

Configure Azure Blob Storage

1. Introduction

<https://learn.microsoft.com/en-us/training/modules/configure-blob-storage/1-introduction>

Introduction

Completed

Azure Blob Storage is an AI-ready service for storing large amounts of unstructured object data. Unstructured data is data that doesn't adhere to a particular data model or definition, such as text or binary data.

In this module, your media company has an extensive library of video clips that are accessed thousands of times a day. The company relies on you to configure Blob Storage for the video data. You plan to use access tiers to reduce cost and improve performance. You're developing a lifecycle management strategy for the older videos. Your plan also includes configuring object replication for failover.

Learning objectives

In this module, you will:

- Understand the purpose and benefits of Azure Blob Storage.
- Create and configure Azure Blob Storage accounts.
- Manage containers and blobs within Azure Blob Storage.
- Optimize blob storage performance and scalability.
- Implement lifecycle management policies to automate data movement and deletion.
- Determine the best pricing plans for your Azure Blob Storage.

Skills measured

The content in the module helps you prepare for [Exam AZ-104: Microsoft Azure Administrator](#).

Prerequisites

Here are some common prerequisites that can be beneficial for understanding and successfully completing this module.

- Basic understanding of cloud computing: Familiarity with cloud computing concepts, such as virtualization, scalability, and pay-as-you-go pricing models, can provide a foundation for understanding how Azure Blob Storage fits into the broader cloud ecosystem.
- Knowledge of Azure fundamentals: Having a basic understanding of Microsoft Azure services and concepts, such as Azure Resource Manager, Azure Storage Accounts, and Azure Virtual Networks, can help you navigate and configure blob storage effectively.
- Understand fundamental storage concepts like file systems, directories, files, and data replication can be beneficial when working with blob storage.
- Experience with Azure portal or Azure CLI: Familiarity with the Azure portal (web-based management interface) or Azure CLI (command-line interface) can help you navigate and configure blob storage resources efficiently.
- Basic programming or scripting skills: While not always required, having some knowledge of programming or scripting languages like PowerShell or Python can be advantageous when automating blob storage configuration tasks.

2. Implement Azure Blob Storage

<https://learn.microsoft.com/en-us/training/modules/configure-blob-storage/2-implement>

Implement Azure Blob Storage


Completed

[Azure Blob Storage](#) is a service that stores unstructured data in the cloud as objects or blobs. Blob stands for Binary Large Object. Blob Storage is also referred to as *object storage* or *container storage*.



Things to know about Azure Blob Storage

Let's examine some configuration characteristics of Blob Storage.



- Blob Storage can store any type of text or binary data. Some examples are text documents, images, video files, and application installers.
- Blob Storage uses three resources to store and manage your data:
 - An Azure storage account
 - Containers in an Azure storage account
 - Blobs in a container
- To implement Blob Storage, you configure several settings:
 - Blob container options.
 - Blob types and upload options.
 - Blob Storage access tiers.
 - Blob lifecycle rules.
 - Blob object replication options.

Things to consider when implementing Azure Blob Storage

There are many common uses for Blob Storage. Consider the following scenarios and think about your own data needs:

- **Consider browser uploads.** Use Blob Storage to serve images or documents directly to a browser.
- **Consider distributed access.** Blob Storage can store files for distributed access, such as during an installation process.
- **Consider streaming data.** Stream video and audio by using Blob Storage.

- **Consider archiving and recovery.** Blob Storage is a great solution for storing data for backup and restore, disaster recovery, and archiving.
- **Consider application access.** You can store data in Blob Storage for analysis by an on-premises or Azure-hosted service.

3. Create blob containers

<https://learn.microsoft.com/en-us/training/modules/configure-blob-storage/3-create-blob-containers>

Create blob containers

Completed

Azure Blob Storage uses a container resource to group a set of blobs. A blob can't exist by itself in Blob Storage. A blob must be stored in a container resource.

Things to know about containers and blobs

Let's look at the configuration characteristics of containers and blobs.

- All blobs must be in a container.
- Containers organize your blob storage.
- A container can store an unlimited number of blobs.
- An Azure storage account can contain an unlimited number of containers.
- You must create a storage container before you can begin to upload data.

Configure a container

In the Azure portal, you configure settings to create a container for an Azure storage account. As you review these details, consider how you might organize containers in your storage account.

The image shows the 'New container' dialog box in the Azure portal. It has a title bar with a plus icon and the text 'Container' and 'Change access level'. Below the title bar, the text 'New container' is displayed. There are two main input fields: 'Name' with a red asterisk and 'Public access level' with an information icon. The 'Name' field contains the text 'container01' and has a green checkmark on the right. The 'Public access level' field is a dropdown menu currently showing 'Private (no anonymous access)'. A blue arrow points from this dropdown to a separate window showing the full list of options: 'Private (no anonymous access)', 'Blob (anonymous read access for blobs only)', and 'Container (anonymous read access for containers and blobs)'. At the bottom of the dialog are 'OK' and 'Cancel' buttons.

- **Name:** Enter a name for your container. The name must be unique within the Azure storage account.
 - The name can contain only lowercase letters, numbers, and hyphens.
 - The name must begin with a letter or a number.
 - The minimum length for the name is three characters.
 - The maximum length for the name is 63 characters.
- **Public access level:** The access level specifies whether the container and its blobs can be accessed publicly. By default, container data is private and visible only to the account owner. There are three access level choices:
 - **Private:** (Default) Prohibit anonymous access to the container and blobs.
 - **Blob:** Allow anonymous public read access for the blobs only.
 - **Container:** Allow anonymous public read and list access to the entire container, including the blobs.

Important

The Blob and Container access levels have no effect unless the storage account's **Allow Blob anonymous access** setting is enabled. When disabled, all containers remain private regardless of their individual access level settings. Microsoft recommends keeping anonymous access disabled at the account level unless serving public content scenarios.

4. Assign blob access tiers

Assign blob access tiers

Completed

Azure Storage supports several [access tiers](#) for blob data. These tiers include Hot, Cool, Cold, and Archive. Each access tier is optimized to support a particular pattern of data usage.



Things to know about blob access tiers

Let's examine characteristics of the blob access tiers.

Hot tier

The Hot tier is optimized for frequent reads and writes of objects in the Azure storage account. A good usage case is data that is actively being processed. The hot tier has the highest storage costs, but the lowest access costs.

Cool tier

The Cool tier is optimized for storing large amounts of infrequently accessed data. This tier is intended for data that remains in the Cool tier for at least 30 days. A usage case for the Cool tier is short-term backup and disaster recovery datasets and older media content. This content shouldn't be viewed frequently, but it needs to be immediately available. Storing data in the Cool tier is more cost-effective. The cool tier has lower storage costs and higher access costs compared to the hot tier.

Cold tier

The Cold tier is also optimized for storing large amounts of infrequently accessed data. This tier is intended for data that can remain in the tier for at least 90 days. The cold tier has lower storage costs and higher access costs compared to the cool tier.

Archive tier

The Archive tier is an offline tier that's optimized for data that can tolerate several hours of retrieval latency. Data must remain in the Archive tier for at least 180 days or be subject to an early deletion charge. Data for the Archive tier includes secondary backups, original raw data, and legally required compliance information. This tier is the most cost-effective option for storing data. Accessing data is more expensive in the Archive tier than accessing data in the other tiers.

To access the blob's content, you can rehydrate it to the hot, cool, or cold tier using two methods: **Copy Blob** (recommended - creates a new blob in an online tier) or **Set Blob Tier** (changes tier in place). Both methods support Standard priority (up to 15 hours) or High priority (within 1 hour for objects under 10 GB, at higher cost). Use High priority for urgent data retrieval in disaster recovery scenarios.

Compare access tiers

The access options for Azure Blob Storage offer a range of features and support levels to help you optimize your storage costs. As you compare the features and support, think about which access options can best support your application needs.

Comparison	Hot access tier	Cool access tier	Cold access tier	Archive access tier
Availability	99.9%	99%	99%	99%
Availability (RA-GRS reads)	99.99%	99.9%	99.9%	99.9%
Latency (time to first byte)	milliseconds	milliseconds	milliseconds	hours
Minimum storage duration	N/A	30 days	90 days	180 days

5. Add blob lifecycle management rules

<https://learn.microsoft.com/en-us/training/modules/configure-blob-storage/5-add-blob-lifecycle-management-rules>

Add blob lifecycle management rules

Completed

Every data set has a unique lifecycle. Early in the lifecycle, users tend to access some of the data in the set, but not all of the data. As the data set ages, access to all of the data in the set tends to dramatically reduce. Some data set stays idle in the cloud and is rarely accessed. Some data expires within a few days or months after creation. Other data is actively read and modified throughout the data set lifetime.

Azure Blob Storage supports [lifecycle management](#) for data sets. It offers a rich rule-based policy for GPv2 accounts and Premium block blob accounts. Legacy Blob Storage accounts are also supported, but GPv2 is recommended for new deployments. You can use lifecycle policy rules to transition your data to the appropriate access tiers, and set expiration times for the end of a data set's lifecycle.

Things to know about lifecycle management

You can use Azure Blob Storage lifecycle management policy rules to accomplish several tasks.

- Transition blobs to a cooler storage tier (Hot to Cool, Hot to Cold, Hot to Archive, Cool to Cold, Cool to Archive, Cold to Archive) to optimize for performance and cost.
- Delete current versions of a blob, previous versions of a blob, or blob snapshots at the end of their lifecycles.
- Automatically transition blobs from Cool back to Hot when accessed. This setting optimizes for unpredictable access patterns without early deletion charges.
- Apply rules to an entire storage account, to select containers, or to a subset of blobs using name prefixes or blob index tags as filters.

Business scenario

Consider a scenario where data is frequently accessed in the early stages of the lifecycle, but only occasionally after two weeks. After the first month, the data set is rarely accessed. In this scenario, the Hot tier of Blob Storage is best during the early stages. Cool tier storage is most appropriate for occasional access. Archive tier storage is the best option after the data ages over a month. To achieve this transition, lifecycle management policy rules are available to move aging data to cooler tiers.

Configure lifecycle management policy rules

In the Azure portal, you create lifecycle management policy rules for your Azure storage account by specifying several settings. For each rule, you create **If - Then** block conditions to transition or expire data based on your specifications. As you review these details, consider how you can set up lifecycle management policy rules for your data sets.

Add a rule ...

✓ Details 2 Base blobs 3 Filter set

Lifecycle management uses your rules to automatically move blobs to cooler tiers or to delete them. If you create multiple rules, the associated actions must be implemented in tier order (from hot to cool storage, then archive, then deletion).

If

Base blobs were *

☒ Last modified

☐ Created

More than (days ago) *

Then

Delete the blob

Move to cool storage
For infrequently accessed data that you want to keep on cool storage for at least 30 days.

Move to cold storage
For rarely accessed data that you want to keep for at least 90 days.

Move to archive storage
Use if you don't need online access and want to keep the object for 180 days or longer.

Delete the blob
Deletes the object per the specified conditions.

- **If:** The **If** clause sets the evaluation clause for the policy rule. When the **If** clause evaluates to true, the **Then** clause is executed. Use the **If** clause to set the time period to apply to the blob data. The lifecycle management feature checks if the data is accessed or modified according to the specified time.
 - **More than (days ago):** The number of days to use in the evaluation condition.
- **Then:** The **Then** clause sets the action clause for the policy rule. When the **If** clause evaluates to true, the **Then** clause is executed. Use the **Then** clause to set the transition action for the blob data. The lifecycle management feature transitions the data based on the setting.
 - **Move to cool storage:** The blob data is transitioned to Cool tier storage.
 - **Move to cold storage:** The blob data is transitioned to Cold tier storage.
 - **Move to archive storage:** The blob data is transitioned to Archive tier storage.
 - **Delete the blob:** The blob data is deleted.

By designing policy rules to adjust storage tiers in respect to the age of data, you can design the least expensive storage options for your needs.

Tip

Expand your knowledge in the [Manage the Azure Blob storage lifecycle](#) training module.

6. Determine blob object replication

<https://learn.microsoft.com/en-us/training/modules/configure-blob-storage/6-determine-blob-object-replication>


Determine blob object replication

Completed

Object replication copies blobs in a container asynchronously according to policy rules that you configure.



Replication includes the blob content, metadata properties, and versions. The following illustration shows an example of asynchronous replication of blob containers between regions.



Things to know about blob object replication

There are several considerations to keep in mind when planning your configuration for blob object replication.

- Object replication requires that [Blob versioning](#) is enabled on both the source and destination accounts. When blob versioning is enabled, you can access earlier versions of a blob. This access lets you recover your modified or deleted data.
- Object replication doesn't support blob snapshots. Any snapshots on a blob in the source account aren't replicated to the destination account.
- Object replication is supported when the source and destination accounts are in the Hot, Cool, or Cold tier. The source and destination accounts can be in different tiers.
- When you configure object replication, you create a replication policy that specifies the source Azure storage account and the destination storage account.
- A replication policy includes one or more rules that specify a source container and a destination container. The policy identifies the blobs in the source container to replicate.

Things to consider when configuring blob object replication

There are many benefits to using blob object replication. Consider the following scenarios and think about how replication can be a part of your Blob Storage strategy.

- **Consider latency reductions.** Minimize latency with blob object replication. You can reduce latency for read requests by enabling clients to consume data from a region that's in closer physical proximity.

- **Consider efficiency for compute workloads.** Improve efficiency for compute workloads by using blob object replication. With object replication, compute workloads can process the same sets of blobs in different regions.
- **Consider data distribution.** Optimize your configuration for data distribution. You can process or analyze data in a single location and then replicate only the results to other regions.
- **Consider costs benefits.** Manage your configuration and optimize your storage policies. After your data is replicated, you can reduce costs by moving the data to the Archive tier by using lifecycle management policies.

7. Manage blobs

<https://learn.microsoft.com/en-us/training/modules/configure-blob-storage/7-upload-blobs>

Manage blobs

Completed

A blob can be any type of data and any size file. Azure Storage offers three types of blobs: *block blob*, *page blob*, and *append blob*.

Things to know about blob types

Let's take a closer look at the characteristics of blob types.

- **Block blobs.** A block blob consists of blocks of data that are assembled to make a blob. Most Blob Storage scenarios use block blobs. Block blobs are ideal for storing text and binary data in the cloud, like files, images, and videos. The block blob type is the default type for a new blob. When you're creating a new blob, if you don't choose a specific type, the new blob is created as a block blob.
- **Append blobs.** An append blob is similar to a block blob because the append blob also consists of blocks of data. The blocks of data in an append blob are optimized for *append* operations. Append blobs are useful for logging scenarios, where the amount of data can increase as the logging operation continues.
- **Page blobs.** A page blob can be up to 8 TB in size. Page blobs are more efficient for frequent read/write operations. Azure Virtual Machines uses page blobs for operating system disks and data disks.

Note

After you create a blob, you can't change its type.

Things to consider when managing blob storage

You can use the portal to upload and manage blobs. This option is good for a few files. After you identify the files to upload, you choose the blob type and block size, and the container folder. You also set the access tier and the encryption scope.

Upload blob

☐ Overwrite if files already exist

Advanced

Blob type ⓘ

Block blob^

Block blob

Page blob

Append blob

☒ Upload .vhd files as page blobs (recommended)

Block size ⓘ

4 MiB

Access tier ⓘ

Hot (Inferred)

Upload to folder


Encryption scope

☒ Use existing default container scope

☐ Choose an existing scope

For larger numbers of files, it's best to use a tool. Review the following options and consider which tools would suit your configuration needs.

- **Azure Storage Explorer.** Upload, download, and manage blobs, files, queues, and tables, as well as Azure Data Lake Storage entities and managed disks. You can also view, edit, and manage resources, preview data, and configure storage permissions and access controls.



- **AzCopy**. An easy-to-use command-line tool for Windows and Linux. You can copy data to and from Blob Storage, across containers, and across storage accounts.
- **Azure Data Box Disk**. A service for transferring on-premises data to Blob Storage when large datasets or network constraints make uploading data over the wire unrealistic. You can use Azure Data Box Disk to request solid-state disks (SSDs) from Microsoft. You can copy your data to those disks and ship them back to Microsoft to be uploaded into Blob Storage.

8. Determine Blob Storage pricing

<https://learn.microsoft.com/en-us/training/modules/configure-blob-storage/8-determine-storage-pricing>


Determine Blob Storage pricing

Completed

Understanding your access patterns and correlating them with your durability and availability needs helps you to best manage your Azure Blob Storage costs. The primary tool for estimating these costs is the Azure pricing calculator. The pricing tool can calculate migration, monthly estimates, and future pricing estimates based on the workload-driven input that you specify. In general, the cost of block blob storage depends on:

- Volume of data stored per month.
- Quantity and types of operations performed, along with any data transfer costs.
- Data redundancy option selected.

You can use the Azure Pricing Calculator to estimate your storage costs.



Things to know about pricing for Blob Storage

Review the following billing considerations for an Azure storage account and Blob Storage.

- **Performance tiers.** The Blob Storage tier determines the amount of data stored and the cost for storing that data. As the performance tier gets cooler, the per-gigabyte cost decreases.
- **Data access costs.** Data access charges increase as the tier gets cooler. For data in the Cool, Cold, and Archive tiers, you're billed a per-gigabyte data access charge for read actions.
- **Transaction costs.** There's a per-transaction charge for all tiers. The charge increases as the tier gets cooler.
- **Geo-replication data transfer costs.** This charge only applies to accounts that have geo-replication configured. Geo-replication data transfer incurs a per-gigabyte charge.
- **Outbound data transfer costs.** Outbound data transfers incur billing for bandwidth usage on a per-gigabyte basis. This billing is consistent with general-purpose Azure storage accounts.
- **Changes to the storage tier.** If you change the account storage tier from Cool to Hot, you incur a charge equal to reading all the data existing in the storage account. Changing the account storage tier from Hot to Cool incurs a charge equal to writing all the data into the Cool tier (GPv2 accounts only).

9. Exercise - Provide storage for the public website

Exercise - Provide storage for the public website

Completed

Exercise scenario

The company website supplies product images, videos, marketing literature, and customer success stories. Customers are located worldwide and demand is rapidly expanding. The content is mission-critical and requires low latency load times. It's important to keep track of the document versions and to quickly restore documents if they're deleted.

Job skills

- Create a storage account with high availability.
- Ensure the storage account has anonymous public access.
- Create a blob storage container for the website documents.
- Enable soft delete so files can be easily restored.
- Enable blob versioning.

Note

Estimated time: 30 minutes. To complete this exercise, you need an [Azure subscription](#).

Launch the exercise, and follow the instructions. When finished, be sure to return to this page so you can continue learning.

[Launch Exercise](#)

10. Module assessment

Module assessment

Completed

11. Summary and resources

<https://learn.microsoft.com/en-us/training/modules/configure-blob-storage/11-summary-resources>

Summary and resources

Completed

In this module, you learned about Azure Blob Storage and how to configure it. You discovered that Blob Storage is Microsoft's object storage solution for the cloud. You learned Azure blob storage is optimized for storing massive amounts of unstructured data like text or binary files. You explored the features of Blob Storage and its use cases. You also learned how to configure Blob Storage, including choosing the appropriate access tiers to reduce cost and improve performance. And, you learned about creating a lifecycle management strategy, and configuring object replication for failover.

The main takeaways from this module are:

- Azure Blob Storage is a powerful solution for storing unstructured data in the cloud, such as text documents, images, and videos.
- Blob Storage offers different access tiers (Hot, Cool, Cold, and Archive) to optimize performance and cost based on the usage patterns of your data.
- You can configure lifecycle management policies to automatically transition data between access tiers and set expiration times for data.
- Object replication allows you to asynchronously copy blobs between containers in different regions, providing redundancy and reducing latency for read requests.

Learn more with Copilot

Copilot can assist you in configuring Azure infrastructure solutions. Copilot can compare, recommend, explain, and research products and services where you need more information. Open a Microsoft Edge browser and choose Copilot (top right) or navigate to copilot.microsoft.com. Take a few minutes to try these prompts and extend your learning with Copilot.

- What are common administration tasks associated with Azure blob storage?
- How is Azure blob storage priced?

Learn more with Azure documentation

- [Azure Blob Storage documentation](#) - Microsoft Azure's official documentation provides comprehensive information on configuring and managing blob storage. You can find detailed guides, tutorials, and examples to help you navigate through different aspects of blob storage configuration.
- [Azure Blob Storage Concepts](#) - This article provides an overview of the key concepts related to Azure Blob Storage, including storage accounts, containers, and blobs. It explains how to create and manage these entities and covers various configuration options.
- [Azure Blob Storage Security](#) - Understanding the security aspects of blob storage is crucial for proper configuration. This article explores authentication, authorization, and encryption options available in Azure Blob Storage. It also covers best practices for securing your blob storage resources.
- [Azure Blob Storage Performance and Scalability](#) - This article delves into performance considerations when configuring blob storage. The module covers the storage account type, and optimizing data transfer.
- [Azure Blob Storage Lifecycle Management](#) - Blob storage lifecycle management allows you to automate the movement and deletion of data based on predefined rules. This article explains how to configure and manage lifecycle policies to optimize storage costs and improve data management.