solution for most organizations that use Azure for their VMs.

Each Azure managed disk is a fully managed block-level storage volume designed for the highest level of redundancy and availability. Azure currently offers different types of managed disks, including Ultra Disks, Premium SSD Disks, Standard SSD Disks, and Standard HDD disks. Each disk type provides varying levels of performance and scalability.

Key features

Some key features and benefits of using managed disks in your Azure environment include the following:

- **High availability, resiliency, and redundancy** Microsoft provides 99.999% availability for VM workloads that use managed disks. Managed disks are designed to maintain multiple replicas—three to be exact, spread across an Azure region. This makes managed disks extremely resilient, and ensures that your workload can continue to process even if there are issues with one or two replicas. Microsoft provides an industry leading 0% annualized failure rate.
- **High scalability** Microsoft currently supports the deployment of 50,000 managed disks per region per subscription, allowing large enterprises to deploy thousands of VMs in a single subscription.
- **Support for large Virtual Machine Scale Sets (VMSS)** You can use managed disks with VMSS. The scalability of managed disks makes it possible to deploy large VMSS consisting of up to 1,000 nodes.

- **Support for availability sets** Azure Managed Disks provides native integration with availability sets. Disks for VMs that are part of an availability set are spread across multiple fault domains with the selected Azure region and isolated from each other.
- **Support for availability zones** You can deploy managed disks across availability zones to improve redundancy. Availability zones provide additional redundancy over availability sets because the power and networking in each availability zone is independent of the others.
- **Support for existing virtual hard disks (VHDs)** You can easily upload existing VHDs up to 32 terabytes (TB) in size to Azure for use as managed disks. This process makes it extremely easy for organizations to migrate their existing workloads to Azure.
- **Role-based access control (RBAC)** Azure Managed Disks supports permission management using Azure RBAC, making it possible to granularly assign permissions to managed disks to administrators based on their roles and responsibilities.
- **Native integration with Azure Backup** You can use Azure Backup to back up managed disks from within the Azure Managed Disks service. You can schedule backups during off-peak hours and retain backups based on your organizational policies. You restore backups from the Azure Backup service.
- **Disk encryption** Managed disks are encrypted by default. They support multiple types of encryption, including Microsoft-managed encryption keys, customer-managed encryption keys, and double encryption with both types of keys. In addition, managed disks support Azure Disk Encryption, which allows you to encrypt the disk inside the VM using BitLocker for Windows or DM-Crypt for Linux VMs.
- **Easy migration for unmanaged disks** You can easily migrate unmanaged disks stored in Azure Storage accounts to managed disks. This increases the resiliency and redundancy of your IaaS VMs and provides significantly higher availability for your workloads.
- **Support for shared disks for clustered applications** You can set up managed disks as shared disks. This allows you to attach them to multiple VMs to host or migrate clustered applications to Azure.
- **Disk bursting for better performance** Managed disks allow you to increase the IOPS available for use for Premium and Standard SSD disks with on-demand or credit-based bursting models. Each model provides different capabilities to maximize the performance of your workloads when needed.
- **Private Link Support** You can use Private Link to import or export managed disks to or from Azure. This enables organizations to securely transfer disk data over a completely private connection.

Key concepts

Now that you have an initial understanding of the Azure Managed Disks service, let's spend some time going through all the different components and features in detail.

Disk roles

In Azure, disks play three primary roles:

- **Operating system (OS) disk** An OS disk is created by default for every VM you create in Azure. This disk contains the OS running on the VM as well as the boot volume. The OS disk supports partitioning with a master boot record (MBR) and GUID partition table (GPT) depending on the OS requirement. By default, most operating systems use partitioning with MBR, which limits the OS disk capacity to 2 TB. However, you can increase this to 4 TB by converting the disk from MBR to GPT.
- **Temporary disk** Microsoft provides a temporary disk as a non-persistent disk for specific VM models in Azure. When selecting the VM size in Azure, you can see the size of the temporary disk provided with that VM type. Any data you store on the temporary disk should be data that you are willing to lose, such as page files, swap files, or temporary logs. Each time a VM undergoes a forced restart, maintenance, or a redeployment, data on the temporary disk is erased. The VM can retain data stored on these disks only during standard reboot operations. Temporary disks are not encrypted by default, although you can enable encryption if needed. These disks are mapped as D: in Windows VMs and /dev/sdb in Linux-based VMs.
- **Data disk** Data disks are optional, and you can use them based on your workload requirements—for example, separating database installation files from data and log files, which can be stored on their own or individual data disks. As mentioned, OS disks have a maximum capacity of 4 TB, so any data-storage requirements that exceed that would require you to use data disks. The maximum disk capacity for a single data disk is currently 32,767 gigabytes (GB) for Standard HDD, Standard SSD, and Premium SSD disks. However, Ultra disks can be scaled up to 65,536 GB. The number and type of data disks that you can use with a VM depends on the size and type of the VM. Be sure to consider this when selecting the size for your VM.

NOTE Every VM has an OS disk. Whether a VM has a temporary disk depends on the VM model. Data disks are optional based on your workload requirements.

Disk types

Azure offers four types of disks:

- Standard HDD disks
- Standard SSD disks
- Premium SSD disks
- Ultra disks

Standard HDD disks

Standard HDD disks are suitable for workloads that are less critical and are not latency sensitive and for dev/test environments. These disks provide write latencies of less than 10 milliseconds (ms) and read latencies of less than 20 ms. Their performance varies depending on numerous factors, including IO size and workload pattern. Standard HDD disks are the least expensive (per gigabyte) disk option in Azure.

Standard SSD disks

Standard SSD disks are a great alternative for customers that want better performance, scalability, availability, and reliability than is possible with Standard HDD disks. Standard SSD disks are a great choice for low-intensity workloads that require consistent performance, such as web servers, low-usage business applications, and low IOPS applications. Standard SSD disks of 512 GB or more support credit-based bursting, making them ideal for applications that require a burst of performance only on rare occasions. All Azure VMs support Standard SSD disks.

Premium SSD disks

Premium SSD disks offer the second highest level of disk performance, with single-digit millisecond latencies, targeted IOPS, and defined throughput 99.9% of the time. They are suitable for high-intensity workloads, such as production applications and databases.

Premium SSD disks come in different sizes, and the level of IOPS support differs depending on the size of the Premium SSD disk. For example, P1 4 GB to P4 32 GB disks provide 120 IOPS, P10 128 GB disks provide 500 IOPS, while P80 32 TB disks provide 20,000 IOPS. Disk throughput and burst performance also increase as the capacity of the Premium SSD disks go up.

A few more features of Premium SSD disks are as follows:

- Premium SSD disks support one-year reservations to help you save on costs. You can set reservations for disks 1 TB and larger.
- Premium SSD disks support on-demand and credit-based bursting models. Bursting enables the Premium SSD to increase its performance in the short term to meet workload requirements.
- Only specific Azure VM types support Premium SSD disks. When you select a VM type, Azure shows you which types of disks that VM type supports. Because Microsoft adds and removes VM SKUs on an ongoing basis, I have not listed the VM types here, because they may change by the time you read this.

Ultra disks

Ultra disks currently provide the highest level of performance in terms of IOPS and disk throughput, with sub-millisecond latency 99.99% of the time. This makes Ultra disks suitable for critical high-performance workloads such as SAP HANA, mission-critical databases, and transaction-heavy applications.

By default, each Ultra disk can be scaled up to 32 TB. However, you can contact Azure support to request an increase of up to 64 TB. In terms of IOPS, each Ultra disk supports a minimum of 300 IOPS per gibibyte (GiB) and currently maxes out at 160,000 IOPS per disk.

Ultra disks allow you to adjust IOPS and throughput performance during runtime. You are permitted four adjustments every 24 hours. Each adjustment can take up to one hour to take effect and requires sufficient performance bandwidth capacity to prevent failures.

At present, Ultra disks have numerous limitations. These include lack of support for the following:

- Availability sets
- Azure Dedicated Host
- Disk snapshots
- Azure Backup
- Azure Site Recovery
- Disk exports
- VM image creation

In addition, Ultra disks cannot be used as OS disks. They can only be set up as data disks. For high-performance workloads that call for the use of an Ultra disk, you will want to set up the OS disk as a Premium SSD disk and leverage Ultra disks for all your workload data.

TIP Review the latest guidance available from Microsoft when planning your deployment, as these limitations may have changed by that time.

Managed disk creation walkthrough

The following sections step you through the process of creating a managed disk using the Azure portal, Azure PowerShell, and the Azure CLI.

IMPORTANT If you are following along, select resources and resource names based on your environment.

IMPORTANT If you are following along, be sure to delete any unwanted resources after you have completed testing to reduce charges levied by Microsoft.

USING AZURE PORTAL

To create a managed disk using the Azure portal, follow these steps:

1. Log in to the Azure portal, type **disks** in the search box, and select the **Disks** option in the list that appears. (See Figure 3-1.)

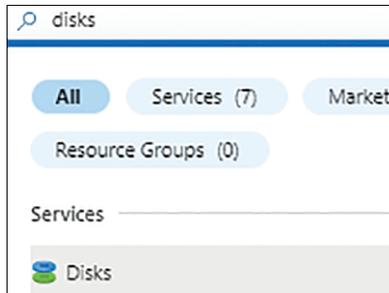


FIGURE 3-1 Searching for the Disks service in the Azure portal.

2. On the Disks page (see Figure 3-2), click **Create**.

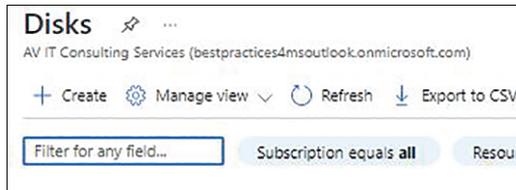


FIGURE 3-2 Creating a new disk.

3. In the **Basics** tab of the Create a Managed Disk wizard (see Figure 3-3), enter the following information:
 - **Subscription** Select the subscription in which you want to create the new managed disk.
 - **Resource Group** Select an existing resource group in which to create the new managed disk or create a new one.
 - **Disk Name** Enter a unique name for the managed disk.
 - **Region** Select the Azure region where you want to host the managed disk.
 - **Availability Zone** Select the availability zone you want to use or leave this option set to **None** (the default).
 - **Source Type** If the disk will be created from source data, such as a snapshot, storage blob, another disk, etc., select the source type.
4. To create a disk that is a different redundancy level, type, size, or performance tier from the default (1,024 GiB Premium SSD LRS), click the **Change Size** link in the **Size** section of the wizard's **Basics** tab.
5. In the Select a Disk Size dialog box, open the **Disk SKU** drop-down list and choose a disk type/redundancy level pairing. (See Figure 3-4.)

NOTE For more on redundancy levels for managed disks, see the section “Disk redundancy” later in this chapter.

The screenshot shows the 'Create a managed disk' wizard in the Basics tab. At the top, there are navigation tabs: Basics, Encryption, Networking, Advanced, Tags, and Review + create. Below the tabs is a paragraph of introductory text about disk availability and encryption. The 'Project details' section contains two dropdown menus: 'Subscription' set to 'Pay-As-You-Go' and 'Resource group' set to 'RG01'. The 'Disk details' section contains five dropdown menus: 'Disk name' set to 'ManagedDisk01', 'Region' set to '(US) East US 2', 'Availability zone' set to 'None', 'Source type' set to 'None', and 'Size' set to '1024 GiB'. The 'Size' dropdown also shows 'Premium SSD LRS' and a 'Change size' link.

FIGURE 3-3 The Basics tab of the Create a Managed Disk wizard.

The screenshot shows the 'Select a disk size' wizard. It starts with the instruction 'Browse available disk sizes and their features.' Below this is a 'Disk SKU' dropdown menu set to 'Premium SSD (locally-redundant storage)'. The main content area is a list of storage options, each with a description: 'Locally-redundant storage (data is replicated within a single datacenter)' with 'Premium SSD' (Best for production and performance sensitive workloads); 'Premium SSD v2' (Best for production and performance-sensitive workloads that consistently require low latency and high IOPS and throughput); 'Standard SSD' (Best for web servers, lightly used enterprise applications and dev/test); 'Standard HDD' (Best for backup, non-critical, and infrequent access); 'Zone-redundant storage (data is replicated to three zones)' with 'Premium SSD' (Best for the production workloads that need storage resiliency against zone failures) and 'Standard SSD' (Best for web servers, lightly used enterprise applications and dev/test that need storage resiliency against zone failures).

FIGURE 3-4 Choose a disk type and redundancy level.

- Click a size option in the list to select it. Alternatively, use the **Custom Disk Size (GiB)** and **Performance Tier** drop-down lists to choose a custom size/tier pairing. Then click **OK**. (See Figure 3-5.)

Select a disk size

Browse available disk sizes and their features.

Disk SKU: Premium SSD (locally-redundant storage)

Size	Disk tier	Provisioned IOPS	Provisioned throughput	Max Shares	Max burst IOPS	Max burst throughput
4 GiB	P1	120	25	3	3500	170
8 GiB	P2	120	25	3	3500	170
16 GiB	P3	120	25	3	3500	170
32 GiB	P4	120	25	3	3500	170
64 GiB	P6	240	50	3	3500	170
128 GiB	P10	500	100	3	3500	170
256 GiB	P15	1100	125	3	3500	170
512 GiB	P20	2300	150	3	3500	170
1024 GiB	P30	5000	200	5	-	-
2048 GiB	P40	7500	250	5	-	-
4096 GiB	P50	10000	250	5	-	-
8192 GiB	P60	16000	300	10	-	-
16384 GiB	P70	16000	750	10	-	-
32768 GiB	P80	20000	900	10	-	-

Custom disk size (GiB): 64

Performance tier: P6 - 240 IOPS, 50 MBps (default)

FIGURE 3-5 Selecting a different disk size and performance tier.

- Back in the **Basics** tab of the Create a Managed Disk wizard, click **Next**.
- In the **Encryption** tab of the Create a Managed Disk wizard (see Figure 3-6), open the **Key Management** drop-down list and choose **Platform-Managed Key**, **Customer-Managed Key**, or **Platform-Managed and Customer-Managed Keys**. Then click **Next**.

NOTE To use customer-managed keys, you must first generate and store the keys in the Azure Key Vault service.

Basics Encryption Networking Advanced Tags Review + create

Azure offers server-side encryption with platform-managed keys by default for managed disks. You may optionally choose to use a customer-managed key. [Learn more](#)

Key management: Platform-managed key

FIGURE 3-6 The Encryption tab of the Create a Managed Disk wizard.

- In the **Networking** tab of the Create a Managed Disk wizard (see Figure 3-7), in the **Network Access** section, leave the **Enable Public Access from All Networks** option button selected and click **Next**.
- In the **Advanced** tab of the Create a Managed Disk wizard (see Figure 3-8), enter the following information and click **Next**:
 - Enable Shared Disk** If you want to use this managed disk as a shared disk, select the **Yes** Option button. Then use the **Max Shares** drop-down list to specify how many VMs will share the disk.

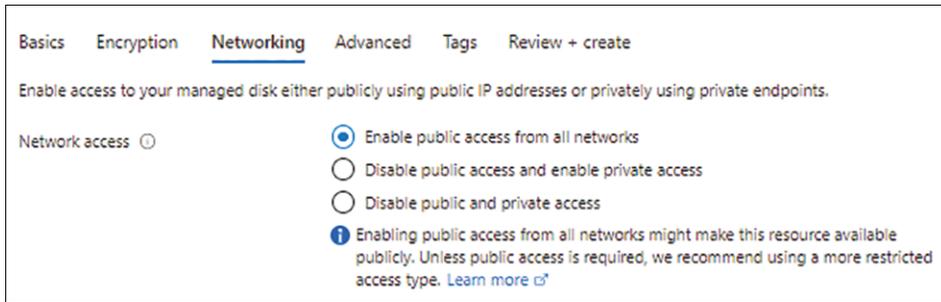


FIGURE 3-7 The Networking tab of the Create a Managed Disk wizard.

NOTE For more on shared disks, see the section “Shared disks” later in this chapter.

- **On-Demand Bursting** If you want this managed disk to be capable of on-demand bursting, select the **Enable On-Demand Bursting** check box.

NOTE The Enable On-Demand Bursting check box is available only if your managed disk is 512 GB or more. This option is covered in more detail later in this chapter.

- **Enable Data Access Authentication Mode** Optionally, select this check box to enable data access authentication. When you enable data access authentication, you can limit who can download the disk to admins who are authorized using Azure AD and authenticated using an approved account.

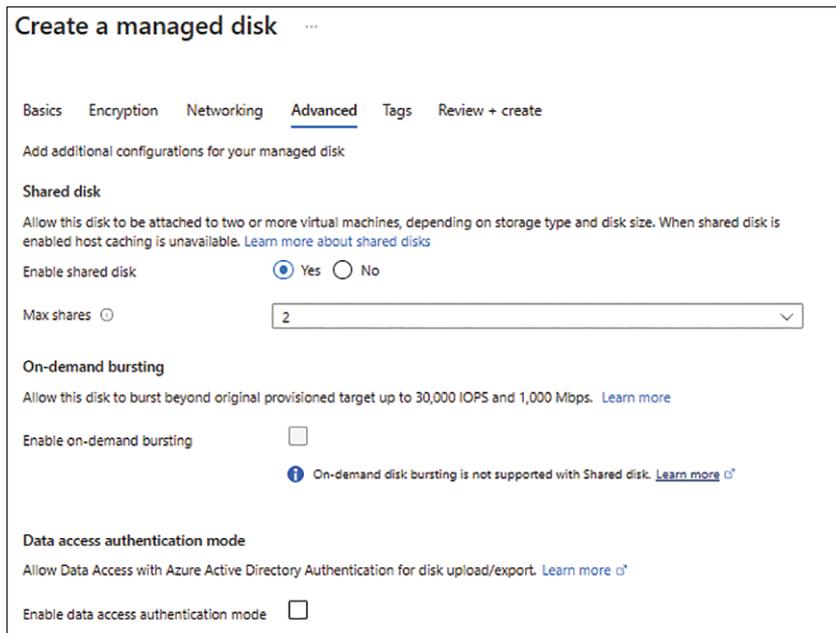


FIGURE 3-8 The Advanced tab of the Create a Managed Disk wizard.

11. In the **Tags** tab (see Figure 3-9), enter any tags you want to associate with the managed disk and click **Next**.

Basics Encryption Networking Advanced **Tags** Review + create

Tags are name/value pairs that enable you to categorize resources and view consolidated billing by applying the same tag to multiple resources and resource groups. [Learn more about tags](#)

Note that if you create tags and then change resource settings on other tabs, your tags will be automatically updated.

Name Value Resource

FIGURE 3-9 The Tags tab of the Create a Managed Disk wizard.

12. In the **Review + Create** tab (see Figure 3-10), review your settings, and click **Create** to create the managed disk.

Basics Encryption Networking Advanced Tags **Review + create**

Basics

Subscription Pay-As-You-Go

Resource group RG01

Region East US 2

Disk name ManagedDisk01

Availability zone None

Source type None

VM architecture x64

Size

Size 64 GiB

Performance tier P6 - 240 IOPS, 50 MBps (default)

Storage type Premium SSD LRS

Encryption

Encryption type Platform-managed key

Advanced

Enable shared disk No

Enable on-demand bursting No

Networking

Network access AllowAll

Tags

(none)

Create < Previous Next > Download a template for automation

FIGURE 3-10 The Review + Create tab of the Create a Managed Disk wizard.

13. After the managed disk is created, click **Go to Resource** to access its page. (See Figure 3-11.)

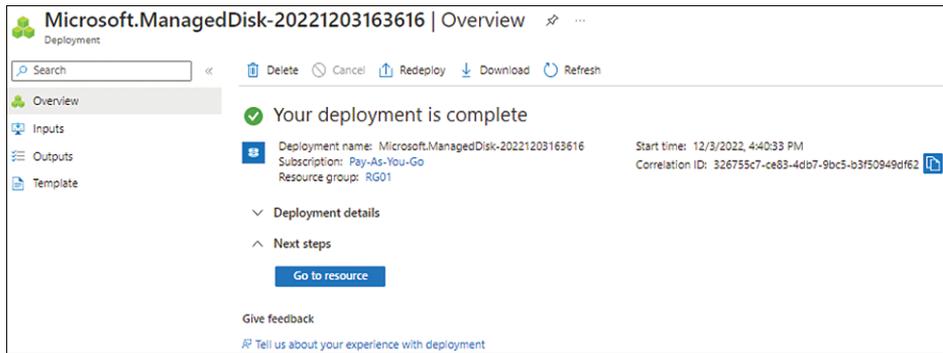


FIGURE 3-11 Managed disk deployment completion.

USING AZURE POWERSHELL

Use the following Azure PowerShell code to create a managed disk:

```
#Define variables
$resourceGroup = "RG01"
$location = "EastUS2"
$vm = "SourceVM"
$MgdDiskName = "ManagedDisk01"

#Create a disk config object - Change the disk redundancy as needed
$MgdDiskConfig = New-AzDiskConfig `
    -Location $location `
    -CreateOption Empty `
    -DiskSizeGB 64 `
    -EncryptionType EncryptionAtRestWithPlatformKey `
    -PublicNetworkAccess true `
    -Architecture X64 `
    -SkuName Standard_LRS/Premium_LRS/StandardSSD_LRS/UltraSSD_LRS/Premium_ZRS/
StandardSSD_ZRS

#Create Data Disk
$MgdDisk = New-AzDisk `
    -ResourceGroupName $resourceGroup `
    -DiskName $MgdDiskName `
    -Disk $mgddiskConfig

#Verify disk
Get-AzDisk `
    -ResourceGroupName $resourceGroup `
```

```

-DiskName $MgdDiskName

#Optional - Attach disk to VM
$Azvm = Get-AzVM `
  -ResourceGroupName $resourceGroup `
  -Name $vm

$Azvm = Add-AzVMDataDisk `
  -VM $vm `
  -Name $MgdDiskName `
  -CreateOption Attach `
  -ManagedDiskId $MgdDisk.Id `
  -Lun 1

Update-AzVM `
  -ResourceGroupName $resourceGroup `
  -VM $Azvm

```

USING AZURE CLI

Use the following code to create a managed disk in the Azure CLI:

```

#Define variables
resourceGroup="RG01"
location="EastUS2"
vm="SourceVM"
MgdDiskName="ManagedDisk01"

#Create managed disk - Change the disk redundancy as needed
az disk create \
    --resource-group $resourceGroup \
    --name $MgdDiskName \
    --size-gb 64 \
    --architecture x64 \
    --encryption-type EncryptionAtRestWithPlatformKey \
    --location $location \
    --public-network-access Enabled \
    --sku Premium_LRS/PremiumV2_LRS/Premium_ZRS/StandardSSD_LRS/StandardSSD_ZRS/
Standard_LRS/UltraSSD_LRS

#Verify disk
mgddisk=$(az disk show \
    --name $MgdDiskName \
    --resource-group $resourceGroup)

#Optional - Attach disk to VM
az vm disk attach \

```

```
--disks $mgddisk \  
--name $MgdDiskName \  
--resource-group $resourceGroup \  
--vm-name $vm
```

Private Link integration

Private Link provides secure connectivity to Azure PaaS services and Azure hosted services from your networks over a private endpoint. A private endpoint is a network interface connected to the Azure PaaS service or Azure hosted service, such as Managed Disks, that is attached to an Azure virtual network. With Private Link and private endpoints, you can safely and securely transfer managed disk files between regions using a private connection on the Microsoft backbone network instead of the public internet. You can also import VHD files from an on-premises environment directly to an empty managed disk in Azure over a private connection. Time-restricted Shared Access Signature (SAS) URLs can provide access to the unused managed disks and snapshots for transfer.

NOTE Another book in this series, *Microsoft Azure Networking: The Definitive Guide*, covers Private Link in detail in Chapter 10.

Private Link integration walkthrough

The following sections step you through the process of creating a private endpoint and integrating Private Link with the managed disk using the Azure portal and the Azure CLI.

IMPORTANT If you are following along, select resources and resource names based on your environment.

IMPORTANT If you are following along, be sure to delete any unwanted resources after you have completed testing to reduce charges levied by Microsoft.

USING AZURE PORTAL

To create a private endpoint and integrate Private Link with a managed disk using the Azure portal, follow these steps:

1. Log in to the Azure portal, type **disk accesses** in the search box, and select the **Disk Access** option from the list that appears. (See Figure 3-12.)

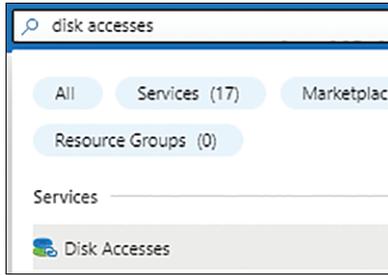


FIGURE 3-12 Searching for disk accesses in the Azure portal.

2. On the Disk Access page, click **Create Disk Access**. (See Figure 3-13.)

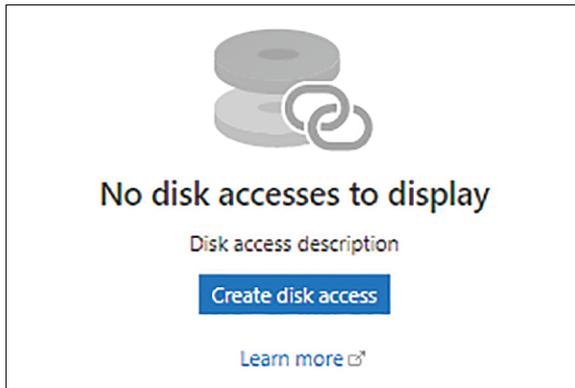


FIGURE 3-13 Create disk access.

3. In the **Basics** tab of the Create a Disk Access wizard (see Figure 3-14), enter the following information:
 - **Subscription** Select the subscription in which you want to create the disk access resource.
 - **Resource Group** Select an existing resource group in which to create the disk access resource or create a new one.
 - **Name** Enter a unique name for the disk access resource.
 - **Region** Select the Azure region where you want to host the disk access resource.Before you continue with the Create a Disk Access wizard, you need to create the private endpoint. You'll do that next.
4. At the bottom of the **Basics** tab, click **Add**.
5. In the Create a Private Endpoint dialog box (see Figure 3-15), enter the following information and click **OK**:
 - **Subscription** Select the subscription you want to use to create the private endpoint.

Create a disk access ...

Basics Tags Review + create

Private links provide protection from a SAS URI being available to anyone by locking down the access to a specific virtual network. This private "tunnel" gives users the security they need when sensitive data is contained on the disks or snapshots. [Learn more](#)

Project details

Select a subscription to manage deployed resources and costs. Use resource groups like folders to organize and manage all your resources.

Subscription *

Resource group * [Create new](#)

Instance details

Name * ✓

Region *

Private endpoint

Name	Subscription	Resource group	Region	Target resource
<i>Click on add to create a private endpoint</i>				
<input type="text"/>				
+ Add				

FIGURE 3-14 The Basics tab of the Create a Disk Access wizard.

- **Resource Group** Select an existing resource group in which to create the private endpoint or create a new one.
- **Location** Select the Azure region where you want to host the private endpoint.
- **Name** Enter a unique name for the private endpoint.
- **Target Resource** Select **Disks**.
- **Virtual Network** Select the virtual network on which to create the private endpoint.
- **Subnet** Select the subnet on which to create the private endpoint.
- **Integrate with Private DNS Zone** Select **Yes** to integrate with a private DNS zone or select **No** if you plan to create a DNS record in your own DNS servers or on the host files of the workloads VMs. In this case, select **Yes**.
- **Private DNS Zone** Select the private DNS zone with which you want to integrate the private endpoint. In this case, leave it set to the default, **privatelink.blob.core.windows.net**.

Create private endpoint ✕

Subscription * ⊙

Resource group * ⊙ [Create new](#)

Location *

Name * ⊙ ✓

Target resource * ⊙

Networking

To deploy the private endpoint, select a virtual network subnet. [Learn more about private endpoint networking](#) ↗

Virtual network * ⊙

Subnet * ⊙ ✓

i If you have a network security group (NSG) enabled for the subnet above, it will be disabled for private endpoints on this subnet only. Other resources on the subnet will still have NSG enforcement.

Private DNS integration

To connect privately with your private endpoint, you need a DNS record. We recommend that you integrate your private endpoint with a private DNS zone. You can also utilize your own DNS servers or create DNS records using the host files on your virtual machines. [Learn more about private DNS integration](#) ↗

Integrate with private DNS zone ⊙ Yes No

Private DNS Zone * ⊙

FIGURE 3-15 The Create Private Endpoint dialog box.

6. Click the **Tags** tab (see Figure 3-16), enter any tags you want to associate with the private endpoint, and click **Next**.
7. In the **Review + Create** tab (see Figure 3-17), review your settings and click **Create** to create the private endpoint.

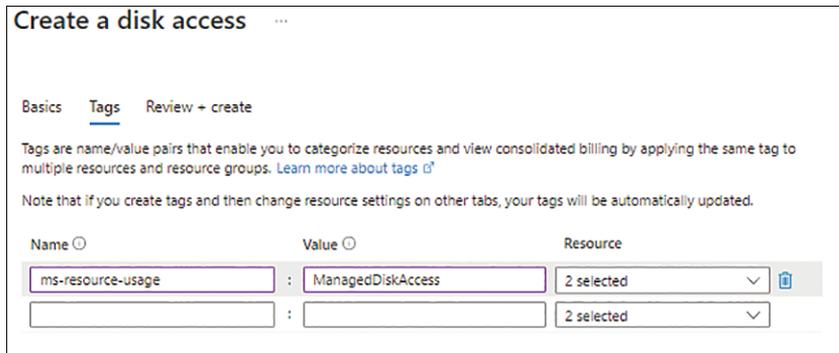


FIGURE 3-16 The Tags tab of the Create a Disk Access wizard.

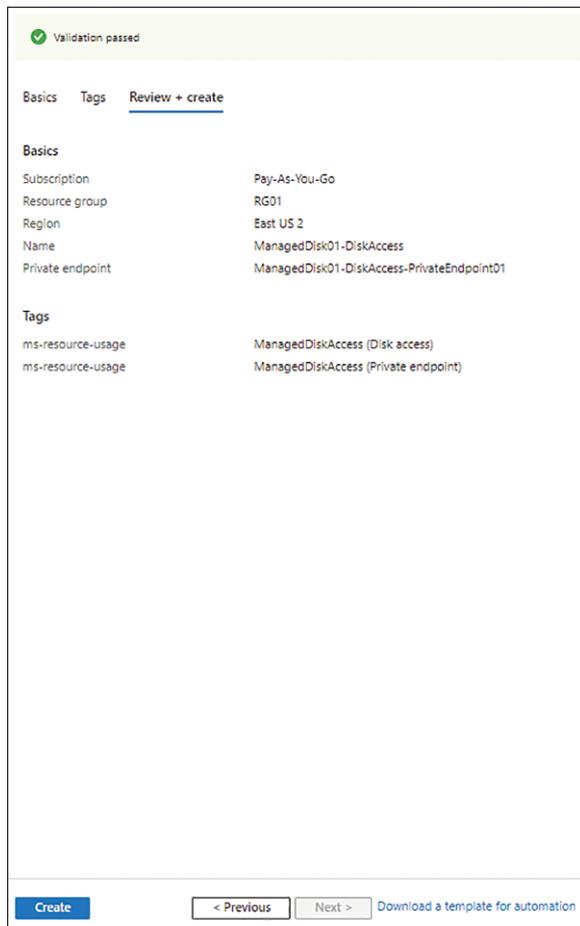


FIGURE 3-17 The Review + Create tab of the Create a Disk Access wizard.

- After the private endpoint is created, click **Go to Resource** to access its page. (See Figure 3-18.)

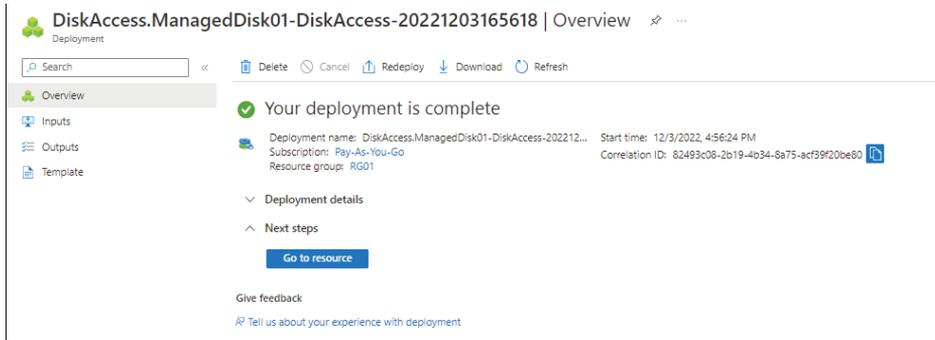


FIGURE 3-18 Private endpoint deployment completion.

9. In the left pane of the page for the managed disk you created earlier, under **Settings**, click **Networking**.
10. On the managed disk's Networking page (see Figure 3-19), perform the following steps and click **Save**:
 - **Network Access** Select the **Disable Public Access and Enable Private Access** option button.
 - **Disk Access** Select the private endpoint you just created.

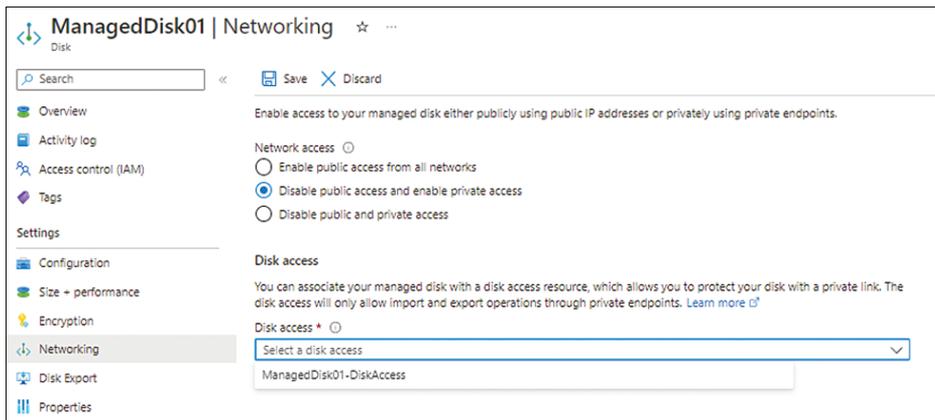


FIGURE 3-19 The managed disk's Networking page.

USING AZURE CLI

Use the following code to create a private endpoint and integrate Private Link with a managed disk in the Azure CLI:

```
#Define variables
resourceGroup="RG01"
```

```

location="EastUS2"
vm="SourceVM"
MgdDiskName="ManagedDisk01"
diskAccess="ManagedDisk01-DiskAccess"
vnet="VNET-01"
subnet="default"
privateEndPoint="ManagedDisk01-DiskAccess-PrivateEndpoint01"
#Create disk access
az disk-access create \
    --name $diskAccess \
    --resource-group $resourceGroup \
    --location $location

diskAccessId=$(az disk-access show \
    --name $diskAccess \
    --resource-group $resourceGroup \
    --query [id] -o tsv)

#Create private endpoint
az network private-endpoint create
    --resource-group $resourceGroup \
    --name $privateEndPoint \
    --vnet-name $vnet \
    --subnet $subnet \
    --private-connection-resource-id $diskAccessId \
    --group-ids disks \
    --connection-name $privateEndPoint

#Create Private DNS zone config
az network private-dns zone create \
    --resource-group $resourceGroup \
    --name "privatelink.blob.core.windows.net"

az network private-dns link vnet create \
    --resource-group $resourceGroup \
    --zone-name "privatelink.blob.core.windows.net" \
    --name $privateEndPoint-DNSLink \
    --virtual-network $vnet \
    --registration-enabled false

az network private-endpoint dns-zone-group create \
    --resource-group $resourceGroup \
    --endpoint-name $privateEndPoint \
    --name $privateEndPoint-ZoneGroup \

```

```

--private-dns-zone "privatelink.blob.core.windows.net" \
--zone-name disks

#Update managed disk with Private Link config
diskAccessId=$(az resource show \
  --name $diskAccess \
  --resource-group $resourceGroup \
  --namespace Microsoft.Compute \
  --resource-type diskAccesses \
  --query [id] -o tsv)

az disk update \
  --name $diskName \
  --resource-group $resourceGroup \
  --network-access-policy AllowPrivate \
  --disk-access $diskAccessId

```

Encryption

Managed disks support two types of disk encryption:

- **Server-Side Encryption (SSE)** SSE manages encryption on the storage layer and is handled by the Azure Storage service. It provides encryption-at-rest and during write operations to the underlying storage, thereby ensuring that disks stored in Azure are not readable in the event of data theft. SSE is enabled by default for all managed disks, snapshots, and images across all Azure regions. SSE supports two types of key management: Azure platform-managed keys or customer-managed keys. You can choose which type of key management you want to use for each managed disk you create.
- **Azure Disk Encryption (ADE)** ADE refers to encryption within the system. It applies to the OS and data disks in an Azure IaaS VM. ADE encryption is performed using BitLocker technology in Windows and DM-Crypt technology in Linux. In both scenarios, the keys are integrated and stored in Azure Key Vault to make it easier for you to manage them.

Managed disk snapshots

Snapshots provide an easy way to back up a point-in-time copy of your managed disk for restore or cloning operations. Snapshots are read-only, crash-consistent copies of the disk. You can use them to create new managed disks without affecting the source managed disk in any way. Snapshots are, by default, stored as standard managed disks, but you can change this during the snapshot creation process.

IMPORTANT While snapshots serve as a great way to create a copy of a managed disk, they are not a replacement for regular backups, and you should not use them as such.

The first time you take a snapshot of a managed disk, it will be a full snapshot. Subsequent snapshots, however, can be incremental. An incremental snapshot captures all changes to the managed disk since the last snapshot of the disk. This reduces your storage footprint. If you need to restore from a single incremental snapshot, Azure automatically identifies all the incremental and full snapshots preceding the current one to reconstruct the entire disk. This makes incremental snapshots extremely cost-effective, making them the preferred option for regular snapshot management.

NOTE If the zone in which the incremental snapshot is created provides ZRS redundancy capabilities, then the incremental snapshot will automatically be saved with ZRS, too, unless specified otherwise.

NOTE If you are using full snapshots on premium storage to scale up VM deployments, we recommend you use custom images on standard storage in the Shared Image Gallery. This will help you achieve a more massive scale with a lower cost. For more on this, see Chapter 2, “Virtual Machine Scale Sets,” in *Microsoft Azure Compute: The Definitive Guide*.

Incremental snapshots can also be useful for disaster recovery between Azure regions—that is, you can identify changes between two snapshots of the same disk, and then transfer only the differential changes to the secondary region instead of the entire snapshot. Then, when you restore/rebuild in the secondary region, you can use the snapshot of the base blob of the managed disk in combination with these differential changes. (See Figure 3-20.) This strategy can reduce time, costs, and network requirements for disaster recovery for managed disks.

NOTE Microsoft provides sample .NET code online to help you test this capability if you are interested in exploring it.

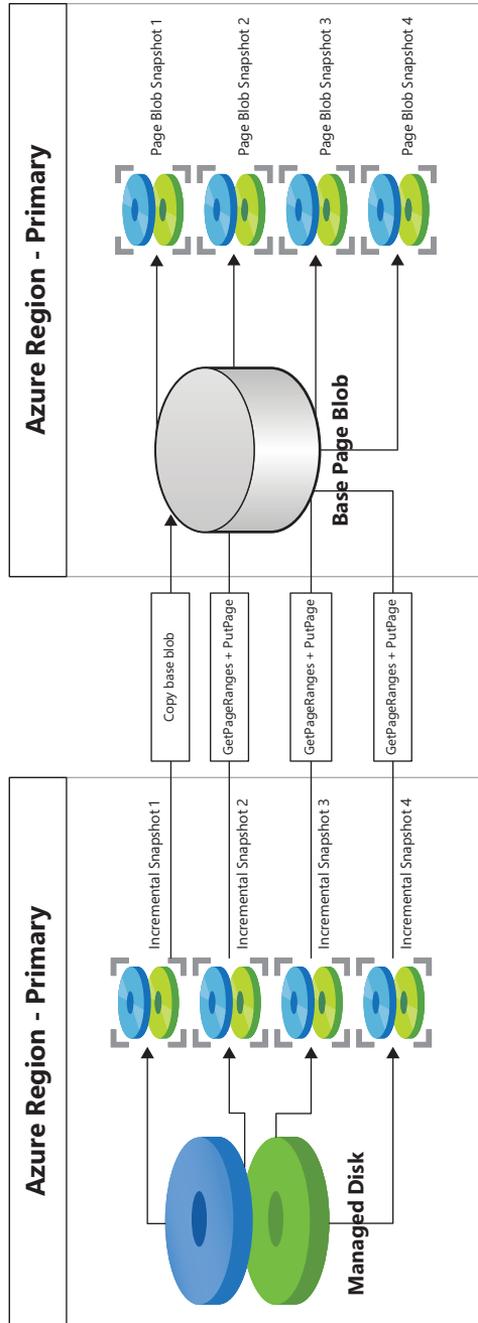


FIGURE 3-20 Incremental snapshots.

Incremental snapshots are a great feature, but they do have some limitations that exist at the time of this writing. By the time you read this, these limitations may have been addressed. Be sure to review Microsoft's latest guidance before finalizing your snapshot management strategy. Some key limitations at present include the following:

- Unlike full snapshots, incremental snapshots always use Standard HDD disks, regardless of the disk type used for the full snapshot.
- A single managed disk supports a maximum of 500 incremental snapshots.
- Each managed disk limits you to creating seven incremental snapshots, with a wait time of 5 minutes between each snapshot.
- The managed disk and snapshots must all be part of the same subscription.
- If you want to move a managed disk to another subscription, you will not be able to do so if the disk has incremental snapshots. You will need to keep this in mind when planning any such migrations.
- Differentials do not work for disks larger than 4 TB.

Managed disk snapshots walkthrough

The following sections step you through the process of creating a snapshot of a managed disk using the Azure portal, Azure PowerShell, and the Azure CLI.

IMPORTANT If you are following along, select resources and resource names based on your environment.

IMPORTANT If you are following along, be sure to delete any unwanted resources after you have completed testing to reduce charges levied by Microsoft.

USING AZURE PORTAL

To create a managed disk snapshot using the Azure portal, follow these steps:

1. In the Overview page for the managed disk you created earlier, click **Create Snapshot**. (See Figure 3-21.)

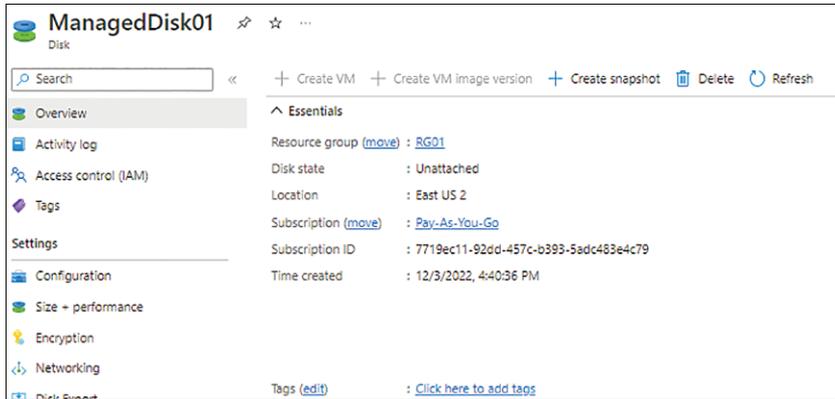


FIGURE 3-21 The Overview page for ManagedDisk01.

2. In the **Basics** tab of the Create Snapshot wizard (see Figure 3-22), enter the following information and click **Next**:
 - **Subscription** Select the subscription in which you want to create the snapshot.
 - **Resource Group** Select an existing resource group in which to create the snapshot or create a new one.
 - **Name** Enter a unique name for the snapshot.
 - **Snapshot Type** Leave this set to the default value of **Full**.

NOTE Figure 3-22 shows a Full button and an Incremental button. Your screen might not reflect that because this is the first time you're creating a snapshot of this managed disk. The next time you create a snapshot, you'll want to choose the Incremental button.

- **Storage Type** Select **Standard HDD**, **Standard SSD**, or **Premium SSD**, depending on your needs. (Remember, this is for the full snapshot; incremental snapshots always use Standard HDD disks.)
3. In the **Encryption** tab of the Create Snapshot wizard (see Figure 3-23), open the Key Management drop-down list and choose **Platform-Managed Key**, **Customer-Managed Key**, or **Platform-Managed and Customer-Managed Keys**. Then click **Next**.

NOTE To use customer-managed keys, you must first generate and store the keys in the Azure Key Vault service.

Create snapshot ...

Basics Encryption Networking Advanced Tags Review + create

A snapshot is a read-only copy of a virtual hard drive (VHD). You can take a snapshot of an OS or data disk VHD to use as a backup, or to troubleshoot virtual machine (VM) issues. [Learn more about snapshots in Azure](#)

Project details

Select the subscription to manage deployed resources and costs. Use resource groups like folders to organize and manage all your resources.

Subscription *

Resource group * [Create new](#)

Instance details

Name *

Region

Snapshot type * Full - make a complete read-only copy of the selected disk.
 Incremental - save on storage costs by making a partial copy of the disk based on the difference between the last snapshot.

Source type

Source subscription

Source disk

Security type

VM generation Generation 1
 Generation 2

VM architecture x64
 Arm64

Storage type *

FIGURE 3-22 The Basics tab of the Create Snapshot wizard.

Basics **Encryption** Networking Advanced Tags Review + create

Configure encryption options for your snapshot. [Learn more](#)

Key management

FIGURE 3-23 The Encryption tab of the Create Snapshot wizard.

- In the **Networking** tab of the Create Snapshot wizard (see Figure 3-24), in the **Network Access** section, select the **Enable Public Access from All Networks** option button.

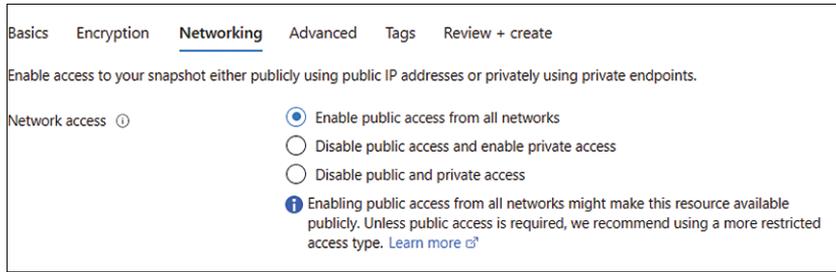


FIGURE 3-24 The Networking tab of the Create Snapshot wizard.

5. The **Advanced** tab of the Create Snapshot wizard (see Figure 3-25) includes an **Enable Data Access Authentication Mode** check box. For this example, leave it unchecked. Then click **Next**.

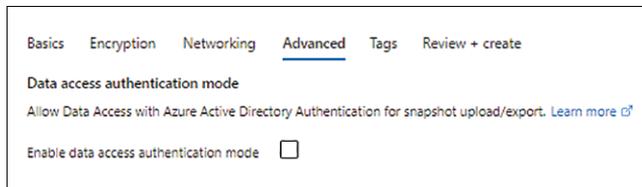


FIGURE 3-25 The Advanced tab of the Create Snapshot wizard.

6. In the **Tags** tab (see Figure 3-26), enter any tags you want to associate with the snapshot and click **Next**.

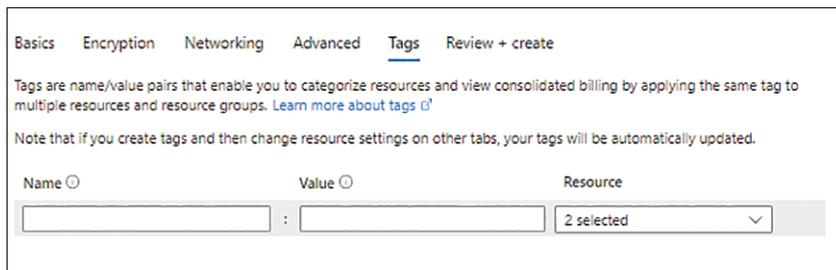


FIGURE 3-26 The Tags tab of the Create Snapshot wizard.

7. In the **Review + Create** tab (see Figure 3-27), review your settings, and click **Create** to create the snapshot.

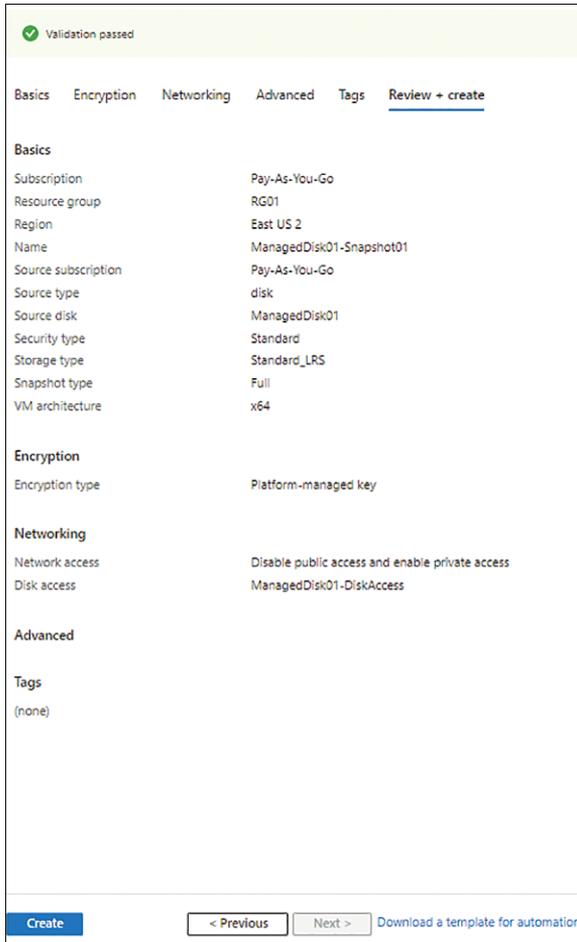


FIGURE 3-27 The Review + Create tab of the Create Snapshot wizard.

8. After the snapshot is created, click **Go to Resource** to access its page. (See Figure 3-28.)

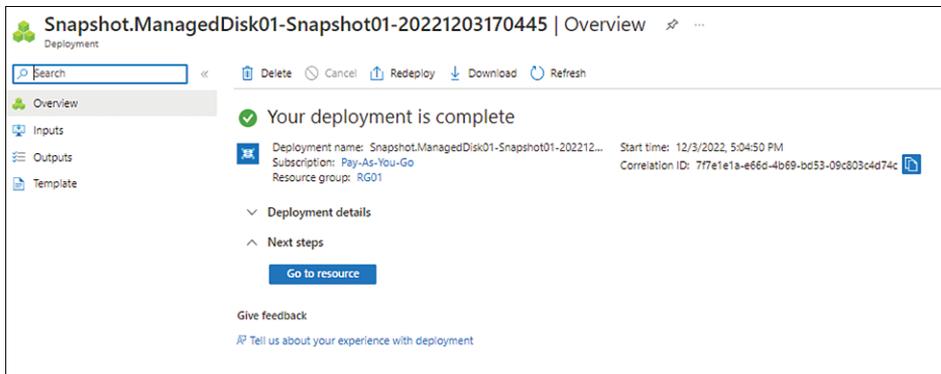


FIGURE 3-28 Snapshot deployment completion.

The snapshot's Overview page displays the properties of the snapshot, as well as Create Disk, Copy Snapshot, Delete, and Refresh options. (See Figure 3-29.)

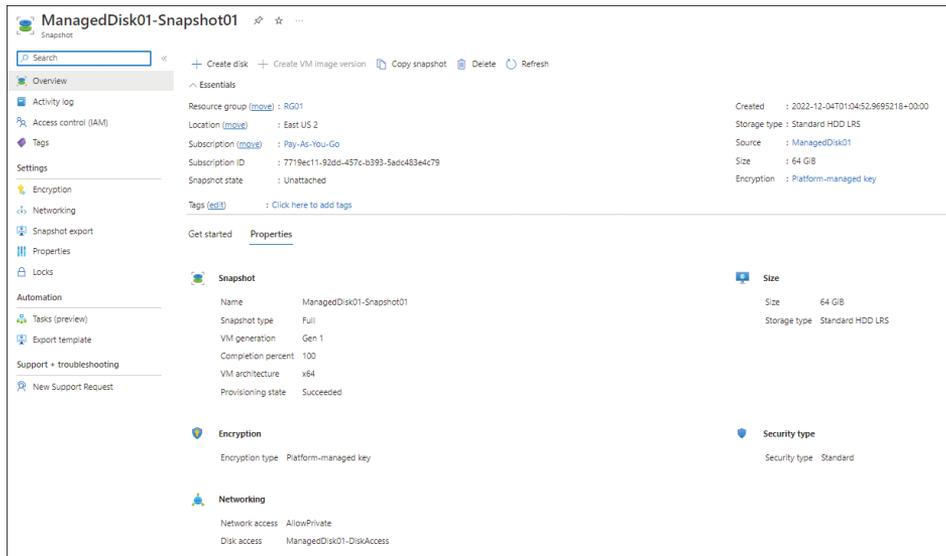


FIGURE 3-29 The new disk snapshot's Overview page.

USING AZURE POWERSHELL

Use the following Azure PowerShell code to create a disk snapshot:

```
#Define variables
$resourceGroup = "RG01"
$location = "EastUS2"
$vm = "SourceVM"
$snapshotName = "SourceVM-Snapshot-20230228"

#get the VM
$vmInfo = Get-AzVM `
    -ResourceGroupName $resourceGroup `
    -Name $vm

#Create the snapshot configuration
$snapshotConfig = New-AzSnapshotConfig `
    -SourceUri $vmInfo.StorageProfile.OsDisk.ManagedDisk.Id `
    -Location $location `
    -CreateOption copy

#Take the snapshot.
New-AzSnapshot `
    -Snapshot $snapshotConfig `
```

```

    -SnapshotName $snapshotName `
    -ResourceGroupName $resourceGroup

#Verify snapshot
Get-AzSnapshot `
    -ResourceGroupName $resourceGroup

```

USING AZURE CLI

Use the following code to create a disk snapshot in the Azure CLI:

```

#Define variables
resourceGroup="RG01"
location="EastUS2"
vm="SourceVM"
snapshotName="SourceVM-Snapshot-20230228"

#get the VM
DiskInfo=$(az vm show \
    --resource-group $resourceGroup \
    --name $vm \
    --query "storageProfile.osDisk.managedDisk.id" \
    -o tsv)

#Take the snapshot.
az snapshot create \
    --resource-group $resourceGroup \
    --source "$DiskInfo" \
    --name $snapshotName

#Verify snapshot
az snapshot list \
    --resource-group $resourceGroup \
    -o table

```

Managed images

Managed images enable you to create hundreds of copies of customized VMs in Azure without having to create multiple copies of the underlying disks associated with each VM or manage any storage accounts to host them. You can easily create managed images out of managed disks; the resulting managed image will contain the configuration of the source VM, including all the managed disks associated with that source VM. This helps you to scale your VM resources using features like VMSS or Azure Virtual Desktop Session Host Pools, where capacity is added as load increases.

The primary difference between managed disks and managed images is that an image is built from a generalized VM and includes all the associated disks, whereas a snapshot is specific

to a single disk and is a point-in-time copy of that disk. Generalizing a VM removes machine and user-specific information from the VM. So, for a VM that has multiple disks using disk spanning, a snapshot currently does not support a coordinated restore of all the disks and, therefore, might not be the right solution.

Managed images walkthrough

The following sections step you through the process of creating a managed image using the Azure portal, Azure PowerShell, and the Azure CLI.

IMPORTANT If you are following along, select resources and resource names based on your environment.

IMPORTANT If you are following along, be sure to delete any unwanted resources after you have completed testing to reduce charges levied by Microsoft.

PREREQUISITE If you are following along, you must create a VM to use to create the managed image. Be sure to stop that VM before starting the following procedure, however. The wizard will generalize this VM and make it unusable after the image is captured. (Optionally, you back up the VM and restore it after the process is complete.)

USING AZURE PORTAL

To create a managed image using the Azure portal, follow these steps:

1. On the Overview page of the VM for which you want to create an image, click **Capture**. (See Figure 3-30.)

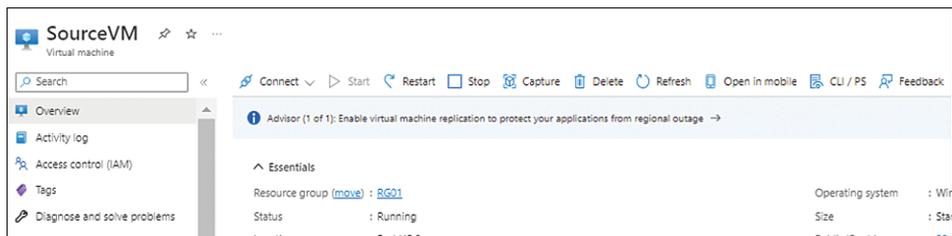


FIGURE 3-30 The Overview page for the VM.

2. In the **Basics** tab of the Create an Image wizard (see Figure 3-31), enter the following information and click **Next**:
 - **Resource Group** Select an existing resource group in which to create the new managed image or create a new one.
 - **Share Image to Azure Compute Gallery** For this walkthrough, select the **No, Capture Only a Managed Image** option button.

- **Automatically Delete this Virtual Machine After Creating the Image** Leave this checkbox unchecked (the default).
- **Zone Resiliency** Select this check box if you want to create a zone redundant image.
- **Name** Enter a unique name for the managed image.

Create an image ...

Basics Tags Review + create

Create an image from this virtual machine that can be used to deploy additional virtual machines and virtual machine scale sets. To create a managed image, you must first generalize this virtual machine. [Learn more](#) [ⓘ]

Project details

Subscription

Resource group * [Create new](#)

Instance details

Region

Share image to Azure compute gallery [ⓘ] Yes, share it to a gallery as a VM image version.
 No, capture only a managed image.

Automatically delete this virtual machine after creating the image [ⓘ]

Zone resiliency [ⓘ]

[ⓘ] Before creating an image, use "generalize" to prepare the Windows guest OS on the virtual machine. If you create an image from a virtual machine that hasn't been generalized, any virtual machines created from that image won't start. [Learn more](#) [ⓘ]

Name * [ⓘ] ✓

FIGURE 3-31 The Basics tab of the Create an Image wizard.

3. In the **Tags** tab (see Figure 3-32), enter any tags you want to associate with the managed image and click **Next**.

Basics Tags Review + create

Tags are name/value pairs that enable you to categorize resources and view consolidated billing by applying the same tag to multiple resources and resource groups. [Learn more about tags](#) [ⓘ]

Note that if you create tags and then change resource settings on other tabs, your tags will be automatically updated.

Name [ⓘ]	Value [ⓘ]
<input type="text" value="ms-resource-usage"/>	<input type="text" value="SourceVM-ManagedImage"/>
<input type="text"/>	<input type="text"/>

FIGURE 3-32 The Tags tab of the Create an Image wizard.

4. In the **Review + Create** tab (see Figure 3-33), review your settings, and click **Create** to create the managed image.



FIGURE 3-33 The Review + Create tab of the Create an Image wizard.

The source VM will be stopped automatically if you haven't turned it off already. (See Figure 3-34.) Azure will then generalize the VM and create the image.

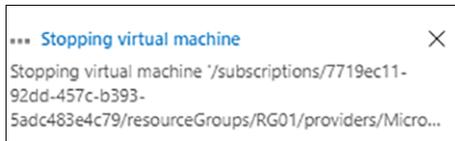


FIGURE 3-34 The VM is stopped (unless you stopped it already).

5. After the managed image is created, click **Go to Resource** to access its page. (See Figure 3-35.)

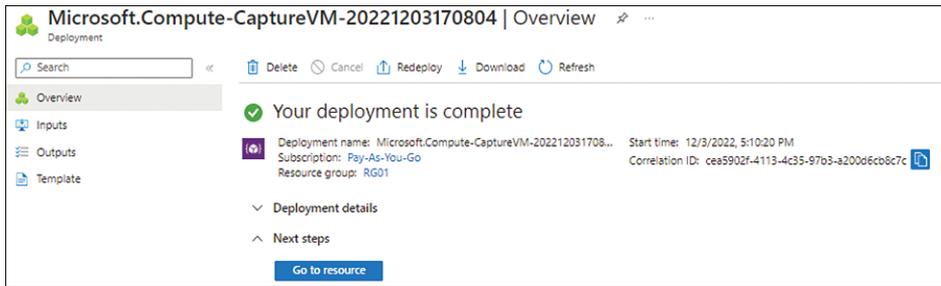


FIGURE 3-35 Managed image deployment completion.

The managed image's Overview page displays the properties of the managed image as well as Create VM, Clone to a VM Image, Delete, and Refresh options. (See Figure 3-36.)

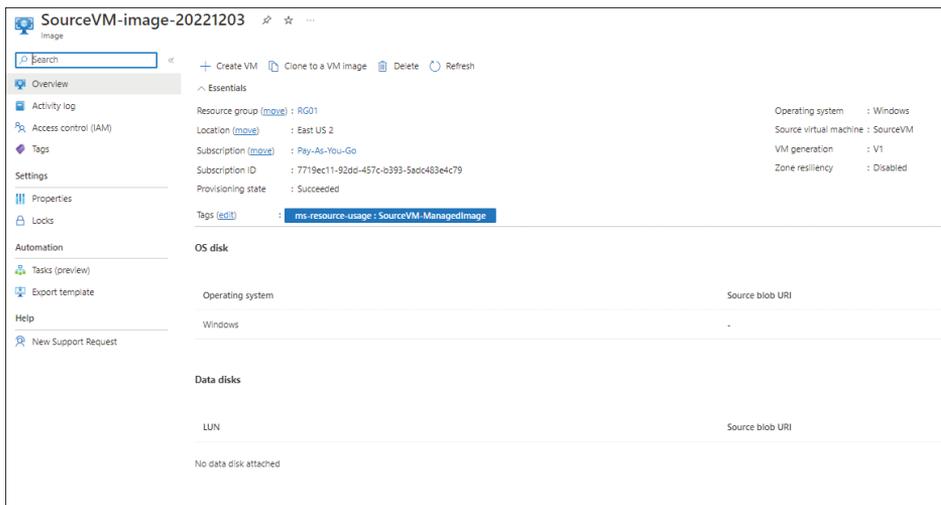


FIGURE 3-36 The new managed image's Overview page.

USING AZURE POWERSHELL

Use the following Azure PowerShell code to create a managed image:

```
#Define variables
$vm = "SourceVM"
$resourcegroup = RG01
$location = "EastUS2"
$imageName = "SourceVM-Image-20221203"
```

```

#VM has been deallocated
Stop-AzVM -ResourceGroupName $resourcegroup -Name $vm -Force

#Set the status of the virtual machine to Generalized.
Set-AzVm -ResourceGroupName $resourcegroup -Name $vm -Generalized

#Create the image configuration.
$vmInfo = Get-AzVM -Name $vm -ResourceGroupName $resourcegroup
$vmImage = New-AzImageConfig -Location $location -SourceVirtualMachineId $vmInfo.Id

#Create the image.
New-AzImage -Image $vmImage -ImageName $imageName -ResourceGroupName $resourcegroup

```

USING AZURE CLI

Use the following code to create a managed image in the Azure CLI:

```

#Define variables
vm="SourceVM"
resourcegroup=$RG01
location="EastUS2"
imageName="SourceVM-Image-20221203"

#VM has been deallocated
az vm deallocate \
  --resource-group $resourcegroup \
  --name $vm

#Set the status of the virtual machine to Generalized.
az vm generalize \
  --resource-group $resourcegroup \
  --name $vm

#Create the image.
az image create \
  --resource-group $resourcegroup \
  --location $location \
  --zone-resilient false \
  --name $imageName --source $vm

```

Performance tiering

When you create a managed disk, Azure automatically assigns a default performance target for that disk. This is based on predefined targets associated with the disk provisioned for the managed disk. This determines the IOPS and throughput available for that managed disk. This

Index

A

ABAC (Attribute-Based Access Control), Azure AD, 25

access

 Azure Files, 118–122

 Azure Queue Storage, 220–221

 Azure AD, 221

 SAS, 222–223

 storage account access keys, 221

 data access authorization, 24–25

 Azure AD, 25–26

 SAS, 29–33

 shared key authorization, 26–28

 storage access tiers, Azure Blob Storage, 49–50

accounts

 Azure Queue Storage, 195–202

 storage accounts, 1

 Azure Files, 82–90

 costs, 3

 failover, 74

 firewalls, 112–116

 names, 2

 premium block blobs, 3

 premium page blobs, 3

 standard general purpose v2 storage accounts, 3

 types of, 3

ACL (Access Control Lists), Azure Data Box, 237

ADE (Azure Disk Encryption), 158

append blobs, 4

archive tier, storage access, 50

archived blobs, rehydrating, 50

at-rest data encryption, Azure Blob Storage, 10–11

authorization, data access, 24–25

 Azure AD, 25–26

 SAS, 29–33

 shared key authorization, 26–28

availability, Azure Data Box, 236

Azure AD (Active Directory), 25

 ABAC, 25

 Azure Queue Storage, 221

 RBAC, 26

 resource scope, 26

Azure Backup integration, 65–70

Azure Blob Storage

 archived blobs, rehydrating, 50

 Azure Backup integration, 65–70

 best practices, 75–77

 blobs, 2

 append blobs, 4

 block blobs, 4

 change feeds, 60

 lifecycle management, 54–58

 point-in-time restores, 61

 premium block blobs, 3

 premium page blobs, 3

 rehydrating, 50

 snapshots, 70–73

 soft-deleting, 59

 versioning, 59

 components, overview, 1–2

 containers, 1, 4

 names, 4

 soft-deleting, 59

 CRC, 12

 data access authorization, 24–25

 Azure AD, 25–26

 SAS, 29–33

 shared key authorization, 26–28

 data protection, 59–65

 data redundancy, 5

 GRS, 8

 GZRS, 9

 LRS, 6

 RA-GRS, 10–11

 RA-GZRS, 10–11

 ZRS, 6–7

 disaster recovery, 73–75

 encryption, 10–11

Azure Blob Storage

Azure Blob Storage, *continued*

- networking
 - endpoints, 41–49
 - NFS 3.0 protocol, 35
 - routing, 33–35
 - SSH File Transfer protocol, 36
 - storage account firewalls, 36
 - overview, 1
 - page blobs, 4
 - premium block blobs, 3
 - static website hosting, 58
 - storage access tiers, 49–50
 - configuring, 51–54
 - early deletion fees, 51
 - rehydrating archived blobs, 50
 - storage accounts, 1
 - Azure CLI setups, 23–24
 - Azure portal setups, 12–21
 - Azure PowerShell setups, 22–23
 - costs, 3
 - names, 2
 - premium block blobs, 3
 - premium page blobs, 3
 - setting up, 12–24
 - standard general purpose v2 storage accounts, 3
 - types of, 3
 - walkthrough, 12–24
 - storage data integrity, 12
 - storage endpoints, 10
 - storage reservations, 58
 - tiers, 5
- Azure CLI (Command-Line Interface)
- Azure Backup integration, 69–70
 - Azure Blob Storage
 - private endpoints, 48–49
 - storage access tiers, 54
 - storage account firewalls, 40–41
 - storage accounts, 23–24
 - Azure Files
 - backups, 135–136
 - file shares, 98–99
 - identity/access, 121–122
 - networking, endpoints, 109–110
 - SMB MultiChannel, 118
 - storage account firewalls, 115–116
 - storage accounts, 89–90
 - Azure Managed Disks
 - creating, 150–151
 - managed images, 172
 - performance tiering, 176
 - Private Link integration, 156–158
 - shared disks, 180
 - snapshots, 167
- Azure Queue Storage
- accounts, 202
 - firewalls, 213
 - Last Sync Time, 227
 - private endpoints, 219–220
 - queues and messages, 208
- blob snapshots, 73
- data protection, 64–65
 - Last Sync Time, 75
 - network routing, 35
 - SAS, 32–33
- Azure Data Box, 232
- ACL, 237
 - availability, 236
 - Azure Data Box Disk, 232–233
 - Azure Data Box Heavy, 233
 - best practices, 249
 - components, 232–233
 - data resiliency, 236–237
 - data security, 234–235
 - data-transfer speeds, 235
 - features, 231–232
 - file attributes, 237
 - import/export workflow, 233–234
 - limitations, 238
 - overview, 231
 - supported client OS, 236
 - supported services, 236
 - timestamps, 237
 - use cases, 248–249
 - walkthrough, 238–248
- Azure Data Explorer, 255–256
- Azure Data Share
- Azure Data Explorer, 255–256
 - best practices, 271
 - data consumers, 252, 256–257
 - data providers, 252, 256
 - data stores, 254
 - data types, 252
 - features, 251–252
 - key concepts, 252–254
 - managed identities, 256–257
 - overview, 251
 - in-place sharing, 253
 - receiving data, 265–271

- sharing data, 257–265
- sharing models, 253
- snapshot-based sharing, 253, 255
- SQL-based sources, 255
- storage services, 254
- Azure Files
 - access, 118–122
 - backups, 127–136
 - best practices, 137–138
 - data protection, 123–136
 - data redundancy, 122
 - GRS, 123
 - GZRS, 123
 - LRS, 122
 - ZRS, 122–123
 - deployment models, 81
 - encryption, 111–112, 124
 - features, 79–81
 - file shares, 90–99
 - cool file shares, 100–101
 - file shares, storage tiers, 99–101
 - hot file shares, 100
 - premium file shares, 100
 - transaction-optimized file shares, 100
 - identity/access, 118–122
 - networking, 101
 - encryption in transit, 111–112
 - endpoints, 103–110
 - FileREST API, 102
 - NFS, 102
 - routing, 110–111
 - SMB, 101–102
 - storage account firewalls, 112–116
 - overview, 79
 - share snapshots, 122
 - SMB MultiChannel, 116–118
 - soft-deleting, 124
 - storage accounts, 82–90
- Azure Managed Disks
 - backups, 181–192
 - best practices, 192
 - bursting, 180
 - creating, 143–151
 - data disks, 141
 - disk redundancy, 176–177
 - disk roles, 140–141
 - disk types, 141–143
 - encryption, 158
 - features, 139–140
 - managed images, 167–172
 - OS disks, 141
 - overview, 139
 - performance tiering, 172–176
 - premium SSD disks, 142
 - Private Link integration, 151–158
 - shared disks, 177–180
 - snapshots, 158–167
 - standard HDD disks, 142
 - standard SSD disks, 142
 - temporary disks, 141
 - ultra disks, 142–143
- Azure portal
 - Azure Blob Storage
 - Azure Backup integration, 66–68
 - blob lifecycle management, 55–57
 - blob snapshots, 71–72
 - data protection, 62–63
 - network routing, 34
 - private endpoints, 42–48
 - SAS, 29–32
 - shared key authorization, 27–28
 - storage access tiers, 51–53
 - storage account firewalls, 37–39
 - storage accounts, 12–21
 - Azure Data Box, 238–248
 - Azure Data Share
 - receiving data, 265–271
 - sharing data, 257–265
 - Azure Files
 - backups, 129–133
 - file shares, 91–96
 - identity/access, 120–121
 - networking, endpoints, 104–108
 - share snapshots, 125–127
 - SMB MultiChannel, 116–117
 - storage account firewalls, 112–114
 - storage accounts, 83–88
 - Azure Managed Disks
 - backups, 181–192
 - creating, 144–149
 - managed images, 168–171
 - performance tiering, 174–175
 - Private Link integration, 151–156
 - shared disks, 179
 - snapshots, 161–166
 - Azure Queue Storage
 - accounts, 196–202
 - firewalls, 209–212

Azure portal

Azure portal, *continued*

- private endpoints, 214–218
 - queues and messages, 203–207
 - SAS, 222–223
- Azure PowerShell
- Azure Blob Storage
 - Azure Backup integration, 68–69
 - blob lifecycle management, 57–58
 - blob snapshots, 72–73
 - data protection, 63–64
 - Last Sync Time, 75
 - network routing, 35
 - private endpoints, 48
 - SAS, 32
 - storage access tiers, 53
 - storage account firewalls, 39–40
 - storage accounts, 22–23
 - Azure Files
 - backups, 134–135
 - file shares, 96–98
 - identity/access, 121
 - networking, endpoints, 108–109
 - share snapshots, 127
 - SMB MultiChannel, 117
 - storage account firewalls, 114–115
 - storage accounts, 89
 - Azure Managed Disks
 - creating, 149–150
 - managed images, 171–172
 - performance tiering, 175
 - shared disks, 179
 - snapshots, 166–167
 - Azure Queue Storage
 - firewalls, 212–213
 - Last Sync Time, 226
 - private endpoints, 218–219
- Azure Queue Storage
- accounts, 195–202
 - Azure AD, 221
 - best practices, 228–230
 - data redundancy, 223–225
 - disaster recovery, 225
 - encryption, 227
 - client-side encryption, 228
 - infrastructure encryption, 227
 - service-level encryption, 228
 - features, 194
 - identity/access, 220–221
 - Azure AD, 221
 - SAS, 222–223
 - storage account access keys, 221
 - Last Sync Time, 226–227
 - networking, 208
 - firewalls, 209–213
 - private endpoints, 213–220
 - secure transfers, 220
 - TLS, 220
 - overview, 193–194
 - queues and messages, 202–208
 - storage account failover, 225–226

B

- backups
- Azure Backup integration, 65–70
 - Azure Files, 127–136
 - Azure Managed Disks, 181–192
- best practices
- Azure Blob Storage, 75–77
 - Azure Data Box, 249
 - Azure Data Share, 271
 - Azure Files, 137–138
 - Azure Managed Disks, 192
 - Azure Queue Storage, 228–230
- Blob Storage. *See* Azure Blob Storage
- blobs, 2
- append blobs, 4
 - archived blobs, rehydrating, 50
 - block blobs, 4
 - change feeds, 60
 - lifecycle management, 54–58
 - page blobs, 4
 - point-in-time restores, 61
 - premium block blobs, 3
 - premium page blobs, 3
 - snapshots, 70–73
 - soft-deleting, 59
 - versioning, 59
- block blobs, 4
- bursting
- Azure Managed Disks, 180
 - credit-based bursting, 180
 - on-demand bursting, 180

C

- change feeds, Azure Blob Storage, 60
- CIFS. *See* SMB
- CLI (Command-Line Interface)
 - Azure Backup integration, 69–70
 - Azure Blob Storage
 - private endpoints, 48–49
 - storage access tiers, 54
 - storage account firewalls, 40–41
 - storage accounts, 23–24
 - Azure Files
 - backups, 135–136
 - file shares, 98–99
 - identity/access, 121–122
 - networking, endpoints, 109–110
 - SMB MultiChannel, 118
 - storage account firewalls, 115–116
 - storage accounts, 89–90
 - Azure Managed Disks
 - creating, 150–151
 - managed images, 172
 - performance tiering, 176
 - Private Link integration, 156–158
 - shared disks, 180
 - snapshots, 167
 - Azure Queue Storage
 - accounts, 202
 - firewalls, 213
 - Last Sync Time, 227
 - private endpoints, 219–220
 - queues and messages, 208
 - blob snapshots, 73
 - data protection, 64–65
 - Last Sync Time, 75
 - network routing, 35
 - SAS, 32–33
- client-OS, Azure Data Box, 236
- client-side encryption, Azure Queue Storage, 228
- configuring
 - storage access tiers, 51–54
 - storage accounts
 - Azure CLI, 23–24
 - Azure portal, 12–24
 - Azure PowerShell, 22–23
- containers, 1, 4
 - names, 4
 - soft-deleting, 59
- cool file shares, Azure Files, 100–101

- cool tier, storage access, 50
- costs, data storage in storage accounts, 3
- CRC (Cyclic Redundancy Checks), 12
- credit-based bursting, 180

D

- data access authorization, 24–25
 - Azure AD, 25–26
 - SAS, 29–33
 - shared key authorization, 26–28
- Data Box. *See* Azure Data Box
- data consumers, Azure Data Share, 252, 256–257
- data disks, 141
- data integrity, Azure Blob Storage, 12
- data protection
 - Azure Blob Storage, 59–65
 - Azure Files, 123–136
- data providers, Azure Data Share, 252, 256
- data redundancy, 5
 - Azure Blob Storage
 - GRS, 8
 - GZRS, 9
 - LRS, 6
 - RA-GRS, 10–11
 - RA-GZRS, 10–11
 - ZRS, 6–7
 - Azure Files, 122
 - GRS, 123
 - GZRS, 123
 - LRS, 122
 - ZRS, 122–123
 - Azure Queue Storage, 223–225
- data resiliency, Azure Data Box, 236–237
- data security, Azure Data Box, 234–235
- data stores, Azure Data Share, 254
- data transfers. *See* Azure Data Box
- data types, Azure Data Share, 252
- data-transfer speeds, Azure Data Box, 235
- deleting, blobs/containers, 59
- deletion fees, storage access tiers, 51
- deployment models, Azure Files, 81
- disaster recovery
 - Azure Blob Storage, 73–75
 - Azure Queue Storage, 225
 - Last Sync Time, 74–75
 - storage account failover, 74
- disk redundancy, Azure Managed Disks, 176–177

E

- early deletion fees, storage access tiers, 51
- encryption
 - ADE, 158
 - Azure Blob Storage, 10–11
 - Azure Files, 111–112, 124
 - Azure Managed Disks, 158
 - Azure Queue Storage, 227
 - client-side encryption, 228
 - infrastructure encryption, 227
 - service-level encryption, 228
 - SSE, 158
- endpoints
 - Azure Blob Storage, 42–49
 - private endpoints, 42–49
 - public endpoints, 42–49
 - Azure Files
 - private endpoints, 103–110
 - public endpoints, 103
 - Azure Queue Storage, 213–220
- export/import workflow, Azure Data Box, 233–234

F

- failover, storage accounts, 74
- file attributes, Azure Data Box, 237
- FileREST API, Azure Files, 102
- file shares, Azure Files, 90–99
 - cool file shares, 100–101
 - file shares, storage tiers, 99–101
 - hot file shares, 100
 - premium file shares, 100
 - transaction-optimized file shares, 100
- firewalls
 - Azure Queue Storage, 209–213
 - storage accounts
 - Azure Blob Storage, 36
 - Azure Files, 112–116

G

- GRS (Geo-Redundant Storage)
 - Azure Blob Storage, 8
 - Azure Files, 123
 - Azure Queue Storage, 224

- GZRS (Geo-Zone Redundant Storage)
 - Azure Blob Storage, 9
 - Azure Files, 123
 - Azure Queue Storage, 224

H

- HDD disks, standard, 142
- hosting websites, Azure Blob Storage, 58
- hot file shares, Azure Files, 100
- hot tier, storage access, 49–50

I

- identity/access
 - Azure Data Share, 256–257
 - Azure Files, 118–122
 - Azure Queue Storage, 220–221
 - Azure AD, 221
 - SAS, 222–223
 - storage account access keys, 221
- images, managed, Azure Managed Disks, 167–172
- import/export workflow, Azure Data Box, 233–234
- infrastructure encryption, Azure Queue Storage, 227
- in-place sharing, Azure Data Share, 253
- Internet routing, Azure Files, 110–111

L

- Last Sync Time
 - Azure Blob Storage, 74–75
 - Azure Queue Storage, 226–227
- lifecycle management, blobs, 54–58
- LRS (Locally Redundant Storage)
 - Azure Blob Storage, 6
 - Azure Files, 122
 - Azure Managed Disks, 176
 - Azure Queue Storage, 223

M

- managing
 - blobs, lifecycle management, 54–58
 - discs. *See* Azure Managed Disks
 - images, Azure Managed Disks, 167–172

messages, Azure Queue Storage, 202–208
 Microsoft routing, Azure Files, 110–111

N

names

- containers, 4
- storage accounts, 2

networking

- Azure Blob Storage
 - endpoints, 41–49
 - NFS 3.0 protocol, 35
 - private endpoints, 42–49
 - public endpoints, 42
 - routing, 33–35
 - SSH File Transfer protocol, 36
 - storage account firewalls, 36
- Azure Files, 100–101
 - encryption in transit, 111–112
 - endpoints, 103–110
 - FileREST API, 102
 - NFS, 102
 - routing, 110–111
 - SMB, 101–102
 - storage account firewalls, 112–116
- Azure Queue Storage, 208
 - firewalls, 209–213
 - private endpoints, 213–220
 - secure transfers, 220
 - TLS, 220
- NFS 3.0 protocol, 35
- routing
 - Azure CLI, 35
 - Azure portal, 34
 - Azure PowerShell, 35
 - walkthrough, 34–35
- SSH File Transfer protocol, 36
- storage account firewalls, 36

- NFS (Network File System)
 - Azure Blob Storage, 35
 - Azure Files, 102

O

- on-demand bursting, 180
- OS (Operating Systems)
 - Azure Data Box, 236
 - disks, 141

P

- page blobs, 4
- partner integrations, Azure Data Box, 237
- performance tiering, Azure Managed Disks, 172–176
- point-in-time restores, 61
- PowerShell
 - Azure Blob Storage
 - Azure Backup integration, 68–69
 - blob lifecycle management, 57–58
 - blob snapshots, 72–73
 - data protection, 63–64
 - Last Sync Time, 75
 - network routing, 35
 - private endpoints, 48
 - SAS, 32
 - storage access tiers, 53
 - storage account firewalls, 39–40
 - storage accounts, 22–23
 - Azure Files
 - backups, 134–135
 - file shares, 96–98
 - identity/access, 121
 - networking, endpoints, 108–109
 - share snapshots, 127
 - SMB MultiChannel, 117
 - storage account firewalls, 114–115
 - storage accounts, 89
 - Azure Managed Disks
 - creating, 149–150
 - managed images, 171–172
 - performance tiering, 175
 - shared disks, 179
 - snapshots, 166–167
 - Azure Queue Storage
 - firewalls, 212–213
 - Last Sync Time, 226
 - private endpoints, 218–219
- premium block blobs, 3
- premium file shares, Azure Files, 100
- premium page blobs, 3
- premium SSD disks, 142
- private endpoints
 - Azure Blob Storage, 42–49
 - Azure Files, 103–110
 - Azure Queue Storage, 213–220
- Private Link integration, Azure Managed Disks, 151–158
- public endpoints
 - Azure Blob Storage, 42
 - Azure Files, 103

Q

queues, Azure Queue Storage, 202–208

R

RA-GRS (Read Only Geo-Redundant Storage)

Azure Blob Storage, 10–11

Azure Queue Storage, 224

RA-GZRS (Read Only Geo-Zone Redundant Storage)

Azure Blob Storage, 10–11

Azure Queue Storage, 224

RBAC (Role-Based Access Control), Azure AD, 26

receiving data, Azure portal, 265–271

redundancy

Azure Blob Storage, 5

GRS, 8

GZRS, 9

LRS, 6

RA-GRS, 10–11

RA-GZRS, 10–11

ZRS, 6–7

Azure Files, 122

GRS, 123

GZRS, 123

LRS, 122

ZRS, 122–123

Azure Managed Disks, 176–177

Azure Queue Storage, 223–225

rehydrating archived blobs, 50

reservations, storage, 58

resiliency, Azure Data Box, 236–237

resource scope, Azure AD, 26

restores, blobs, 61

routing, network

Azure Blob Storage, 33–35

Azure Files, 110–111

S

SAS (Shared Access Signature), 29

Azure CLI, 32–33

Azure portal, 29–32

Azure PowerShell, 32

Azure Queue Storage, 222–223

walkthrough, 29–33

secure transfers, Azure Queue Storage, 220

security

data protection, Azure Blob Storage, 59–65

data security, Azure Data Box, 234–235

encryption

ADE, 158

Azure Blob Storage, 10–11

Azure Files, 111–112, 124

Azure Managed Disks, 158

SSE, 158

firewalls, Azure Queue Storage, 209–213

storage account firewalls

Azure Blob Storage, 36

Azure Files, 112–116

TLS, Azure Queue Storage, 220

service-level encryption, Azure Queue Storage, 228

setting up

storage access tiers, 51–54

storage accounts

Azure CLI, 23–24

Azure portal, 12–24

Azure PowerShell, 22–23

share snapshots, Azure Files, 124–127

shared disks, Azure Managed Disks, 177–180

shared key authorization, 26–27

sharing data, Azure Data Share, 257–265

sharing models, Azure Data Share, 253

SMB (Server Message Blocks), Azure Files, 101–102

SMB MultiChannel, Azure Files, 116–118

snapshots

Azure Managed Disks, 158–167

blobs, 70–73

sharing, Azure Data Share, 253, 255

soft-deleting

Azure Files, 124

blobs/containers, 59

SQL-based sources, Azure Data Share, 255

SSD disks

premium SSD disks, 142

standard SSD disks, 142

SSE (Server-Side Encryption), 158

SSH File Transfer protocol, Azure Blob Storage, 36

standard general purpose v2 storage accounts, 3

standard HDD disks, 142

standard SSD disks, 142

static website hosting, Azure Blob Storage, 58

storage access tiers

Azure Blob Storage, 49–50

configuring, 51–54

early deletion fees, 51

rehydrating archived blobs, 50

- storage accounts, 1
 - access keys, Azure Queue Storage, 221
 - Azure Files, 82–90
 - Azure Queue Storage, access keys, 221
 - costs, 3
 - failover
 - Azure Blob Storage, 74
 - Azure Queue Storage, 225–226
 - firewalls, 36–41, 112–116
 - names, 1
 - premium page blobs, 3
 - setting up
 - Azure CLI, 23–24
 - Azure portal, 12–21
 - Azure PowerShell, 22–23
 - standard general purpose v2 storage accounts, 3
 - types of, 3
 - walkthrough, 12–24
- storage data integrity, Azure Blob Storage, 12
- storage endpoints, 10
- storage firewalls, Azure Queue Storage, 209–213
- storage reservations, 58
- storage services, Azure Data Share, 254

T

- temporary disks, 141
- tiers, Azure Blob Storage, 5

- timestamps, Azure Data Box, 237
- TLS (Transport Layer Security), Azure Queue Storage, 220
- transaction-optimized file shares, Azure Files, 100
- transferring data. *See* Azure Data Box

U

- ultra disks, 142–143

V

- versioning, blobs, 59

W

- website hosting, Azure Blob Storage, 58

X - Y - Z

- ZRS (Zone-Redundant Storage)

- Azure Blob Storage, 6–7
- Azure Files, 122–123
- Azure Managed Disks, 176
- Azure Queue Storage, 224