

Azure File Sync


Azure File Sync to centralize your organization's file shares in Azure Files, while keeping the flexibility, performance, and compatibility of an on-premises file server. Azure File Sync transforms Windows Server into a quick cache of your Azure file share. You can use any protocol that's available on Windows Server to access your data locally, including SMB, NFS, and FTPS. You can have as many caches as you need across the world.

Prepare Windows Server to use with Azure File Sync

Open Server Manager.

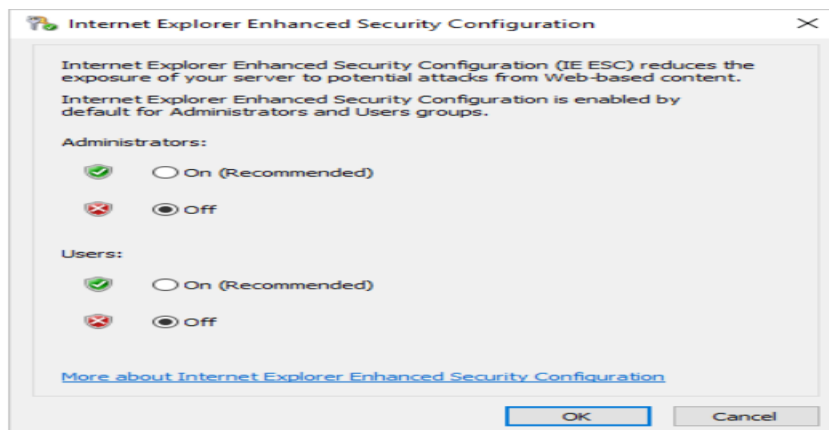
Click Local Server:

"Local Server" on the left side of the Server Manager UI



On the Properties subpane, select the link for IE Enhanced Security Configuration.

In the Internet Explorer Enhanced Security Configuration dialog box, select Off for Administrators and Users:

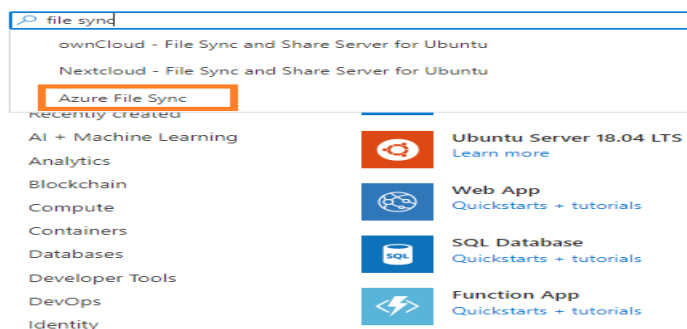


Deploy the Storage Sync Service

The deployment of Azure File Sync starts with placing a Storage Sync Service resource into a resource group of your selected subscription. You will create a trust relationship between your servers and this resource and a server can only be registered to one Storage Sync Service. As a result, it is recommended to deploy as many storage sync services as you need to separate groups of servers. Keep in mind that servers from different storage sync services cannot sync with each other.

[Home](#) >

New ...





Home > New >

Azure File Sync

Microsoft



Azure File Sync [Add to Favorites](#)

Microsoft
★★★★☆ 4.2 (19 ratings)

Create


[Overview](#) [Plans](#) [Usage Information + Support](#) [Reviews](#)

Azure File Sync extends on-premises files servers into Azure providing clo

Azure File Sync provides:

- Multi-site access - provide write access to the same data across Wi
- Cloud tiering - store only recently accessed data on local servers
- Integrates with Azure backup - no need to back up your data on pr
- Fast disaster recovery - restore file metadata immediately and reca

Media



Name: A unique name (per region) for the Storage Sync Service.

Subscription: The subscription in which you want to create the Storage Sync Service. Depending on your organization's configuration strategy, you might have access to one or more subscriptions. An Azure subscription is the most basic container for billing for each cloud service (such as Azure Files).

Resource group: A resource group is a logical group of Azure resources, such as a storage account or a Storage Sync Service. You can create a new resource group or use an existing resource group for Azure File Sync. (We recommend using resource groups as containers to isolate resources logically for your organization, such as grouping HR resources or resources for a specific project.)


Location: The region in which you want to deploy Azure File Sync. Only supported regions are available in this list.

Home > New > Azure File Sync >

Deploy Azure File Sync

[Basics](#) [Tags](#) [Review + create](#)

Azure File Sync in combination with Azure file shares allows you to centralize your organization's file shares in Azure, while keeping the flexibility, performance, and compatibility of an on-premises file server. [Learn more](#)



Deploying this storage sync service resource will allow you to transform your Windows Server into a quick cache for Azure file shares with optional cloud tiering and multi-server sync functionality. Keep in mind that servers registered to different storage sync service resources cannot exchange data with each other. It's best to register all servers to the same storage sync service if they will ever have a need to sync the same Azure file share.

Subscription *	<input type="text" value="Pay-As-You-Go"/>
Resource group *	<input type="text" value="rg-asegk-00"/> Create new
Storage sync service name *	<input type="text" value="asegkfilesync"/>
Region *	<input type="text" value="South India"/>

Create to deploy the Storage Sync Service.

Install the Azure File Sync agent

The Azure File Sync agent is a downloadable package that enables Windows Server to be synced with an Azure file share.

<https://go.microsoft.com/fwlink/?linkid=858257>

Operating system requirements

Azure File Sync is supported with the following versions of Windows Server:

Version Supported SKUs Supported deployment options

Windows Server 2019 Datacenter, Standard, and IoT Full and Core

Windows Server 2016 Datacenter, Standard, and Storage Server Full and Core

Windows Server 2012 R2 Datacenter, Standard, and Storage Server Full and Core

Recommended system resources

Namespace size - files & directories (millions) Typical capacity (TiB) CPU Cores

Recommended memory (GiB)

3	1.4	2	8 (initial sync)/ 2 (typical churn)
5	2.3	2	16 (initial sync)/ 4 (typical churn)
10	4.7	4	32 (initial sync)/ 8 (typical churn)
30	14.0	8	48 (initial sync)/ 16 (typical churn)
50	23.3	16	64 (initial sync)/ 32 (typical churn)
100*	46.6	32	128 (initial sync)/ 32 (typical churn)

Register Windows Server with Storage Sync Service

Registering Windows Server with a Storage Sync Service establishes a trust relationship between your server (or cluster) and the Storage Sync Service. A server can only be registered to one Storage Sync Service and can sync with other servers and Azure file shares associated with the same Storage Sync Service.

Microsoft Azure File Sync - Server Registration

Choose a Storage Sync Service

Azure Subscription

Subscription ID:

Resource Group

Storage Sync Service

Register

Create a sync group

A sync group defines the sync topology for a set of files. Endpoints within a sync group are kept in sync with each other. A sync group must contain one cloud endpoint, which represents an Azure file share and one or more server endpoints. A server endpoint represents a path on a registered server. A server can have server endpoints in multiple sync groups. You can create as many sync groups as you need to appropriately describe your desired sync topology.

A cloud endpoint is a pointer to an Azure file share. All server endpoints will sync with a cloud endpoint, making the cloud endpoint the hub. The storage account for the Azure file share must be located in the same region as the Storage Sync Service. The entirety of the Azure file share will be synced, with one exception: A special folder, comparable to the hidden "System Volume Information" folder on an NTFS volume, will be provisioned. This directory is called ".SystemShareInformation". It contains important sync metadata that will not sync to other endpoints. Do not use or delete it!

NB: You can make changes to any cloud endpoint or server endpoint in the sync group and have your files synced to the other endpoints in the sync group. If you make a change to the cloud endpoint (Azure file share) directly, changes first need to be discovered by an Azure File Sync change detection job. A change detection job is initiated for a cloud endpoint only once every 24 hours.

Sync group name: The name of the sync group to be created. This name must be unique within the Storage Sync Service, but can be any name that is logical for you.

Subscription: The subscription where you deployed the Storage Sync Service in Deploy the Storage Sync Service.

Storage account: If you select Select storage account, another pane appears in which you can select the storage account that has the Azure file share that you want to sync with.

Azure file share: The name of the Azure file share with which you want to sync.

Create a server endpoint

A server endpoint represents a specific location on a registered server, such as a folder on a server volume. A server endpoint must be a path on a registered server (rather than a mounted share), and to use cloud tiering (Cloud tiering, an optional feature of Azure File Sync, decreases the amount of local storage required while keeping the performance of an on-premises file server. When enabled, this feature stores only frequently accessed (hot) files on your local server. Infrequently accessed (cool) files are split into namespace (file and folder structure) and file content. The namespace is stored locally and the file content stored in an Azure file share in the cloud. When a user opens a tiered file, Azure File Sync seamlessly recalls the file data from the file share in Azure.), the path must be on a non-system volume. Network attached storage (NAS) is not supported.

To add a server endpoint, go to the newly created sync group and then select Add server endpoint. In the Add server endpoint pane, enter the following information to create a server endpoint:

Registered server: The name of the server or cluster where you want to create the server endpoint.

Path: The Windows Server path to be synced as part of the sync group.

Cloud Tiering: A switch to enable or disable cloud tiering. With cloud tiering, infrequently used or accessed files can be tiered to Azure Files.

Volume Free Space: The amount of free space to reserve on the volume on which the server endpoint is located. For example, if volume free space is set to 50% on a volume that has a single server endpoint, roughly half the amount of data is tiered to Azure Files. Regardless of whether cloud tiering is enabled, your Azure file share always has a complete copy of the data in the sync group.

Initial download mode: This is an optional selection, starting with agent version 11, that can be helpful when there are files in the Azure file share but not on the server. Such a situation can exist, for instance, if you create a server endpoint to add another branch office server to a sync group or when you disaster-recover a failed server. If cloud tiering is enabled, the default is to only recall the namespace, no file content initially. That is useful if you believe that rather user access requests should decide what file content is recalled to the server. If cloud tiering is disabled, the default is that the namespace will download first and then files will be recalled based on last-modified timestamp until the local capacity has been reached. You can however change the initial download mode to namespace only. A third mode can only be used if cloud tiering is disabled for this server endpoint. This mode avoids recalling the namespace first. Files will only appear on the local server if they had a chance to fully download. This mode is useful if for instance an application requires full files to be present and cannot tolerate tiered files in it's namespace.

To add the server endpoint, select Create. Your files are now kept in sync across your Azure file share and Windows Server.

Configure firewall and virtual network settings

Portal

If you'd like to configure your Azure File sync to work with firewall and virtual network settings, do the following:

From the Azure portal, navigate to the storage account you want to secure.

Select the Firewalls and virtual networks button on the left menu.

Select Selected networks under Allow access from.

Make sure your servers IP or virtual network is listed under the appropriate section.

Make sure Allow trusted Microsoft services to access this storage account is checked.

Select Save to save your settings.

Security

All Azure Files data is encrypted at rest. For encryption in transit, Azure provides a layer of encryption for all data in transit between Azure Datacenters using MACSec. Through this, encryption exists when data is transferred between Azure datacenters. Unlike Azure Files using the SMB protocol, file shares using the NFS protocol do not offer user-based authentication. Authentication for NFS shares is based on the configured network security rules. Due to this, to ensure only secure connections are established to your NFS share, you must use either service endpoints or private endpoints. If you want to access shares from on-premises then, in addition to a private endpoint, you must setup a VPN or ExpressRoute. Requests that do not originate from the following sources will be rejected:

A private endpoint

Azure VPN Gateway

Point-to-site (P2S) VPN

Site-to-Site

ExpressRoute

A restricted public endpoint

SMB shares

Azure file shares mounted with SMB offer more Azure Files features and have no Azure Files feature restrictions since it is generally available.

Features

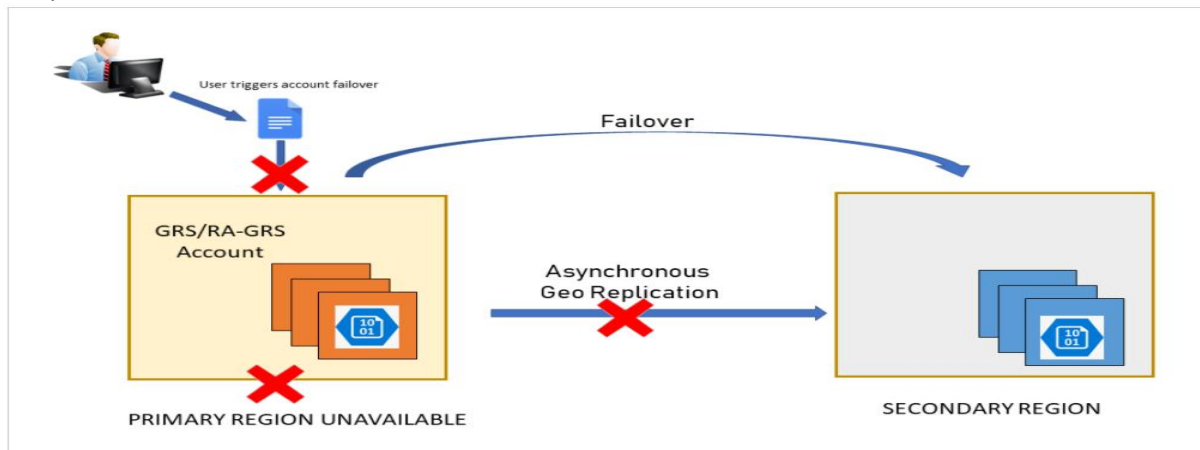
Azure file sync

Identity-based authentication

Azure Backup support
Snapshots
Soft delete
Encryption-in-transit and encryption-at-rest

Disaster recovery and storage account failover

Azure Storage supports account failover for geo-redundant storage accounts. With account failover, you can initiate the failover process for your storage account if the primary endpoint becomes unavailable. The failover updates the secondary endpoint to become the primary endpoint for your storage account. Once the failover is complete, clients can begin writing to the new primary endpoint.



Account failover is available for general-purpose v1, general-purpose v2, and Blob storage account types with Azure Resource Manager deployments. Account failover is supported for all public regions but is not available in sovereign or national clouds at this time.

Choose the right redundancy option

Azure Storage maintains multiple copies of your storage account to ensure durability and high availability. Which redundancy option you choose for your account depends on the degree of resiliency you need. For protection against regional outages, configure your account for geo-redundant storage, with or without the option of read access from the secondary region:

Geo-redundant storage (GRS) or geo-zone-redundant storage (GZRS) copies your data asynchronously in two geographic regions that are at least hundreds of miles apart. If the primary region suffers an outage, then the secondary region serves as a redundant source for your data. You can initiate a failover to transform the secondary endpoint into the primary endpoint.

Read-access geo-redundant storage (RA-GRS) or read-access geo-zone-redundant storage (RA-GZRS) provides geo-redundant storage with the additional benefit of read access to the secondary endpoint. If an outage occurs in the primary endpoint, applications configured for read access to the secondary and designed for high availability can continue to read from the secondary endpoint. Microsoft recommends RA-GZRS for maximum availability and durability for your applications.

Snapshots for Azure Files

Share snapshots capture the share state at that point in time.

Limits

The maximum number of share snapshots that Azure Files allows today is 200. After 200 share snapshots, you have to delete older share snapshots in order to create new ones.

There is no limit to the simultaneous calls for creating share snapshots. There is no limit to amount of space that share snapshots of a particular file share can consume.

It is not possible to mount share snapshots on Linux. This is because the Linux SMB client does not support mounting snapshots like Windows does.

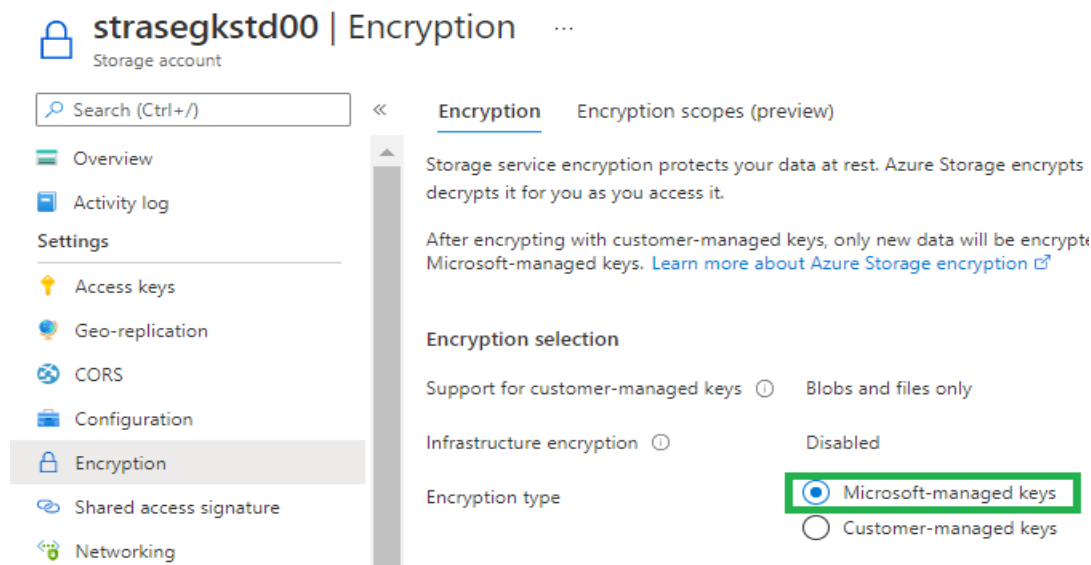
Encryption key model for the storage account

To determine whether a storage account is using Microsoft-managed keys or customer-managed keys for encryption, use one of the following approaches.

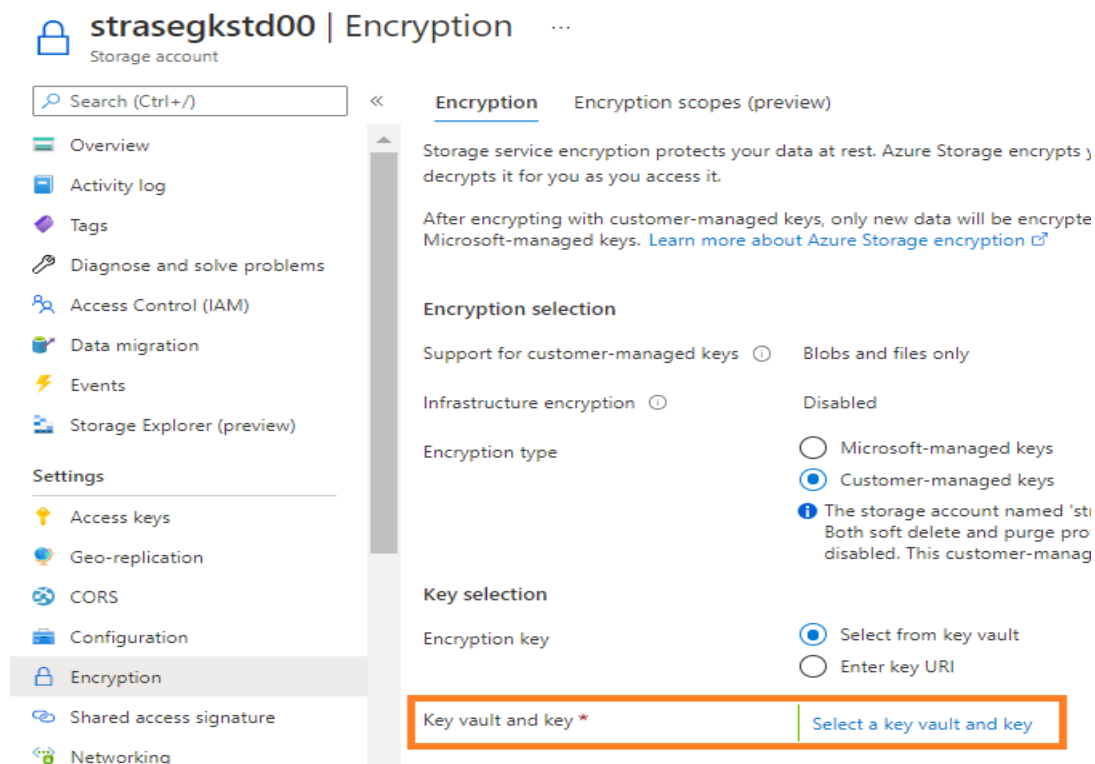
In the Azure portal, navigate to your storage account.

Select the Encryption setting and note the setting.

The following image shows a storage account that is encrypted with Microsoft-managed keys:



And the following image shows a storage account that is encrypted with customer-managed keys:



Configure encryption with customer-managed keys stored in Azure Key Vault
Azure Storage encrypts all data in a storage account at rest. By default, data is encrypted with Microsoft-managed keys. For additional control over encryption keys, you can manage your own keys. Customer-managed keys must be stored in Azure Key Vault or Key Vault Managed Hardware Security Model (HSM) (preview).

Configure a key vault

You can use a new or existing key vault to store customer-managed keys. The storage account and the key vault must be in the same region, but they can be in different subscriptions.

Using customer-managed keys with Azure Storage encryption requires that both soft delete and purge protection be enabled for the key vault. Soft delete is enabled by default when you create a new key vault and cannot be disabled. You can enable purge protection either when you create the key vault or after it is created.

To enable purge protection on an existing key vault, follow these steps:

Navigate to your key vault in the Azure portal.

Under Settings, choose Properties.

In the Purge protection section, choose Enable purge protection.

Add a key

add a key in the key vault(Storage encryption supports RSA and RSA-HSM keys of sizes 2048, 3072 and 4096).

Create key vault

Key vault name *  


Region * 


Pricing tier *  


Recovery options

Soft delete protection will automatically be enabled on this key vault. This feature allows you to recover or permanently delete a key vault and secrets for the duration of the retention period. This protection applies to the key vault and the secrets stored within the key vault.

To enforce a mandatory retention period and prevent the permanent deletion of key vaults or secrets prior to the retention period elapsing, you can turn on purge protection. When purge protection is enabled, secrets cannot be purged by users or by Microsoft.


Soft-delete  Enabled

Days to retain deleted vaults * 

Purge protection 

Disable purge protection (allow key vault and objects to be purged during retention period)

Enable purge protection (enforce a mandatory retention period for deleted vaults and vault objects)

 Once enabled, this option cannot be disabled

Select key from Azure Key Vault

Subscription *

Key vault * [Create new](#)

Key * [Create new](#)

Distributed File System (DFS)

Azure File Sync supports interop with DFS Namespaces (DFS-N) and DFS Replication (DFS-R).

DFS Namespaces (DFS-N): Azure File Sync is fully supported on DFS-N servers. You can install the Azure File Sync agent on one or more DFS-N members to sync data between the server endpoints and the cloud endpoint.

DFS Replication (DFS-R): Since DFS-R and Azure File Sync are both replication solutions, in most cases, we recommend replacing DFS-R with Azure File Sync. There are however several scenarios where you would want to use DFS-R and Azure File Sync together:

You are migrating from a DFS-R deployment to an Azure File Sync deployment.

Not every on-premises server that needs a copy of your file data can be connected directly to the internet.

Branch servers consolidate data onto a single hub server, for which you would like to use Azure File Sync.

For Azure File Sync and DFS-R to work side by side:

Azure File Sync cloud tiering must be disabled on volumes with DFS-R replicated folders.

Server endpoints should not be configured on DFS-R read-only replication folders.

To migrate a DFS-R deployment to Azure File Sync:

Create a sync group to represent the DFS-R topology you are replacing.

Start on the server that has the full set of data in your DFS-R topology to migrate. Install Azure File Sync on that server.

Register that server and create a server endpoint for the first server to be migrated. Do not enable cloud tiering.

Let all of the data sync to your Azure file share (cloud endpoint).

Install and register the Azure File Sync agent on each of the remaining DFS-R servers.

Disable DFS-R.

Create a server endpoint on each of the DFS-R servers. Do not enable cloud tiering.

Ensure sync completes and test your topology as desired.

Retire DFS-R.

Cloud tiering may now be enabled on any server endpoint as desired.