

Describe Azure storage services

1. Introduction

<https://learn.microsoft.com/en-us/training/modules/describe-azure-storage-services/1-introduction>

Introduction

Completed

In this module, you'll be introduced to the Azure storage services. You'll learn about the Azure Storage Account and how that relates to the different storage services that are available. You'll also learn about blob storage tiers, data redundancy options, and ways to move data or even entire infrastructures to Azure.

Learning objectives

After completing this module, you'll be able to:

- Compare Azure storage services.
- Describe storage tiers.
- Describe redundancy options.
- Describe storage account options and storage types.
- Identify options for moving files, including AzCopy, Azure Storage Explorer, and Azure File Sync.
- Describe migration options, including Azure Migrate and Azure Data Box.

2. Describe Azure storage accounts


<https://learn.microsoft.com/en-us/training/modules/describe-azure-storage-services/2-accounts>

Describe Azure storage accounts

Completed

The following video introduces the services available with Azure Storage.

A storage account provides a unique namespace for your Azure Storage data that's accessible from anywhere in the world over HTTP or HTTPS. Data in this account is secure, highly available, durable, and massively scalable.



When you create your storage account, you'll start by picking the storage account type. The type of account determines the storage services and redundancy options and has an impact on the use cases. These redundancy options are covered later in this module:

- Locally redundant storage (LRS)
- Geo-redundant storage (GRS)
- Read-access geo-redundant storage (RA-GRS)
- Zone-redundant storage (ZRS)
- Geo-zone-redundant storage (GZRS)
- Read-access geo-zone-redundant storage (RA-GZRS)

Type	Supported services	Redundancy Options	Usage
Standard general-purpose v2	Blob Storage (including Data Lake Storage), Queue Storage, Table Storage, and Azure Files	LRS, GRS, RA-GRS, ZRS, GZRS, RA-GZRS	Standard storage account type for blobs, file shares, queues, and tables. Recommended for most scenarios using Azure Storage. If you want support for network file system (NFS) in Azure Files, use the premium file shares account type.
Premium block blobs	Blob Storage (including Data Lake Storage)	LRS, ZRS	Premium storage account type for block blobs and append blobs. Recommended for scenarios with high transaction rates or that use smaller objects or require consistently low storage latency.
Premium file shares	Azure Files	LRS, ZRS	Premium storage account type for file shares only. Recommended for high-scale or high-performance applications. Use this account type if you want a storage account that supports both Server Message Block (SMB) and NFS file shares.
Premium page blobs	Page blobs only	LRS	Premium storage account type for page blobs only.

Azure Storage Account Types

Standard GP v2	Premium Block Blobs	Premium File Shares	Premium Page Blobs
Services Blob (+ Data Lake) Queue Storage Table Storage Azure Files	Services Blob (+ Data Lake)	Services Azure Files	Services Page blobs only
Redundancy LRS, GRS, RA-GRS ZRS, GZRS, RA-GZRS	Redundancy LRS, ZRS	Redundancy LRS, ZRS	Redundancy LRS
Best For Most scenarios. Recommended default for blobs, files, queues, and tables.	Best For High transaction rates, small objects, or low latency requirements.	Best For High-scale or high-performance apps. SMB + NFS support.	Best For Page blob workloads (e.g., VM disks).

Standard GP v2 is recommended for most scenarios. Premium types offer lower latency for specialized workloads.

Storage account endpoints


One of the benefits of using an Azure storage account is having a unique namespace in Azure for your data. Every storage account must have a unique account name within Azure. The combination of the account name and the Azure Storage service endpoint forms the endpoints for your storage account.

When naming your storage account, keep these rules in mind:

- Storage account names must be between 3 and 24 characters in length and may contain numbers and lowercase letters only.
- Your storage account name must be unique within Azure. No two storage accounts can have the same name. This supports the ability to have a unique, accessible namespace in Azure.

The following table shows the endpoint format for Azure Storage services.

Storage service	Endpoint
Blob Storage	https://<storage-account-name>.blob.core.windows.net
Data Lake Storage Gen2	https://<storage-account-name>.dfs.core.windows.net
Azure Files	https://<storage-account-name>.file.core.windows.net
Queue Storage	https://<storage-account-name>.queue.core.windows.net
Table Storage	https://<storage-account-name>.table.core.windows.net



3. Describe Azure storage redundancy

<https://learn.microsoft.com/en-us/training/modules/describe-azure-storage-services/3-redundancy>

Describe Azure storage redundancy

Completed

Azure Storage keeps multiple copies of data to protect against hardware failures, outages, and regional disasters. Redundancy choices determine availability, durability, and cost.

When selecting a redundancy option, balance lower cost against higher availability. Key decision factors include:


- How your data is replicated in the primary region.
- Whether your data is replicated to a second region that is geographically distant to the primary region, to protect against regional disasters.
- Whether your application requires read access to the replicated data in the secondary region if the primary region becomes unavailable.

Redundancy in the primary region

Data in Azure Storage is always replicated three times in the primary region. Primary-region options are locally redundant storage (LRS) and zone-redundant storage (ZRS).

Locally redundant storage


Locally redundant storage (LRS) replicates your data three times within a single data center in the primary region. LRS provides at least 11 nines of durability (99.999999999%) of objects over a given year.



LRS is the lowest-cost option and protects against server rack and drive failures in one datacenter. Because all replicas stay in that datacenter, a datacenter-level disaster can still cause data loss. For stronger resilience, use ZRS, GRS, or GZRS.

Zone-redundant storage

For Availability Zone-enabled Regions, zone-redundant storage (ZRS) replicates your Azure Storage data synchronously across three Azure availability zones in the primary region. ZRS offers durability for Azure Storage data objects of at least 12 nines (99.9999999999%) over a given year.



With ZRS, data stays available for read and write operations even if one zone is unavailable. Azure performs networking updates, such as DNS repointing, during recovery.

Microsoft recommends ZRS for high availability in-region scenarios and for requirements that keep replication within a country or region.

Redundancy in a secondary region

For higher durability, you can replicate data to a secondary region that is geographically distant from the primary region.

When you create a storage account, you choose the primary region. Azure assigns the paired secondary region based on region pairs.

Secondary-region options are geo-redundant storage (GRS) and geo-zone-redundant storage (GZRS). GRS uses LRS in both regions, while GZRS uses ZRS in the primary region and LRS in the secondary region.


By default, secondary-region data isn't readable unless failover occurs. If the primary region is unavailable, you can fail over so the secondary region becomes primary.

Important

Because data is replicated to the secondary region asynchronously, a failure that affects the primary region may result in data loss if the primary region can't be recovered. The interval between the most recent writes to the primary region and the last write to the secondary region is known as the recovery point objective (RPO). The RPO indicates the point in time to which data can be recovered. Azure Storage typically has an RPO of less than 15 minutes, although there's currently no SLA on how long it takes to replicate data to the secondary region.


Geo-redundant storage

GRS copies data synchronously three times in the primary region (LRS), then asynchronously to the secondary region (also LRS). It provides at least 16 nines of durability over a year.



Geo-zone-redundant storage

GZRS combines zone-level resilience in the primary region with geo-replication to a secondary region. Data is copied across three availability zones in the primary region and replicated to the paired secondary region using LRS. Microsoft recommends GZRS for workloads that need maximum consistency, availability, and disaster recovery resilience.



GZRS is designed to provide at least 16 nines (99.99999999999999%) of durability of objects over a given year.

Read access to data in the secondary region

GRS and GZRS protect against regional outages by replicating data to a secondary location. To read secondary-region data before failover, enable read-access geo-redundant storage (RA-GRS) or read-access geo-zone-redundant storage (RA-GZRS).

Important

Remember that the data in your secondary region may not be up-to-date due to RPO.

4. Describe Azure storage services


<https://learn.microsoft.com/en-us/training/modules/describe-azure-storage-services/4-describe-azure-storage-services>

Describe Azure storage services

Completed

Azure Storage includes these core data services:

- **Azure Blobs:** A massively scalable object store for text and binary data. Also includes support for big data analytics through Data Lake Storage Gen2.
- **Azure Files:** Managed file shares for cloud or on-premises deployments.
- **Azure Queues:** A messaging store for reliable messaging between application components.
- **Azure Disks:** Block-level storage volumes for Azure VMs.
- **Azure Tables:** NoSQL table option for structured, non-relational data.



Benefits of Azure Storage

Azure Storage provides these benefits:

- **Durable and highly available.** Redundancy options protect data from hardware failures, outages, and regional incidents.
- **Secure.** All data written to an Azure storage account is encrypted by the service. Azure Storage provides you with fine-grained control over who has access to your data.
- **Scalable.** Azure Storage is designed to be massively scalable to meet the data storage and performance needs of today's applications.
- **Managed.** Azure handles hardware maintenance, updates, and critical issues for you.

- **Accessible.** Data can be accessed globally over HTTP or HTTPS by using REST APIs, SDKs, Azure CLI, Azure PowerShell, the Azure portal, or Azure Storage Explorer.

Azure Blobs

Azure Blob storage is an unstructured object storage service for large volumes of text or binary data. It's designed for scenarios such as large uploads, media content, log files, and analytics datasets. Developers store objects as blobs while Azure handles the underlying storage infrastructure.

Blob storage is ideal for:

- Serving images or documents directly to a browser.
- Storing files for distributed access.
- Streaming video and audio.
- Storing data for backup and restore, disaster recovery, and archiving.
- Storing data for analysis by an on-premises or Azure-hosted service.

Accessing blob storage

Objects in blob storage can be accessed from anywhere in the world via HTTP or HTTPS. Users or client applications can access blobs via URLs, the Azure Storage REST API, Azure PowerShell, Azure CLI, or an Azure Storage client library. The storage client libraries are available for multiple languages, including .NET, Java, Node.js, Python, PHP, and Ruby.


Blob storage tiers

Data access patterns change over time, so Azure Blob storage offers access tiers that balance storage cost against retrieval speed. The available access tiers include:

- **Hot access tier:** Optimized for storing data that is accessed frequently (for example, images for your website).
- **Cool access tier:** Optimized for data that is infrequently accessed and stored for at least 30 days (for example, invoices for your customers).
- **Cold access tier:** Optimized for storing data that is infrequently accessed and stored for at least 90 days.
- **Archive access tier:** Appropriate for data that is rarely accessed and stored for at least 180 days, with flexible latency requirements (for example, long-term backups).

Considerations for access tiers:

- Hot, cool, and cold access tiers can be set at the account level. The archive access tier isn't available at the account level.
- Hot, cool, cold, and archive tiers can be set at the blob level, during or after upload.
- Cool and cold tiers have lower storage costs but higher access costs and lower availability SLAs.
- Archive tier has the lowest storage cost and highest rehydration latency and access cost.



Azure Files

Azure Files offers fully managed file shares in the cloud that are accessible via the industry standard Server Message Block (SMB) or Network File System (NFS) protocols. Azure file shares can be mounted concurrently by cloud or on-premises deployments. SMB Azure file shares are accessible from Windows, Linux, and macOS clients. NFS Azure file shares are accessible from Linux or macOS clients. Additionally, SMB Azure file shares can be cached on Windows Servers with Azure File Sync for fast access near where the data is being used.

Azure Files key benefits

- **Shared access:** SMB and NFS protocol support enables compatibility with existing applications.
- **Fully managed:** No server hardware or OS patching to manage.
- **Scripting and tooling:** Manage shares with Azure CLI, Azure PowerShell, the Azure portal, and Storage Explorer.
- **Resiliency:** Built for high availability.
- **Familiar programmability:** Applications can use standard file I/O APIs plus Azure SDKs and REST APIs.

Azure Queues

Azure Queue storage stores large numbers of messages for asynchronous processing. Queues are accessed by authenticated HTTP/HTTPS calls, can hold millions of messages, and support messages up to 64 KB.

Queue storage is commonly paired with Azure Functions so messages trigger background actions.

Azure Disks

Azure Disk storage (managed disks) provides block-level volumes for Azure VMs. They are virtualized and managed by Azure for improved resiliency and simpler operations.

Azure Tables

Azure Table storage is a NoSQL store for large amounts of structured, non-relational data, accessible through authenticated calls from cloud and hybrid environments.

5. Identify Azure data migration options

<https://learn.microsoft.com/en-us/training/modules/describe-azure-storage-services/6-identify-azure-data-migration-options>

Identify Azure data migration options

Completed

Now that you understand the different storage options within Azure, it's important to also understand how to get your data into Azure. Azure supports both real-time migration of infrastructure, applications, and data using Azure Migrate as well as asynchronous migration of data using Azure Data Box.

Azure Data Migration Options

Azure Migrate	Azure Data Box
Online Real-time migration hub Single portal to assess and migrate on-premises infrastructure to Azure Discovery and assessment Server migration Database migration Web app migration ISV tool integration	Physical Offline bulk data transfer Ship a secure storage device with up to 80 TB capacity One-time bulk migration Periodic large uploads Data export from Azure NIST 800-88r1 disk wipe End-to-end tracking in portal

Azure Migrate

Azure Migrate is a service that helps you migrate from an on-premises environment to the cloud. Azure Migrate functions as a hub to help you manage the assessment and migration of your on-

premises datacenter to Azure. It provides the following:

- **Unified migration platform:** A single portal to start, run, and track your migration to Azure.
- **Range of tools:** A range of tools for assessment and migration. Azure Migrate tools include Azure Migrate: Discovery and assessment and Azure Migrate: Server Migration. Azure Migrate also integrates with other Azure services and tools, and with independent software vendor (ISV) offerings.
- **Assessment and migration:** In the Azure Migrate hub, you can assess and migrate your on-premises infrastructure to Azure.

Integrated tools

In addition to working with tools from ISVs, the Azure Migrate hub includes integrated options for discovery and assessment, server migration, database migration, and web app migration.

Azure Data Box

Azure Data Box is a physical migration service that helps transfer large amounts of data in a quick, inexpensive, and reliable way. The secure data transfer is accelerated by shipping you a proprietary Data Box storage device that has a maximum usable storage capacity of 80 terabytes. The Data Box is transported to and from your datacenter via a regional carrier. A rugged case protects and secures the Data Box from damage during transit.

You can order the Data Box device via the Azure portal to import or export data from Azure. Once the device is received, you can quickly set it up using the local web UI and connect it to your network. Once you're finished transferring the data (either into or out of Azure), simply return the Data Box. If you're transferring data into Azure, the data is automatically uploaded once Microsoft receives the Data Box back. The entire process is tracked end-to-end by the Data Box service in the Azure portal.

Use cases

Data Box is ideally suited for moving very large data sets (often tens of terabytes or more) when network bandwidth is limited.

Common Data Box scenarios include:

- One-time bulk migration of on-premises data to Azure.
- Periodic large data uploads where online transfer is too slow.
- Exporting large data sets from Azure for recovery or regulatory needs.

Once the data from your import order is uploaded to Azure, the disks on the device are wiped clean in accordance with NIST 800-88r1 standards. For an export order, the disks are erased once the device reaches the Azure datacenter.


6. Identify Azure file movement options

<https://learn.microsoft.com/en-us/training/modules/describe-azure-storage-services/7-identify-azure-file-movement-options>

Identify Azure file movement options

Completed

In addition to large-scale migration using services like Azure Migrate and Azure Data Box, Azure also has tools designed to help you move or interact with individual files or small file groups. Among those tools are AzCopy, Azure Storage Explorer, and Azure File Sync.



AzCopy

AzCopy is a command-line utility that you can use to copy blobs or files to or from your storage account. With AzCopy, you can upload files, download files, copy files between storage accounts, and even synchronize files. AzCopy can even be configured to work with other cloud providers to help move files back and forth between clouds.

Important

Synchronizing blobs or files with AzCopy is one-direction synchronization. When you synchronize, you designate the source and destination, and AzCopy will copy files or blobs in that direction. It doesn't synchronize bi-directionally based on timestamps or other metadata.

Azure Storage Explorer

Azure Storage Explorer is a standalone app that provides a graphical interface to manage files and blobs in your Azure Storage Account. It works on Windows, macOS, and Linux operating systems and uses AzCopy on the backend to perform all of the file and blob management tasks. With Storage Explorer, you can upload to Azure, download from Azure, or move between storage accounts.

Azure File Sync

Azure File Sync is a tool that lets you centralize your file shares in Azure Files and keep the flexibility, performance, and compatibility of a Windows file server. It's almost like turning your Windows file server into a miniature content delivery network. Once you install Azure File Sync on your local Windows server, it will automatically stay bi-directionally synced with your files in Azure.

With Azure File Sync, you can:

- Use any protocol that's available on Windows Server to access your data locally, including SMB, NFS, and FTPS.
- Have as many caches as you need across the world.
- Replace a failed local server by installing Azure File Sync on a new server in the same datacenter.
- Configure cloud tiering so the most frequently accessed files are replicated locally, while infrequently accessed files are kept in the cloud until requested.

7. Module assessment

<https://learn.microsoft.com/en-us/training/modules/describe-azure-storage-services/8-knowledge-check>

Module assessment

Completed

8. Summary

Summary

Completed

In this module, you learned about Azure storage services. You learned about Azure storage accounts and how they relate to different storage services. You explored storage blobs, redundancy options, and ways to migrate and move your data both into and within Azure.

Learning objectives

You should now be able to:

- Compare Azure storage services.
- Describe storage tiers.
- Describe redundancy options.
- Describe storage account options and storage types.
- Identify options for moving files, including AzCopy, Azure Storage Explorer, and Azure File Sync.
- Describe migration options, including Azure Migrate and Azure Data Box.

Additional resources

The following resources provide more information on topics in this module or related to this module.

- [Store data in Azure](#) is a Microsoft Learn course that covers more information about storing data in Azure.
- [Microsoft Certified: Azure Data Fundamentals](#) is an entire certification, with associated training that dives deeper into data fundamentals on Azure.

Explore with Copilot

Tip

Try one of these prompts in Copilot Chat:

- "Map four workload patterns to the right Azure storage service and justify each choice for performance, durability, and access requirements."
- "Compare LRS, ZRS, GRS, and GZRS for a business-critical workload and recommend one option based on RPO, availability, and budget constraints."

- "Create a migration checklist that uses Azure Migrate, Data Box, AzCopy, Storage Explorer, or File Sync where each tool fits best."