

# Storage Essentials

Click-to-Run™ Step-by-Step Deployment Guide



## Benefits of Storage Essentials

The Storage Essentials Solution is designed to help you deploy and configure three of the most common Storage scenarios: Storage Archiving, Azure File Share and Azure File Sync.

The Storage Archiving is particularly well-suited for the storage of infrequently accessed archival data. It is built on Azure managed service Blob Storage. This service runs in a high availability environment, patched and supported, allowing you to focus on your solution instead of the environment they run in.

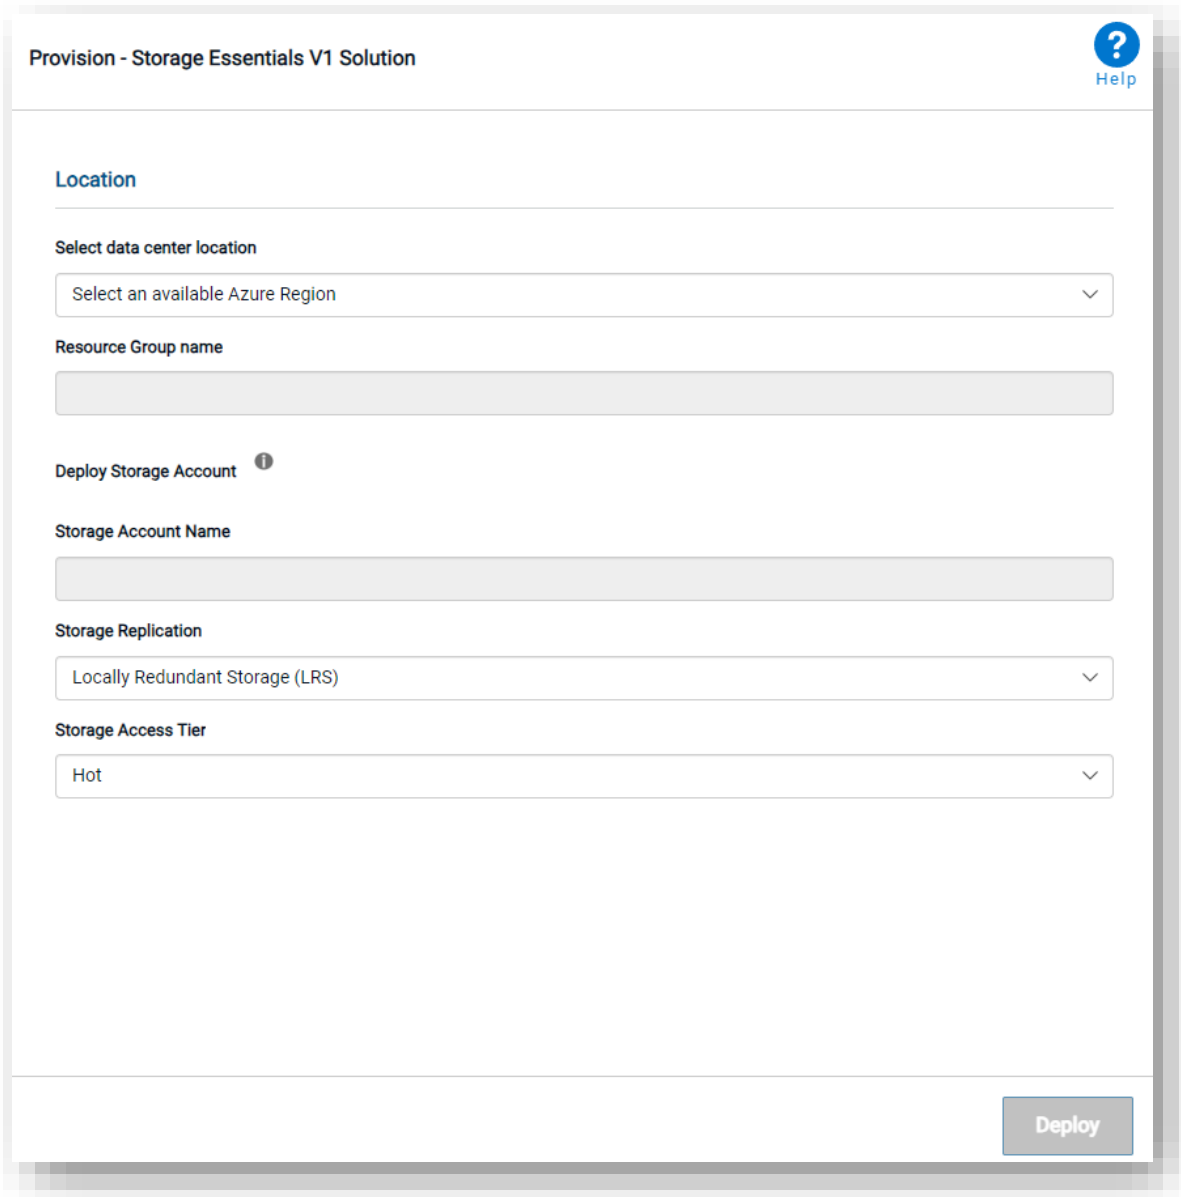
The Azure File Share is designed for centralizing your organization's file shares in Azure Files without giving up the flexibility, performance, and compatibility of an on-premises file server.

Transform your Windows Servers into a quick cache of your Azure file share and access them through SMB or NFS shares on Windows Server. This solution is useful for scenarios in which data needs to be accessed and modified far away from an Azure datacenter, such as in a branch office scenario. Data may be replicated between multiple Windows Server endpoints, such as between multiple branch offices.

Finally, Azure File Sync enables you 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 an on-premises (or cloud) Windows Server into a quick cache of your Azure file share. Azure File Sync allows you to cache a number of Azure file shares on an on-premises Windows Server or cloud VM. These files are stored in Azure file shares.

# Storage Essentials Deployment and Considerations

Purchase the Storage Essentials Click-to-Run™ Solution through StreamOne and proceed to configure and deploy the solution.



Provision - Storage Essentials V1 Solution Help

**Location**

Select data center location

Select an available Azure Region

Resource Group name

Deploy Storage Account ⓘ

Storage Account Name

Storage Replication

Locally Redundant Storage (LRS)

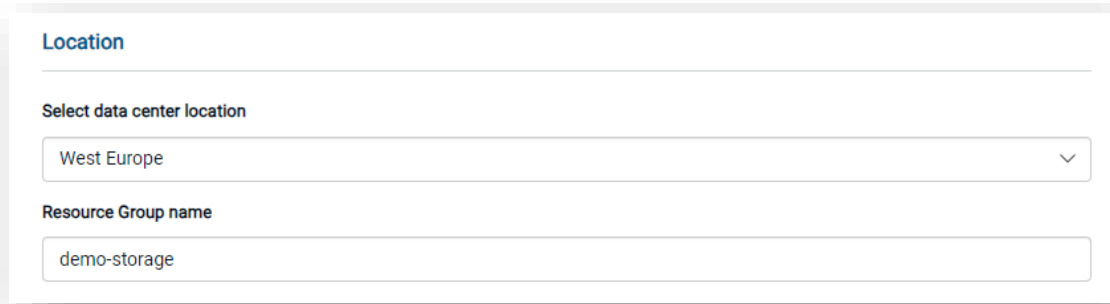
Storage Access Tier

Hot

Deploy

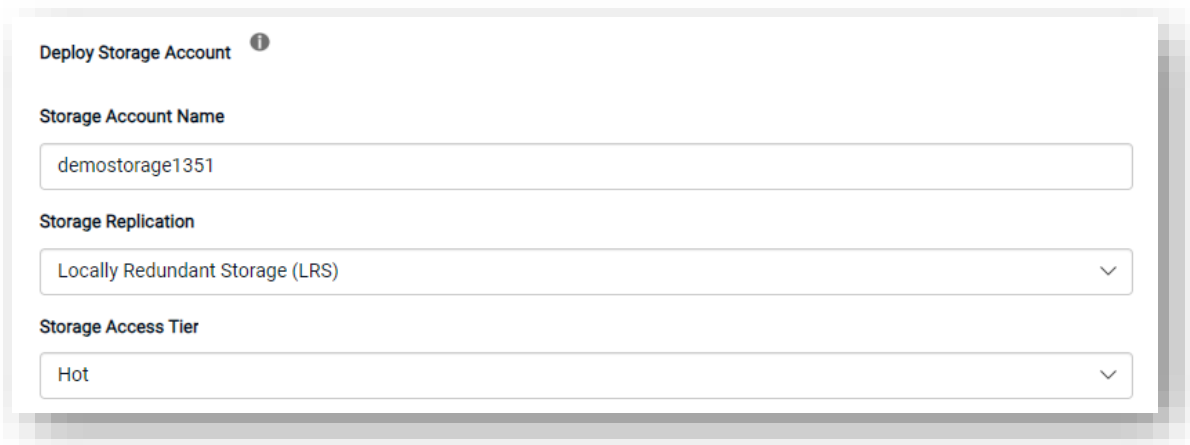
Let's get started with the User Interface!

Please select a location and a resource group name:



The screenshot shows a form titled "Location". It contains two sections: "Select data center location" with a dropdown menu showing "West Europe", and "Resource Group name" with a text input field containing "demo-storage".

You will then be able to define a Storage Account Name, chose a Storage Replication Tier and a Storage Access Tier:



The screenshot shows a form titled "Deploy Storage Account" with an information icon. It contains three sections: "Storage Account Name" with a text input field containing "demostorage1351", "Storage Replication" with a dropdown menu showing "Locally Redundant Storage (LRS)", and "Storage Access Tier" with a dropdown menu showing "Hot".

Storage Replication Tier:

- **Locally redundant storage (LRS)** copies your data synchronously three times within a single physical location in the primary region. LRS is the least expensive replication option but is not recommended for applications requiring high availability.
- **Zone-redundant storage (ZRS)** copies your data synchronously across three Azure availability zones in the primary region. For applications requiring high availability, Microsoft recommends using ZRS in the primary region, and replicating to a secondary region.
- **Geo-redundant storage (GRS)** copies your data synchronously three times within a single physical location in the primary region using LRS. It then copies your data asynchronously to a single physical location in the secondary region.
- **Geo-zone-redundant storage (GZRS)** copies your data synchronously across three Azure availability zones in the primary region using ZRS. It then copies your data asynchronously to a single physical location in the


secondary region.

- For read access to the secondary region, enable **Read-access geo-redundant storage (RA-GRS)** or **Read-access geo-zone-redundant storage (RA-GZRS)**.

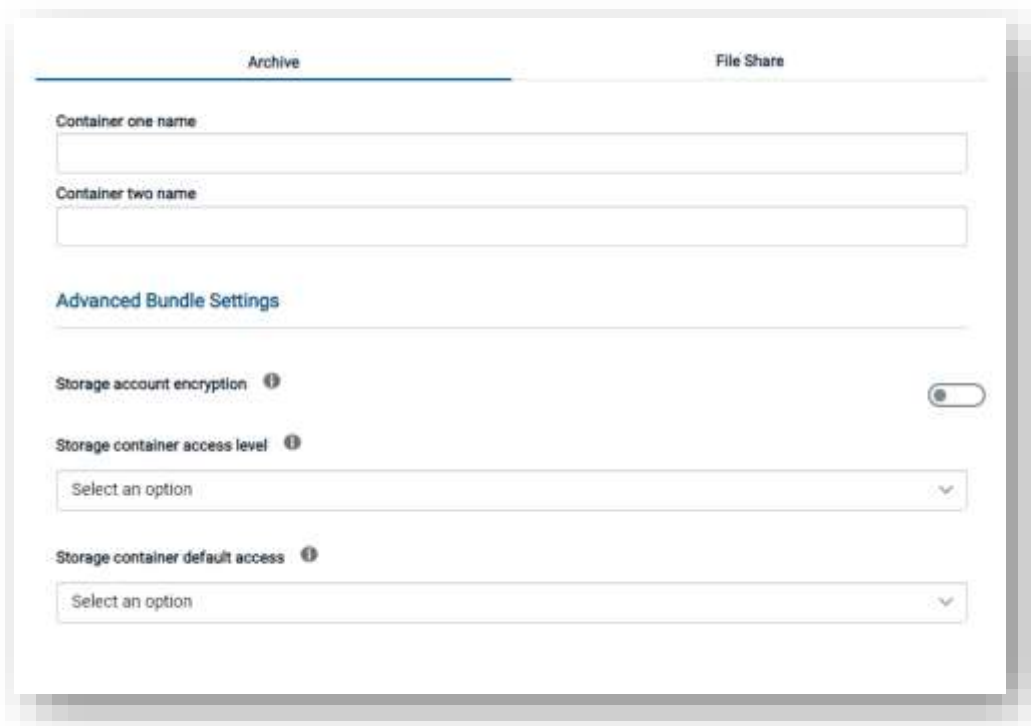
Access Tier: This sets the default tier of the Blob containers.

- **Hot** - Optimized for storing data that is accessed frequently.
- **Cold** - Optimized for storing data that is infrequently accessed and stored for at least 30 days.

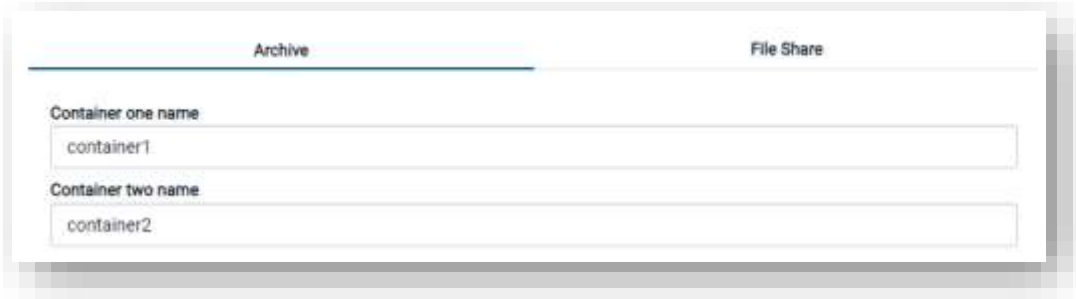
You can then select which options you want to enable. If you enable Azure File Share, the toggle for Azure File Sync will appear:



If you have enabled Archive, you can fill the Archive tab:



First choose names for the Containers that will be created:

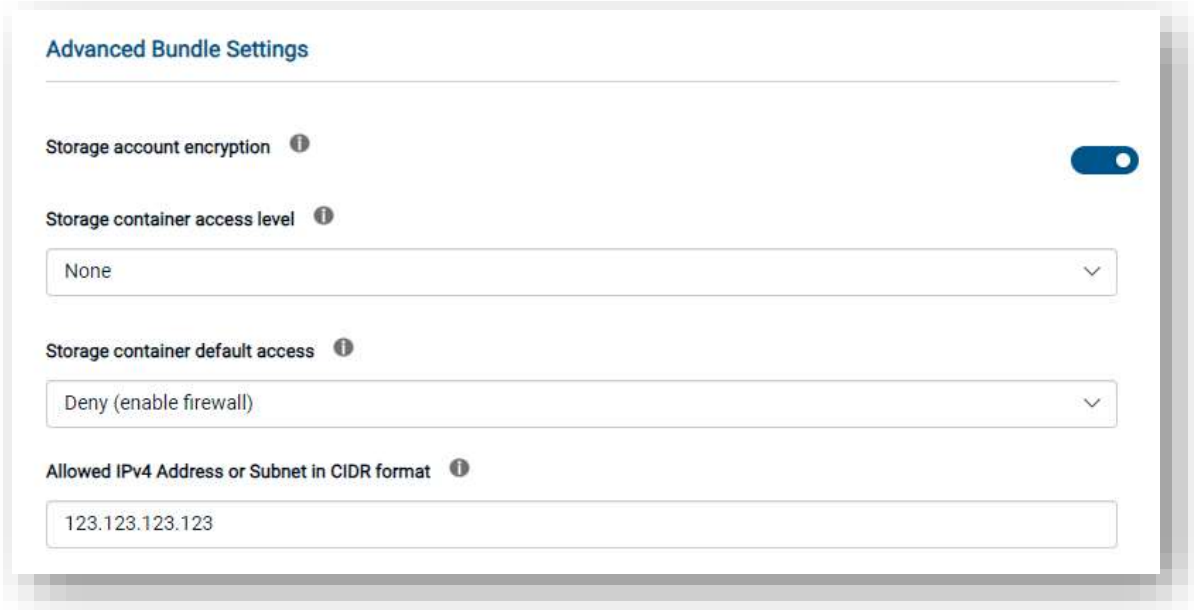


Archive File Share

Container one name  
container1

Container two name  
container2

Then you can select any Advanced Setting you would like:



Advanced Bundle Settings

Storage account encryption ⓘ

Storage container access level ⓘ  
None ▾

Storage container default access ⓘ  
Deny (enable firewall) ▾

Allowed IPv4 Address or Subnet in CIDR format ⓘ  
123.123.123.123

Enable Encryption: This sets encryption on the storage account.

- True – Enable Encryption
- False – Disable Encryption (Not recommended)

Container Access: By default, a container and any blobs within it may be accessed only by a user that has been given appropriate permissions. To grant anonymous users read access to a container and its blobs, you can set the container public access level. When you grant public access to a container, then anonymous users can read blobs within a publicly accessible container without authorizing the request.

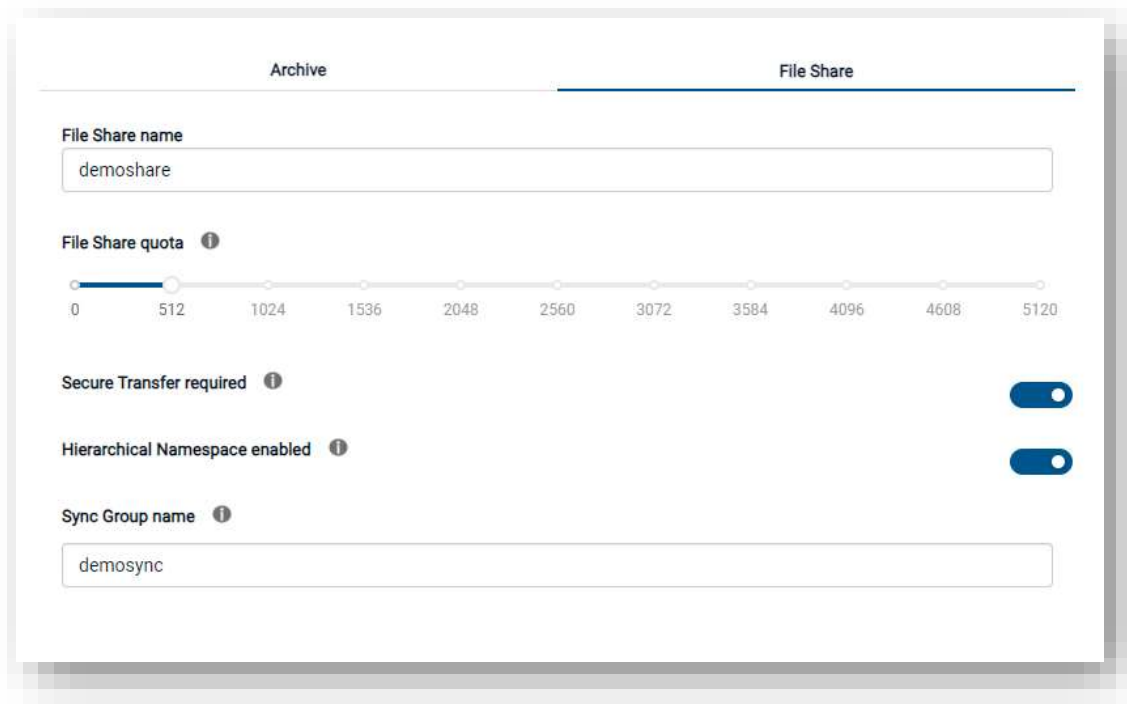
- None - The container and its blobs can be accessed only by the storage account owner. This is the default for all new containers.
- Blob - Blobs within the container can be read by anonymous request, but container data is not available. Anonymous clients cannot enumerate the blobs within the container.
- Container - All container and blob data can be read by anonymous request. Clients can enumerate blobs within the container by anonymous request but cannot enumerate containers within the storage account.

**Default Access:** If you select Deny, this will allow you to limit access to the Azure Storage Account to a public IP address/IP range that you specify. If you need to add additional IP addresses/ranges, you can do so in the storage account properties post-deployment.

- Deny – If Deny, specify an IP/IP range that you want to be able to access the storage account.
- Allow – If Allow, access to the storage account will not be limited to an IP address/IP range.

**Limit IP Access:** If setting Default Access to Deny, please specify your public IP address or IP range that you want to be able to access the blob's inside of the storage account. Examples: 97.96.157.22 or 97.96.157.0/24.

If you have enabled File Share, you can fill the File Share tab:



The screenshot shows the 'File Share' configuration tab. It includes the following fields and controls:

- File Share name:** A text input field containing 'demoshare'.
- File Share quota:** A slider control with a value of 512. The scale ranges from 0 to 5120 with increments of 512.
- Secure Transfer required:** A toggle switch that is currently turned on.
- Hierarchical Namespace enabled:** A toggle switch that is currently turned on.
- Sync Group name:** A text input field containing 'demosync'.

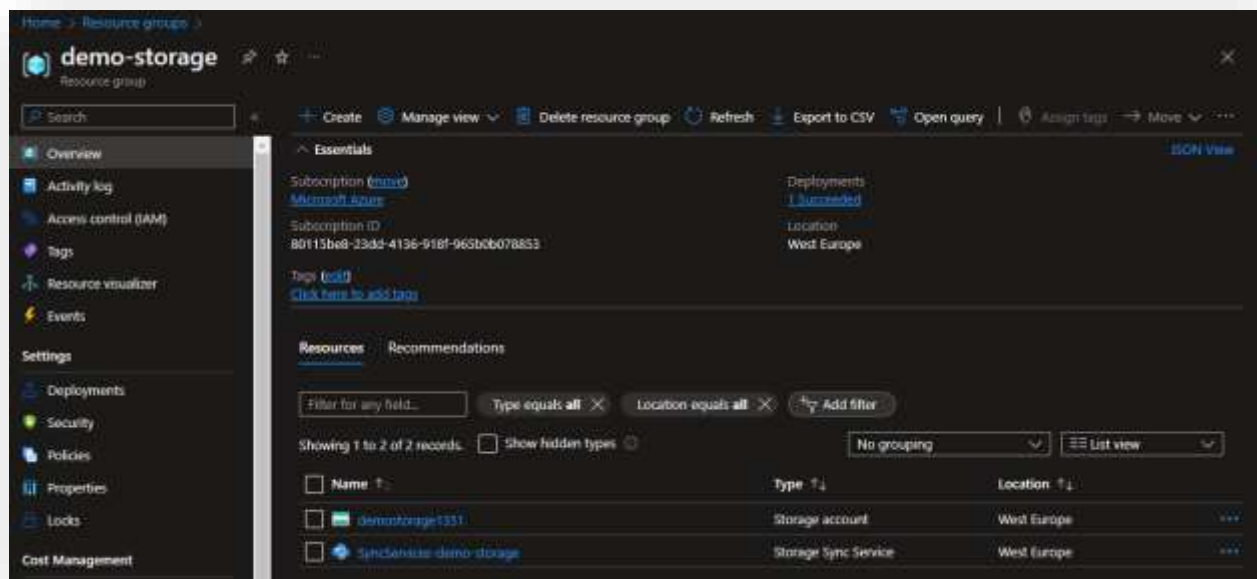
You can define a File Share name and select a File Share Quota up to 5 TB.

You should enable **Secure Transfer** if you only want to accept connections from secure protocols. If you are using the Azure File Sync agent this will utilize encryption in transit. If you needed to access the storage account directly using an insecure protocol (e.g., legacy application) then you have the option to turn this feature off. It is recommended that Secure Transfer be enabled unless you have a specific use.

Hierarchical Namespaces allow for the collection of object/files in a storage account in a hierarchy of directories (like a filesystem) improving the capability of providing the scalability and cost-effectiveness of object storage, with file system semantics that are familiar to analytics engines and frameworks. This setting also enables file level ACLs. It is recommended that Hierarchical Namespaces be enabled to take benefits.

If you have enabled File Sync, you can define the first Sync Group name to use.

You can then click on Deploy Now!



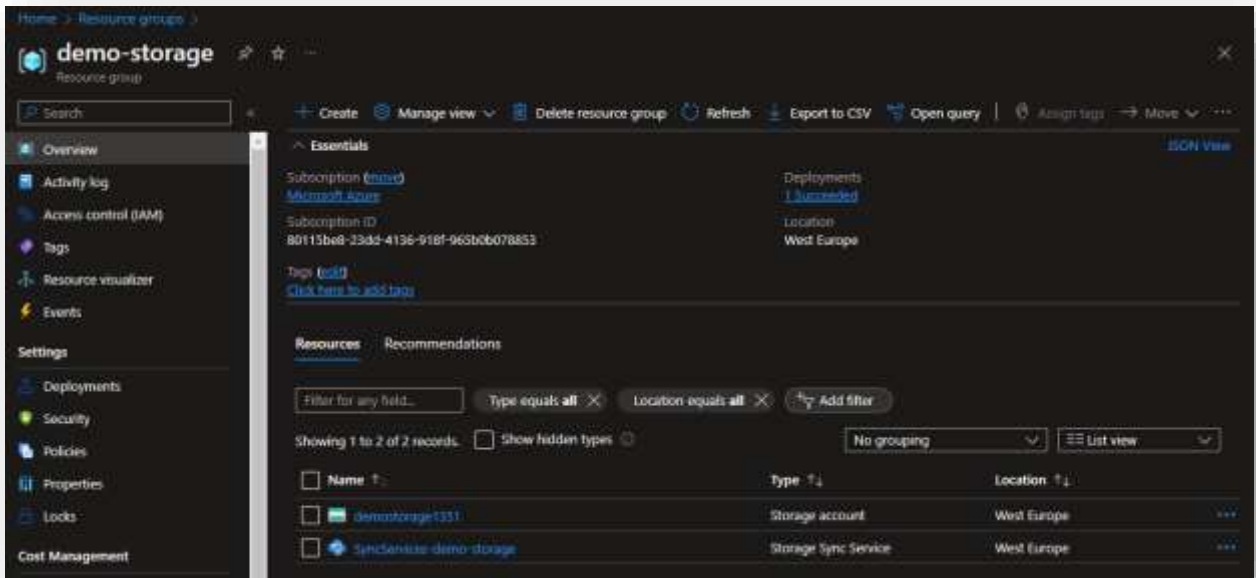
Once the deployment finished, you can start using your services!



# Post Deployment:

## Azure File Sync:


After the deployment is complete you should notice that a new resource group has been created along with the storage account and Storage Sync Service.



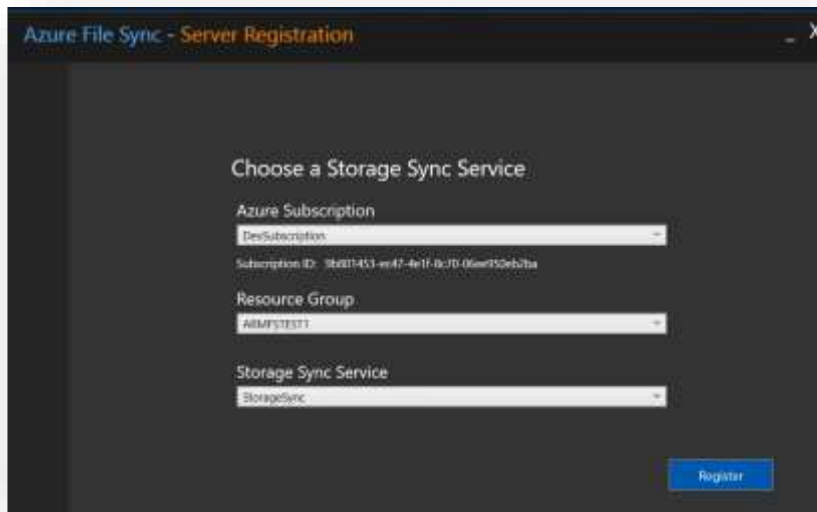
Visit <https://www.microsoft.com/en-us/download/details.aspx?id=57159> to download the appropriate version of the File Sync Agent and complete the wizard.



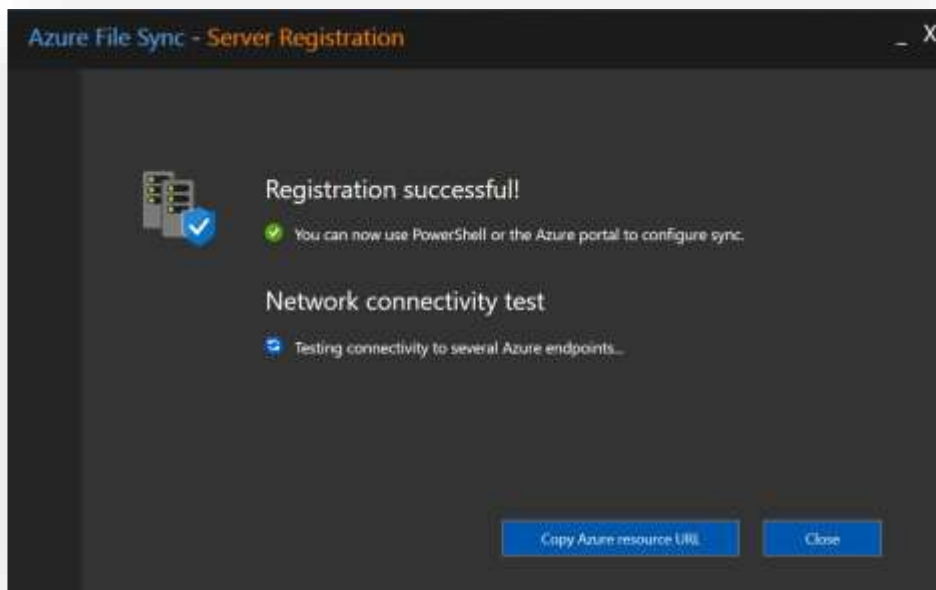
After installation, the setup for the File Sync Agent will start. You can sign into your azure subscription.



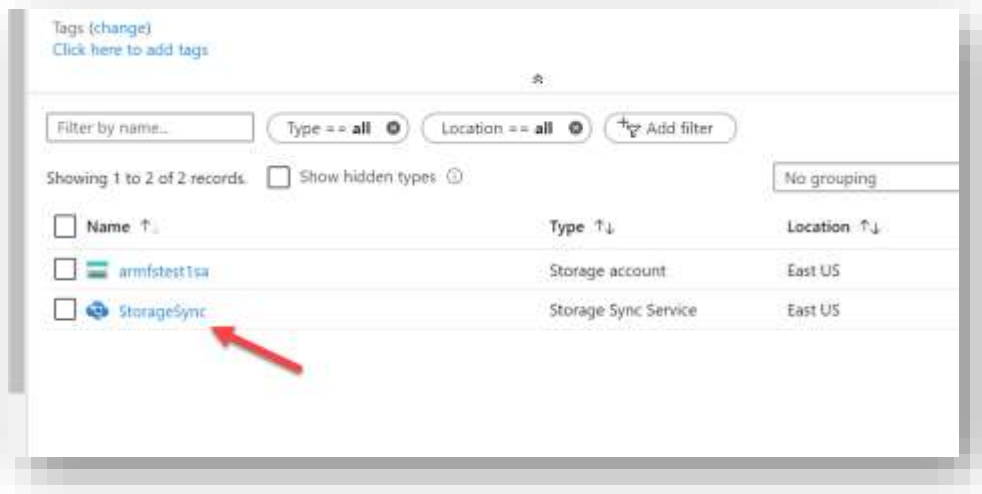
I then used the drop-down boxes to select the appropriate subscription, resource group, and then I selected Storage Sync for the Storage Sync Service. Then you can click on Register.



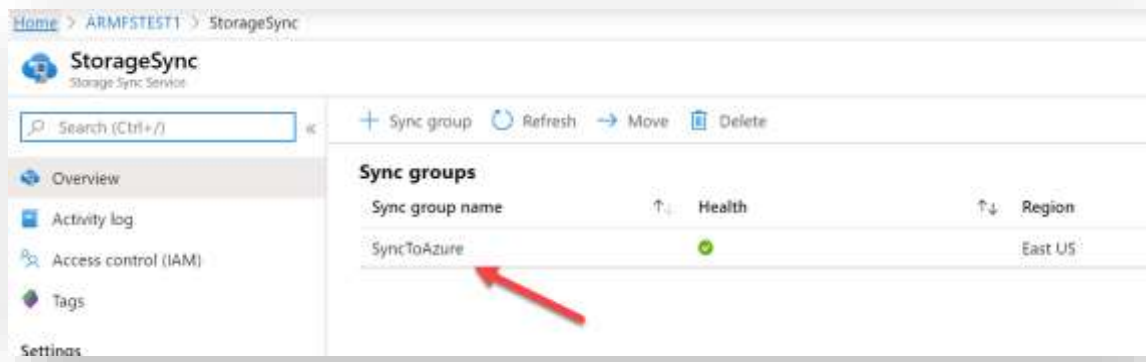
The next screen will let you know if the registration was successful.



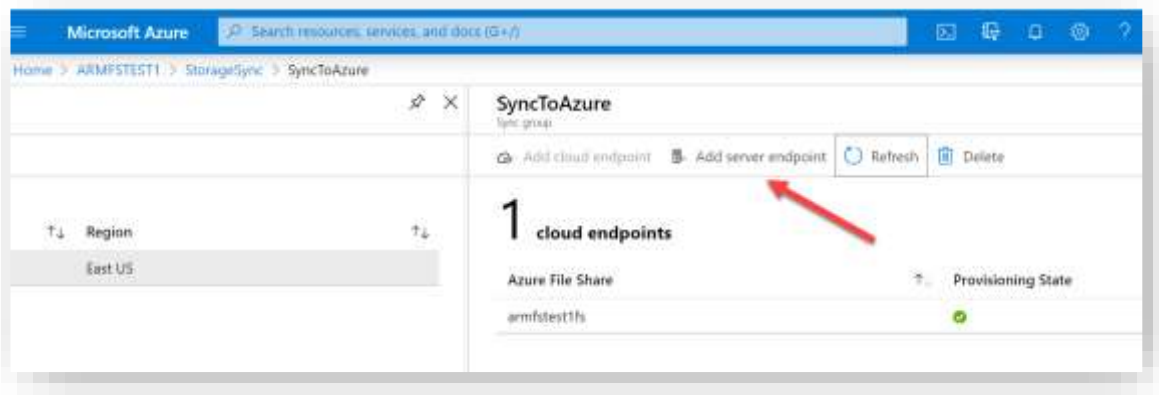
In the Azure Portal you can click on the Storage Sync Service.



Next click on the Sync Group that we want to configure. During the deployment in our case we created the Sync Group named “SyncToAzure”.



Next click on “Add server endpoint” to complete the linking of the Server Registration to the Sync Group.



You will be able to drop down the Registered Server to find the server that you completed the wizard on. You will also specify a path on the server that you would like to sync. For the purposes of this demo we will use the C:\Users directory. When you are done you can click on Create.

### Add server endpoint ✕

A server endpoint integrates an entire volume or a subfolder of a volume from a registered server as a location to sync. The following considerations apply:

- Servers must be registered to the storage sync service that contains this sync group before you can add a location on them here.
- A specific location on the server can only sync with one sync group. Syncing the same location or even a part of it – with a different sync group doesn't work.
- Make sure that the path you specify for this server is correct.


[Learn more](#)

Registered Server

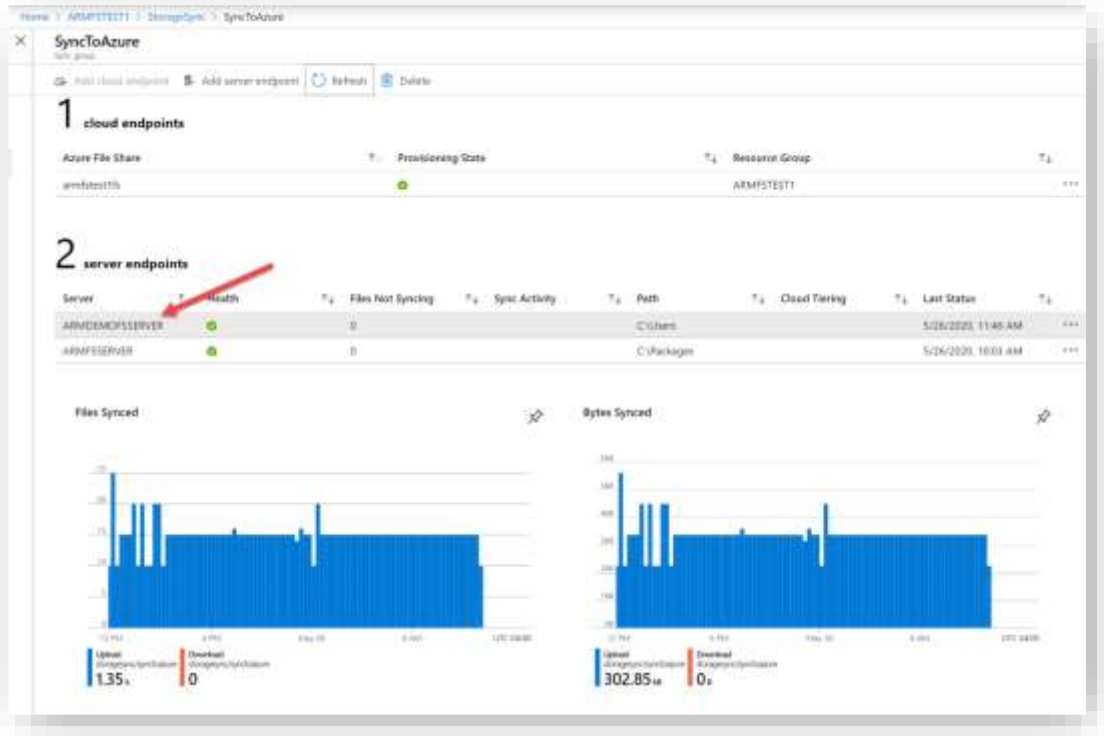
Path

Cloud Tiering

Offline Data Transfer

 You may be trying to create a server endpoint on the server's system volume. Please note that you will not be able to enable cloud tiering on the system volume.

You can monitor the sync status on the main Sync Group page. You want to wait for this first sync to complete before you install the Azure File Sync Agent on other servers.



If you desire, you can also browse into the storage account to find the files as well.

