

# Windows Server 2016/2019 Hybrid scenario: File Sharing

---

White paper and guide

# Table of Contents

---

<b>1: Hybrid file sharing scenario</b>	<b>3</b>
Background	3
Extending on-premises file solutions to the Microsoft Azure cloud	3
When to use a hybrid solution	3
Challenges	4
Network file sharing	4
<b>2: Scenario setup</b>	<b>5</b>
Microsoft Azure signup	5
Free trial subscription	5
Windows Server hybrid file sharing architecture	7
Windows Server host	7
Software required	8
Scenario host configuration	8
Windows Server virtual machine creation	9
Windows Server domain controller	9
Windows Server file server	9
File share setup	11
Azure File Sync solution	12
Azure File Sync terminology	12
Azure File Sync setup	14
Backup and recovery – on-premises	21
Backup and recovery (files) – Azure	25
<b>3: Appendix</b>	<b>35</b>
Windows 10 client running Windows Admin Center (WAC)	35

# 1: Hybrid file sharing scenario

---

## Background

In recent years, small and medium businesses (SMBs) have taken an increasing interest in modernizing their datacenters. The immense growth of stored data, the need for greater data security, the cost of datacenter leases, hardware updates, and services, and the pressures of regulatory compliance are top-of-mind issues for key IT and business decision makers. They are taking a deeper look at their current on-premises datacenter deployments. Despite some initial reticence, SMBs are discovering the benefits of the cloud and hybrid solutions.

In a [January 2018 Microsoft online survey](#) of more than 1,700 non-Microsoft IT professionals from several countries, 67% of the respondents are using or were planning to deploy a hybrid cloud. Control of where data is stored is cited as the top use case by 71% of the respondents, with 69% noting cloud-based back-up and disaster recovery as a close second.

For SMBs, specifically, Microsoft Windows Server 2016 is an acknowledged leader among on-premises infrastructure solutions. It was designed to enable a hybrid deployment of workloads and data to the cloud for two scenarios: file sharing and database. [The new Microsoft Windows Server 2019 extends the capabilities of Windows Server 2016 even further.](#)

In terms of hybrid capabilities with Azure over the 2016 edition, Windows Server 2019 has added storage migration service, predictive analytics with System Insights, and Azure network adaptor to easily connect Azure virtual networks. Both Windows Server 2016 and 2019 offer the ability to synchronize file servers to Azure and virtual machine (VM) protection that replicates workloads running on physical and virtual machines from a primary site to a secondary location—which are important to a discussion about hybrid cloud deployment. We will refer to these capabilities in the course of the discussion.

For this white paper, the file sharing scenario capabilities of Windows Server 2016/2019 will be discussed. The cloud platform used may be from any provider, or in a multi-cloud deployment; but the cloud platform of preference here is Microsoft Azure.

## Extending on-premises file solutions to the Microsoft Azure cloud

When businesses move workloads and data to the cloud, their on-premises datacenters often continue to play an important role. The term *hybrid cloud* normally refers to a combination of public cloud and private or community cloud solutions. In this series of scenarios, we are considering the connection of on-premises datacenters to specific Azure services to bring value and security to the SMB workloads. This also helps to create an integrated IT environment that spans both. Some organizations use hybrid cloud as a path to migrate their entire datacenter to the cloud over time. Other organizations use cloud services to extend their existing on-premises infrastructure.

This paper describes some considerations and best practices for managing file sharing in a hybrid cloud solution deploying, centralizing, and synchronizing file shares with Azure File Sync.

## When to use a hybrid solution

**There are many triggers to considering a hybrid solution, including:**

- As a transition strategy during a longer-term migration to a fully cloud native solution.
- When regulations or policies do not permit moving specific data or workloads to the cloud.
- For disaster recovery and fault tolerance, by replicating data and services between on-premises and cloud environments.

- To reduce latency between an on-premises datacenter and remote locations, by hosting part of the architecture in Azure.

## Challenges

- Creating a consistent environment to meet complex needs in terms of security, management, and development, and in avoiding the duplication of work.
- Creating a reliable, low-latency, and secure data connection between on-premises and cloud environments.
- Replicating data and modifying applications and tools to use the correct data stores within each environment.
- Securing and encrypting data that is hosted in the cloud but accessed from on-premises, or vice versa.

## Network file sharing

In a hybrid cloud architecture, it is common for an organization to keep newer files on-premises while archiving older files to the cloud. This is sometimes called file tiering, and, in Azure this is called *cloud tiering*. There is seamless access to both sets of files, on-premises and cloud-hosted. This approach helps to minimize network bandwidth usage and access times for newer files, which are likely to be accessed the most often. At the same time, there are many benefits of cloud-based storage for archived data.

In Azure, an on-premises tier is accomplished using [Azure Files](#). Azure Files offers fully managed file shares that can be accessed with the standard [Server Message Block](#) (SMB) protocol, a more robust version of Common Internet File System, or CIFS, protocol, mainly used in network-attached storage (NAS) systems. Azure Files can be mounted as a file share on a local computer or used together with existing applications that access local or network share files.

This white paper provides a walk-through of the hardware, software, and services required to implement hybrid file sharing using Azure File Sync to synchronize file shares in Azure Files with on-premises Windows Servers. Azure File Sync also provides cloud tiering between an on-premises file server and Azure Files. This lets users keep only the newest and most recently accessed files locally.

### File sharing scenario pre-requisites

This file sharing hybrid scenario has a number of pre-requisites.

1. A current Microsoft Azure subscription
2. A physical Windows Server 2016 host machine
3. Two Windows Server 2016 Virtual Machines:
  - A domain controller
  - A file server
4. A Windows 10 (1803 version or later) virtual machine client

There are a number of choices for the Azure subscription. If the user already has a current subscription, Microsoft Azure Signup section can be safely skipped. A free trial subscription is also available and provides a credit and one month of usage. This subscription can be upgraded to a pay-as-you-go (PAYG) subscription at the end of the trial, if required.



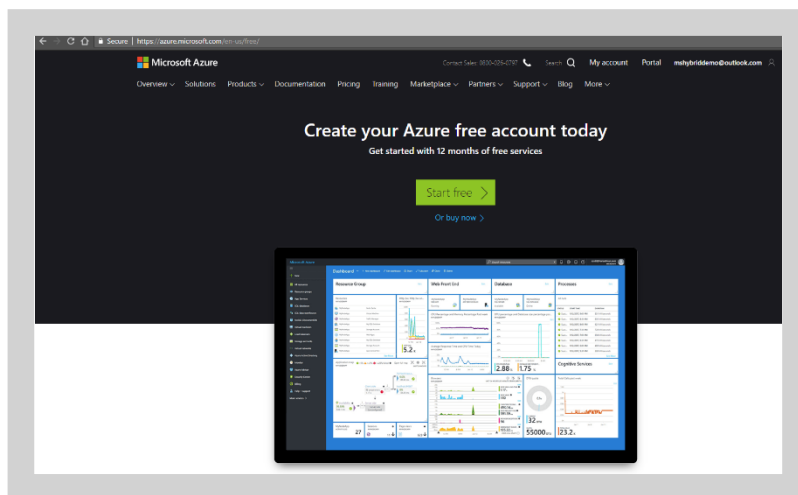
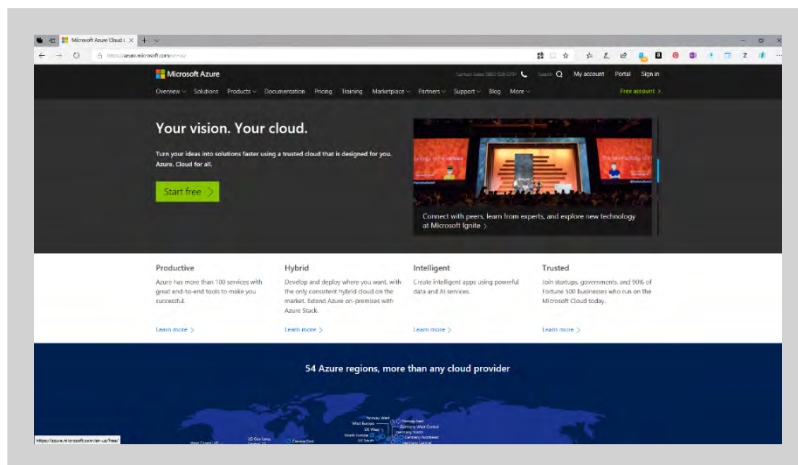
## 2: Scenario setup

### Microsoft Azure signup

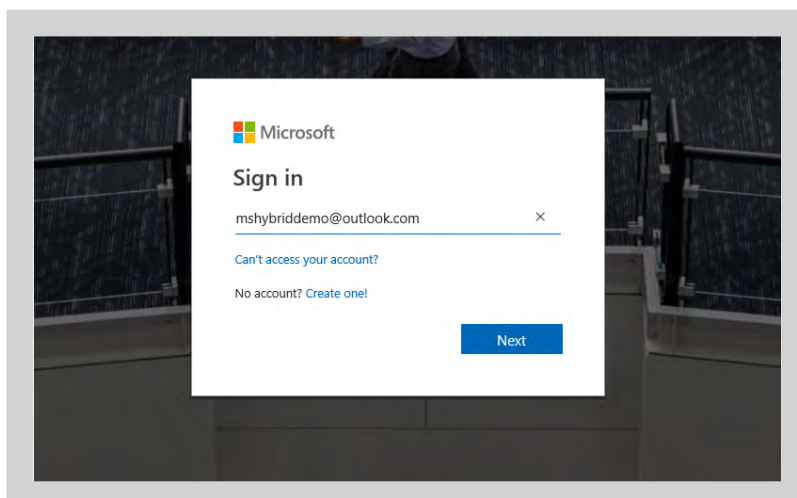
The first step in adopting the hybrid on-premises cloud capabilities of Windows Server 2016/2019 is to subscribe to Microsoft Azure. Microsoft offers a free trial subscription that provides an opportunity to experience the benefits of the Azure cloud for 12 months of popular services, a credit for US \$200, or the equivalent in foreign currency, to explore any Azure service for 30 days, and access to more than 25 free services.

### Free trial subscription

Use any current version of Microsoft Edge or Internet Explorer (IE) web browser (or any other browser) and navigate to <https://azure.com>, and click *Start Free*. You will need to click *Start Free* again on the *Create your Azure free account today* page.



Sign in with an existing Microsoft account or create a new one. (For example, we have used [mshybriddemo@outlook.com](mailto:mshybriddemo@outlook.com) as the identity for this series of hybrid scenarios.) (Note that a Microsoft account can have only one free trial of Microsoft Azure.) If you need to sign up for a new Microsoft account, navigate to <https://outlook.com> and choose to create a free account.



Sign in and complete the free trial account using your own business details. Once completed, the Azure Portal will be available at <https://portal.azure.com>. This will start your 12-month free trial subscription, with the 30-day free Azure service and the US \$200 (or equivalent) credit.

Once the free month has elapsed, you will still be eligible for more than 25 free services for one year.

**Here is a listing of some of those free services:**

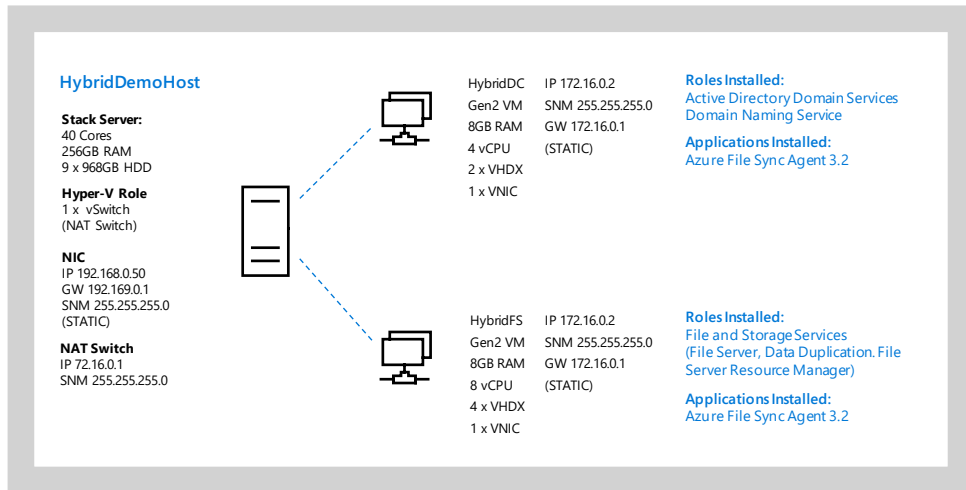
Free services for 12 months ⓘ				
Usage period: 8/20/2018 - 9/19/2018				
METER	USAGE/LIMIT		STATUS	
Storage (GB) - Cosmos DB	0%	0 / 5 GB	Not in use	
100 Request Units (Hours) - Cosmos DB	0%	0 / 2,976 Hours	Not in use	
Standard IO - Files (GB) - Locally Redun...	0%	0 / 5 GB	Not in use	
Premium Managed Disks/P6 (Units) - Lo...	0%	0 / 2 Units	Not in use	
Standard Managed Disk/Snapshots (GB)...	0%	0 / 1 GB	Not in use	
Standard Managed Disk Operations (in ...	0%	0 / 200 10K Operations	Not in use	
Standard IO - Hot Block Blob (GB) - Loca...	0%	0 / 5 GB	Not in use	
B1s - B5 Series	0%	0 / 750 1 Hour	Not in use	
B1s - B5 Series Windows	0%	0 / 750 1 Hour	Not in use	
Data Transfer Out (GB)	0%	0 / 15 GB	Not in use	
Standard S0 Database Days - SQL Datab...	0%	0 / 31 Days	Not in use	
Standard IO - File Read Operation Units ...	0%	0 / 4 10,000s	Not in use	
Standard IO - File Write Operation Units...	0%	0 / 1 10,000s	Not in use	
Standard IO - File Protocol Operation U...	0%	0 / 4 10,000s	Not in use	
Standard IO - File List Operation Units (i...	0%	0 / 1 10,000s	Not in use	
Standard IO - Hot Block Blob Read Oper...	0%	0 / 2 10,000s	Not in use	
Standard IO - Hot Block Blob Write Ope...	0%	0 / 1 10,000s	Not in use	
IP Address Hours - Public IP Addresses	0%	0 / 1,500 Hours	Not in use	
Compute Hours - Red Hat Enterprise Lin...	0%	0 / 750 Hours	Not in use	

The 30-day free trial and the 12-month free service credit are sufficient to successfully deploy and evaluate all the hybrid scenarios in this series.

You should retain your account and Azure subscription ID details. Without them, you will not be able to access your Azure subscription.

## Windows Server hybrid file sharing architecture

A general discussion of setting up the host server and virtual machines follows according to the architecture below.



## Windows Server host

Windows Server 2016 (as well as Windows Server 2019) introduces new features for tenants and reduces the total cost of ownership. The hosted server makes the IaaS solution more reliable and robust, and improves infrastructure security, even for service providers using other solutions as the IaaS platform.

The scenario in this white paper was created using a twin Xeon CPU server (20 cores) with 256GB of RAM and 9 HDDs all 1 TB in size.

In addition to the installation of the Host OS, virtual machines will need to be set up.

Once Windows Server 2016/2019 is installed with the default settings, a single [storage pool](#) will be created from the remaining disks. This is followed by creating a 5TB volume for this scenario that is labelled D: (the C: system OS disk is 250GB). Both Disks C: and D: should be formatted with NTFS. The configuration entails a Microsoft Hyper-V role that uses either [Server Manager](#) or [PowerShell](#) and the creation of an internal switch using [Hyper-V Manager](#) or [PowerShell](#). Internet access should be made available for the VMs in your trial scenario. There are a number of ways to do this, including using the [NetNat](#) feature. Next, create an internal virtual switch using [Hyper-V Manager](#) or [PowerShell](#).

There are a number of ways of providing internet access for your scenario virtual machines. This scenario document uses the newly created NetNat feature in the Windows Server.

Using this feature, a NetNat object is created inside the new virtual switch that provides network address translation out of the virtual network to the internet (provided the host has internet access).

To [create a NetNat](#) object, open a PowerShell session as an administrator and run Get-NetAdapter to find the interface index of a new Virtual Switch, a logical switching fabric. This internal switch will need to be renamed like *vEthernet* (as a SwitchName) and an Interface Description of Hyper-V Virtual Ethernet Adapter needs to be created. Take note of its *ifIndex* to use in the next step.

Create a NAT gateway using **New-NetIPAddress** in PowerShell:

```
New-NetIPAddress -IPAddress <NAT Gateway IP> -PrefixLength <NAT Subnet Prefix Length>
-InterfaceIndex <ifIndex>
```

You'll need some network information:

- **IPAddress:** NAT Gateway IP specifies the IPv4 or IPv6 address to use as the NAT gateway IP. A common gateway IP is 172.16.0.1
- **PrefixLength:** NAT Subnet Prefix Length defines the NAT local subnet size (subnet mask). A common PrefixLength is 24 -- this is a subnet mask of 255.255.255.0
- **InterfaceIndex:** ifIndex is the interface index of the virtual switch, which you determined in the previous step.

Based on the above information, you could run the following to create the NAT Gateway:

```
New-NetIPAddress -IPAddress 172.16.0.1 -PrefixLength 24 -InterfaceIndex 24
```

The final step is to configure the NAT network using [New-NetNat](#).

```
New-NetNat -Name <NATOutsideName> -InternalIPInterfaceAddressPrefix <NAT subnet prefix>
```

You will need the network info and NAT gateway from previous steps:

- **Name:** NATOutsideName describes the name of the NAT network. You'll use this to remove the NAT network.
- **InternalIPInterfaceAddressPrefix:** NAT subnet prefix describes both the NAT Gateway IP prefix from above as well as the NAT Subnet Prefix Length from above.

Run the following to setup the NAT network for the scenario described above:

```
New-NetNat -Name HybridDemoNat -InternalIPInterfaceAddressPrefix 172.16.0.0/24
```

## Software required

If you are building this scenario as a proof of concept and using evaluation media, you will need to download the ISO files from the Microsoft Evaluation Center [here](#).

You will need a [Windows Server ISO](#) and a [Windows 10 ISO](#) (for the Windows Admin Center client).

At the end of host setup, you will need to configure the scenario host.

## Scenario host configuration

The host in this example has the following configuration:

Name - HybridDemoHost  
 CPU 2 x 20 Core, 40 logical processors  
 RAM - 256GB RAM  
 HDD - 8 x 900GB HDD  
 Windows Server 2016 Datacenter Edition  
 Hyper-V Role  
 Physical NIC  
 IP 192.168.0.50 (STATIC)  
 Gateway 192.168.0.1  
 Subnet Mask 255.255.255.0



NAT Switch (vSwitch)  
Gateway IP address 172.16.0.1  
Subnet Mask 255.255.255.0

## Windows Server virtual machine creation

After the host has been built, a Small Business environment should be built. In this case, a single domain controller and a file server; in production you should have at least two domain controllers. The following steps are needed to create these machines.

## Windows Server domain controller

The example SMB company in this scenario has a single-domain Active Directory Domain Services forest, **hybriddemo.cloud**. Because this domain will be extended into Microsoft Azure Active Directory in other scenarios, it is not necessary in the file sharing scenario.

Due to the small number of devices in the scenarios, static IP addressing is used rather than DHCP.

### The VM that runs the domain controller has the following configuration:

Name HybridDC  
Version - Gen2 VM  
RAM - 8GB RAM Static  
Processors – 4 vCPU  
1 x OS VHDX (127GB Dynamically Expanding Disk)  
1 x Data VHDX (127GB Dynamically Expanding Disk)  
(For production, consider fixed disks)  
  
1 x vNIC connected to the NAT Switch  
Network Configuration  
IP 172.16.0.2 (STATIC)  
SNM 255.255.255.0  
GW 172.16.0.1

Install the Windows Server 2016/2019 OS on the VM and add the Active Directory Domain Services role.

You can do this using [Server Manager](#) or [PowerShell](#).

The next step is to promote the **HybridDC** server to a domain controller as the first domain controller in a new forest.

### The example uses the following settings:

- Domain Name - **Hybriddemo.cloud**
- Volume for the sysvol files - D: (best practice for AD files, also NTFS formatted)

Having promoted the server to a domain controller, restart and log in as an administrative user. (The best practice is to disable the default administrator user and create a new admin user. In this scenario, **Labadmin** is used.)

The next stage is to create a second VM for the file server and join it to the **hybriddemo.cloud** domain.

## Windows Server file server

The file server will host a storage pool with a storage space and virtual disks to manage the file shares.

Install Windows Server 2016/2019 Standard to the VM configured as below.

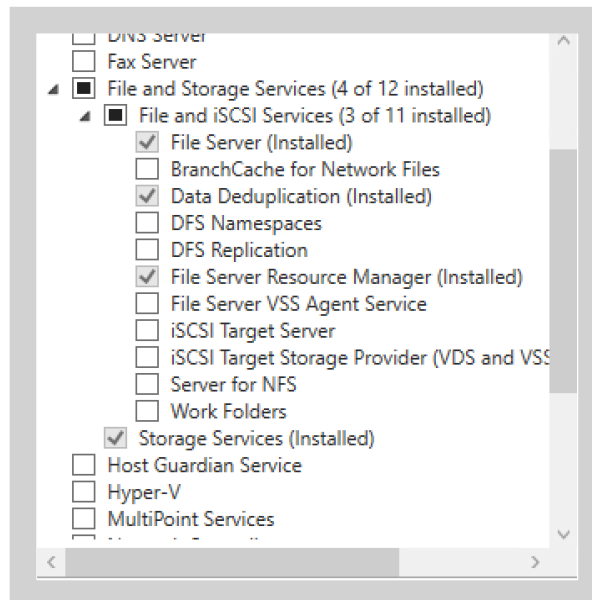
The VM to run the File server has the following configuration:

Name - HybridFS  
Version - Gen2 VM  
RAM - 8GB RAM Static  
Processors - 8 vCPU  
1 x OS VHDX (127GB Dynamically Expanding Disk)  
3 x Data VHDX (127GB Dynamically Expanding Disk)  
(For production consider fixed disks)  
  
1 x VNIC  
IP address 172.16.0.3 (STATIC)  
Subnet Mask 255.255.255.0  
Gateway 172.16.0.1

Having installed the operating system, join the file server to the **hybriddemo.cloud** domain, either by using the [GUI](#) or [PowerShell](#), you will need credentials for the administrative user you created (or another user with permission to create computer objects in the domain).

Once the computer has restarted, log in with the domain admin credentials. (This is not a best practice, but useful for this scenario.)

The next step is to add the necessary roles and services using either Server Manager or PowerShell for the roles required shown below:



The file server resource manager is not required, but it is very useful on any corporate file server to create file screens and quotas. It does not affect the hybrid elements of this scenario.


In addition, install the Windows Server Backup (WSB) feature.

Once completed, we are ready to start the hybrid file sharing scenario.

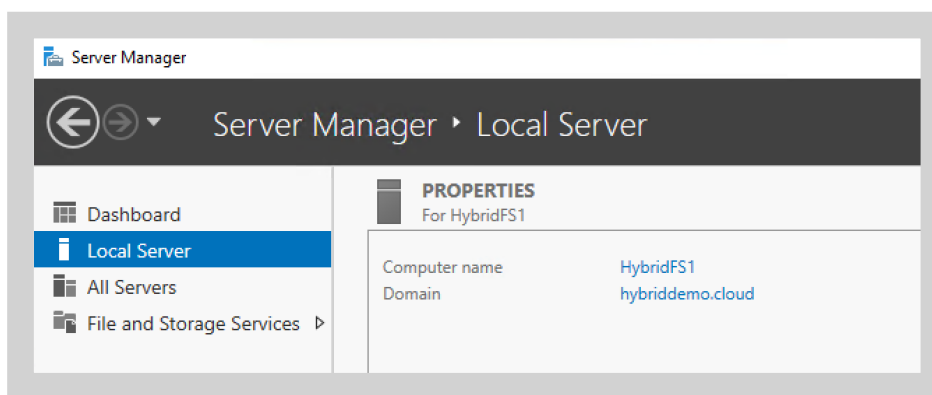
## File share setup

There are several ways to create a file share in Windows Server. To simplify the process, this document will cover the GUI version using Server Manager.

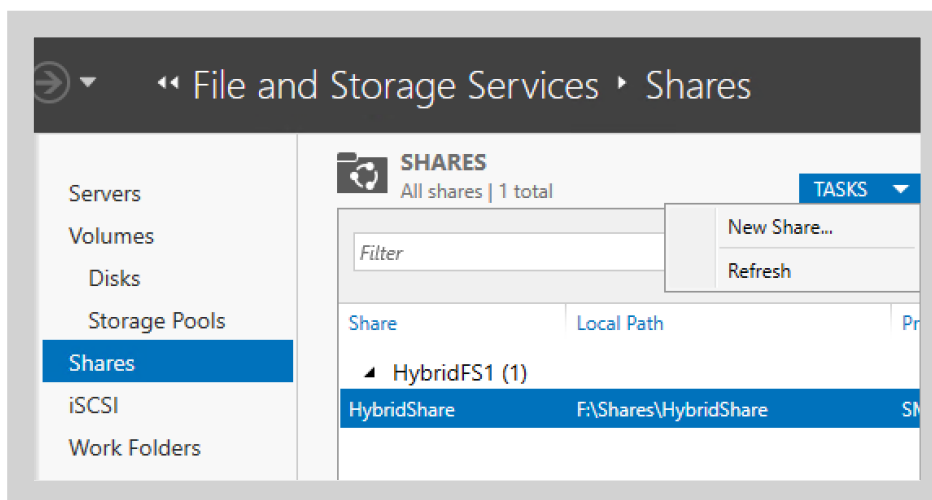
First, create a folder named **HybridShare** in the F: drive under the Shares folder:



Select **File and Storage Services**:



Select **Tasks and New Share**:



Complete the wizard with the following settings:

Confirm that the following are the correct settings, and then click Create.

SHARE LOCATION	
Server:	HybridFS1
Cluster role:	Not Clustered
Local path:	F:\shares\hybridshare
SHARE PROPERTIES	
Share name:	hybridshares
Protocol:	SMB
Access-based enumeration:	Disabled
Caching:	Enabled
BranchCache:	Disabled
Encrypt data:	Disabled

Complete the wizard.

If you have some files you want to use, transfer them now to this HybridShare folder.

If you need to create some random files, then you can use the following **FSUTIL** command:

```
fsutil file createnew test.txt 52428800
```

That will create a 52MB file called *test.txt*

This completes the on-premises setup. In production, you would need to set file permissions for different users and groups. This is not required in this scenario.

The next step is to create an Azure File Sync solution to connect to the file share we have just created.

## Azure File Sync solution

Azure Files deployment is one of the core services within Azure Storage. Azure Files provides an SMB2.1 and 3.0 file share service to on-premises and cloud-based clients. A recent extension to this is Azure File Sync. Below are some details and a glossary of the terms we will use in setting up the file sync solution for our *hybridfs1* server.

Azure File Sync turns Windows Server into a fast cache of Azure file share. You can use SMB, NFS, and FTPS to access your data locally. You can have as many caches as you need across the world.

## Azure File Sync terminology

### Storage Sync Service

The Storage Sync Service is the Azure resource for Azure File Sync. It is a peer of the storage account resource and can be deployed to Azure resource groups. A subscription can have multiple Storage Sync Service resources deployed.



### Sync group

A sync group defines the sync topology for a set of files. Endpoints within a sync group stay synchronized with each other. If, for example, you have two distinct sets of files that you want to manage with Azure File Sync, you would create two sync groups and add different endpoints to each sync group. A Storage Sync Service can host as many sync groups as you need.

### Registered server

The registered server object represents a trust relationship between your server and the Storage Sync Service. You can register as many servers to a Storage Sync Service instance as you want. A server can be registered with only one Storage Sync Service at a time. However, it is recommended to [deploy as many storage sync services as needed](#) to separate groups of servers.

### 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. The Azure File Sync agent has three main components:

- **FileSyncSvc.exe:** The background Windows service that is responsible for monitoring changes on server endpoints, and for initiating sync sessions to Azure.
- **StorageSync.sys:** The Azure File Sync file system filter that is responsible for tiering files to Azure Files (when cloud tiering is enabled).
- **PowerShell management cmdlets:** PowerShell cmdlets that you use to interact with the Microsoft StorageSync Azure resource provider (Windows Server 2016).

### Server endpoint

A server endpoint represents a specific location on a registered server, such as a folder on a server volume. Multiple server endpoints can exist on the same volume if their namespaces do not overlap (for example, F:\share1 and F:\share2). Cloud tiering policies can be created individually for each server endpoint.

If you add a server location that has an existing set of files as a server endpoint to a sync group, those files are merged with any other files that are already on other endpoints in the sync group.

### Cloud endpoint

A cloud endpoint is an Azure file share that is part of a sync group. The entire Azure file share syncs, and an Azure file share can be a member of only one cloud endpoint. If you add an Azure file share that has an existing set of files as a cloud endpoint to a sync group, the existing files are merged with any other files that are already on other endpoints in the sync group.

### Cloud tiering

Cloud tiering is an optional feature of Azure File Sync in which infrequently used or accessed files greater than 64 KiB in size can be tiered to Azure Files. When a file is tiered, the Azure File Sync file system filter *StorageSync.sys* replaces the file locally with a pointer, or reparse point. The reparse point represents a URL to the file in Azure Files. A tiered file has the “offline” attribute set in NTFS so third-party applications can identify tiered files. When a user opens a tiered file, Azure File Sync seamlessly recalls the file data from Azure Files without the user needing to know that the file is not stored locally on the system. This functionality is also known as Hierarchical Storage Management (HSM).

More information on Azure File Sync can be found [here](#).

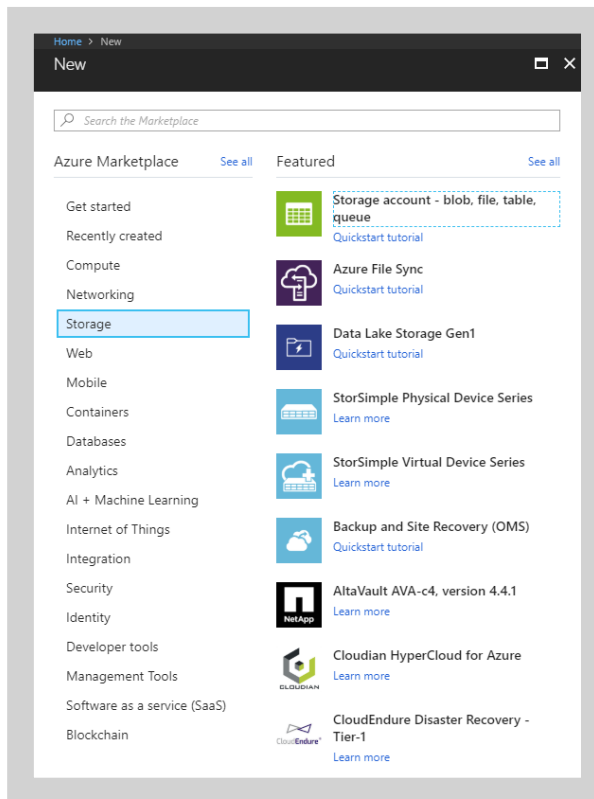
The next section describes the steps necessary to configure the Azure File Sync agent.

## Azure File Sync setup

First, create a storage account in Azure.

Log into your Azure subscription at <https://portal.azure.com>.

Click **Create a resource** and choose **Storage**, as below:



Home > New > Create storage account

## Create storage account

The cost of your storage account depends on the usage and the options you choose below. [Learn more](#)

\* Name <sup>i</sup>  
hybriddemosa ✓  
core.windows.net

Deployment model <sup>i</sup>  
**Resource manager** Classic

Account kind <sup>i</sup>  
StorageV2 (general purpose v2) ▾

\* Location  
East US ▾

Replication <sup>i</sup>  
Locally-redundant storage (LRS) ▾

Performance <sup>i</sup>  
**Standard** Premium

Access tier (default) <sup>i</sup>  
Cool **Hot**

\* Secure transfer required <sup>i</sup>  
**Disabled** Enabled

\* Subscription  
Free Trial ▾

\* Resource group  
☒ Create new ☐ Use existing  
hybridfsRG ✓

Virtual networks  
Configure virtual networks <sup>i</sup>  
Disabled Enabled

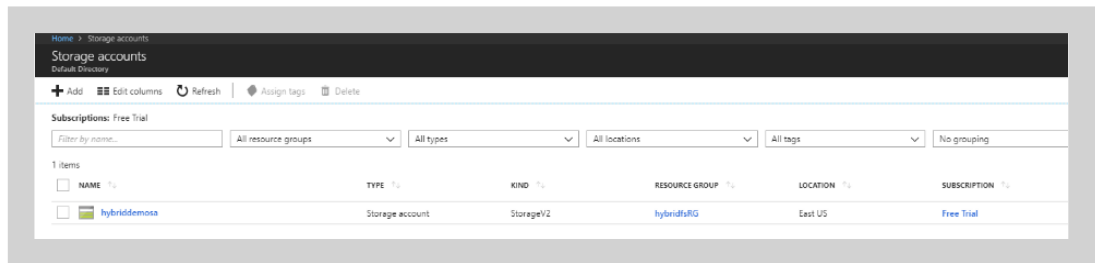
Data Lake Storage Gen2 (preview)  
Hierarchical namespace <sup>i</sup>  
Disabled Enabled

**Create** Automation options

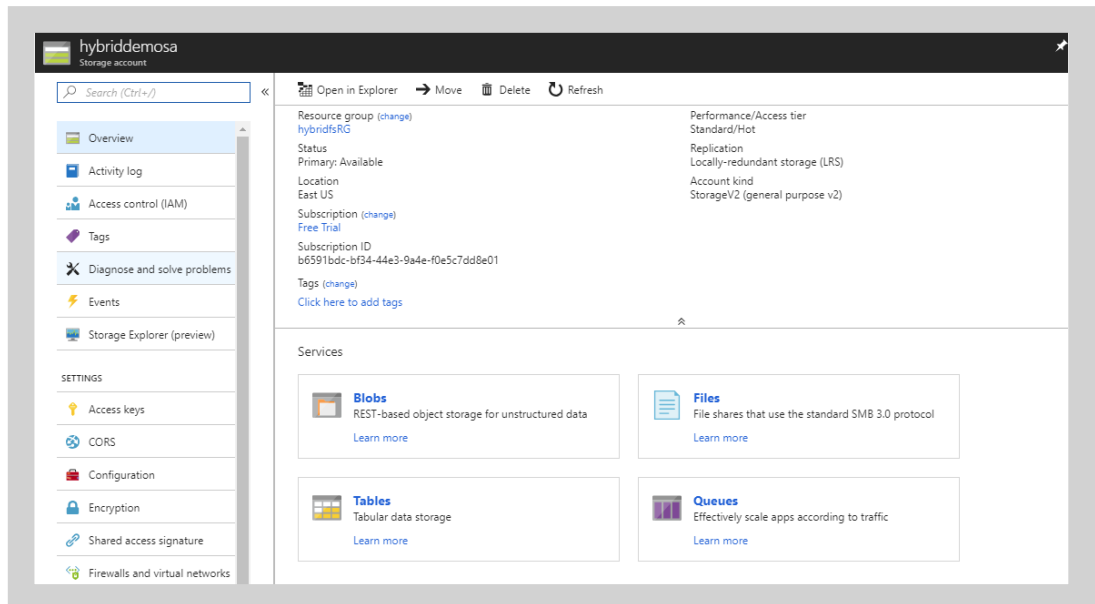
Click *Create storage account* and enter a storage account name (must start with a lowercase letter and only contain lowercase letters and numbers). Be sure to select the same options as the graphic below.

Choose a locally redundant storage and resource manager account, and enter a region close to your location and a memorable resource group name. (Note the location of the Azure File Sync agent should be in the same location as the storage account).

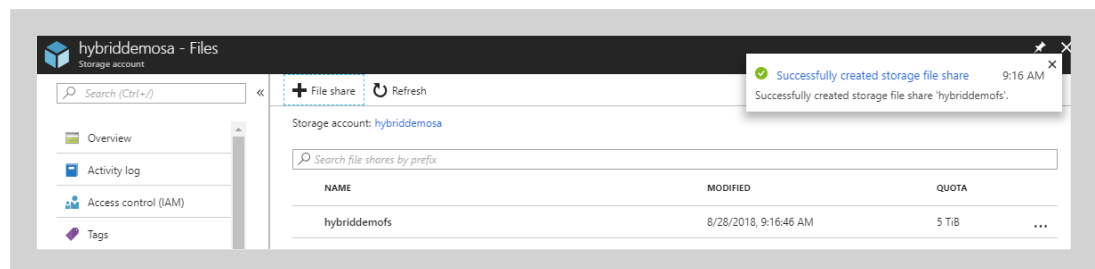
Complete the deployment of the storage account as shown:



When complete, click on the storage account and create an Azure file service:

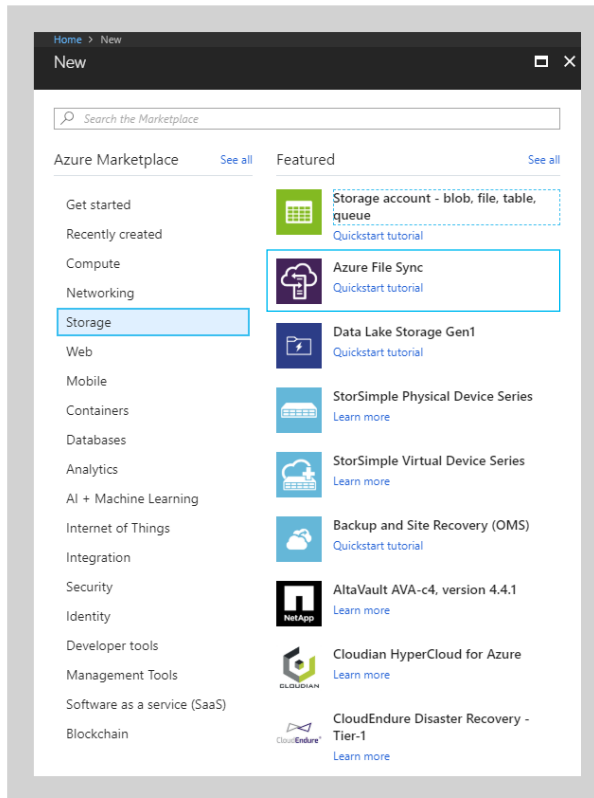


When you select *Files*, go ahead and create a new file share:

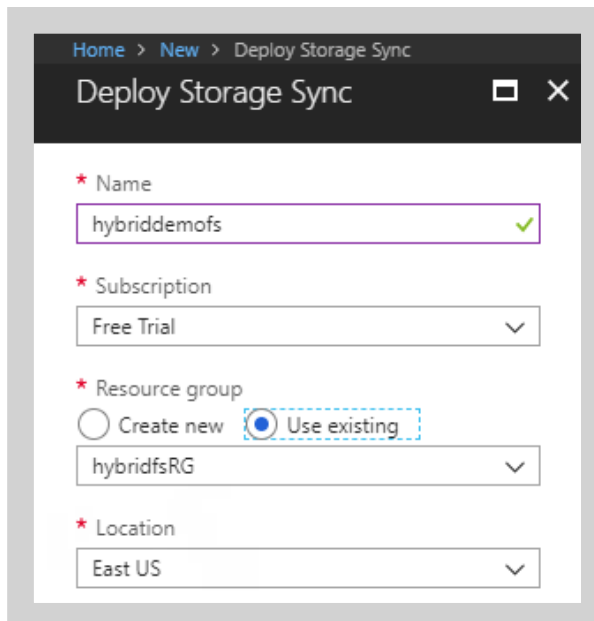




Once the file share is completed, create a new resource in the same region and the same resource group. Select **Storage** and **Azure File Sync**:



Complete the **Deploy Storage Sync** creation blade:



Having created the sync service as above, go ahead and create a new sync group:

Select the storage account you created earlier and the file share in the file service. In the sync group, you will then see a pre-provisioned cloud endpoint.

AZURE FILE SHARE	PROVISIONING STATE	RESOURCE GROUP
hybriddemos	<span style="color: green;">●</span>	hybriddemosrg

Before you can add a server endpoint, you need to register the *hybridfs1* server with the sync service.

To register a server:

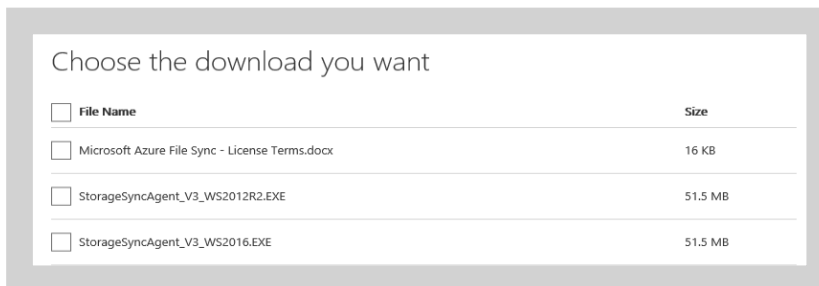
- 1 Download the [Azure File Sync agent](#) and install it on all servers you want to sync.
- 2 After finishing the agent install, use the server registration utility that opens to register the server to this Storage Sync Service. [Learn more](#)

To register the server, download the agent and install it. The process of installation will register the server and make it available to become a server endpoint.

You will need the Microsoft Account credentials for your Azure subscription.

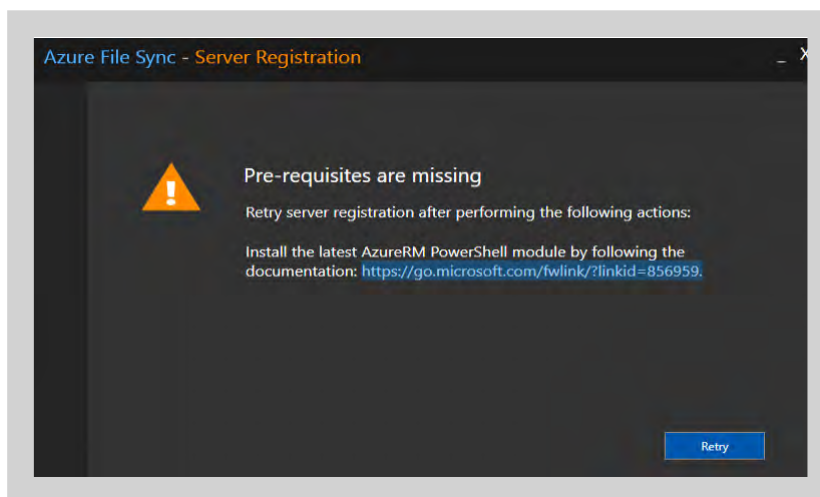
The agent is available in two varieties.

One for Windows Server 2012 R2 and one for Windows Server 2016.

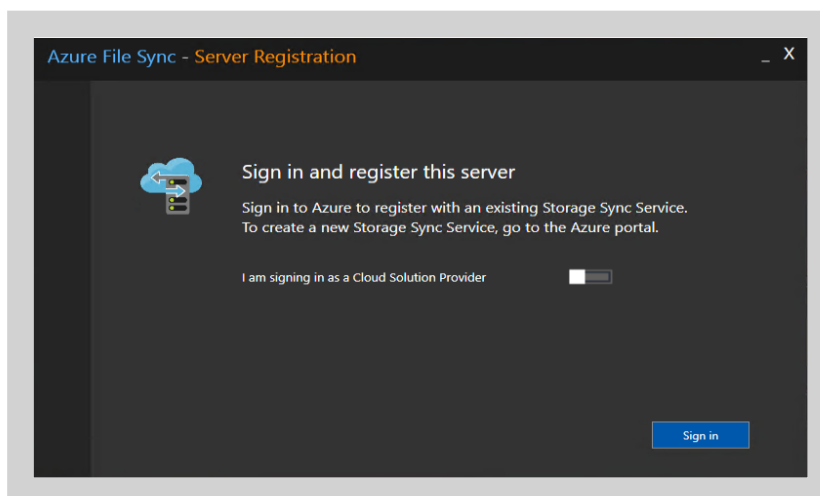


<input type="checkbox"/> File Name	Size
<input type="checkbox"/> Microsoft Azure File Sync - License Terms.docx	16 KB
<input type="checkbox"/> StorageSyncAgent_V3_WS2012R2.EXE	51.5 MB
<input type="checkbox"/> StorageSyncAgent_V3_WS2016.EXE	51.5 MB

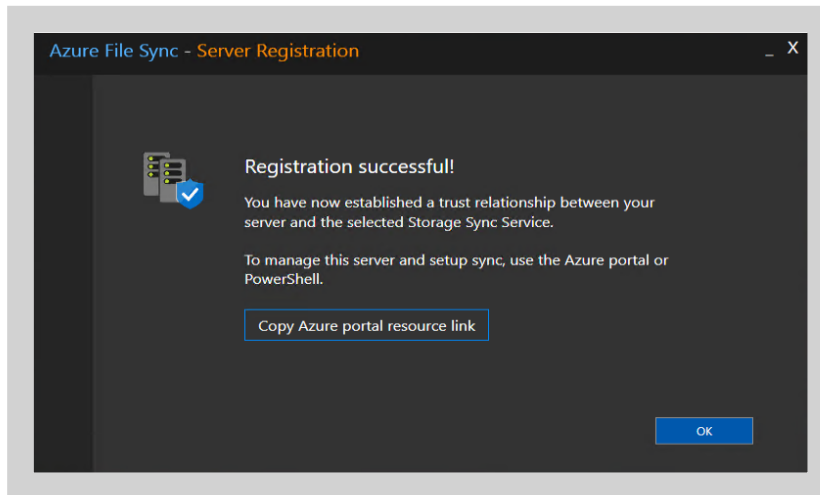
Note that the AzureRM module for PowerShell is a pre-requisite for the agent to successfully install.



Having installed the latest AzureRM PowerShell module, you can continue.



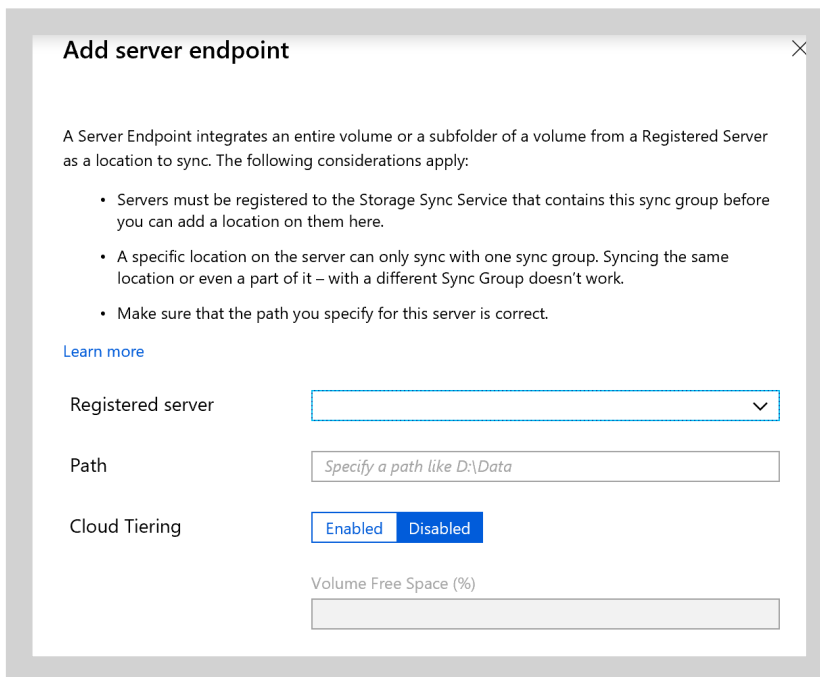
The wizard completes once you have entered the correct credentials.



Having registered the server, it is now possible to create a server endpoint.

Open the Azure portal at <https://portal.azure.com> and navigate to your Azure File Sync agent and the correct Sync Group.

#### Add a server endpoint:



Ensure that you select the correct server and enter the correct path (local) for the file share. If you want to enable the cloud tiering option, you must ensure the file share is not on the system drive.

The setup is now complete. The files in the Azure file share and the local file share will now synchronize.

Note that files added to the Azure file share will only synchronize once every 12 hours or so while the files in the local file share will be synchronized every few minutes. This happens because the Azure file share does not support journaling and a scheduled job will carry out the synchronization periodically.



Now that we have a fully functioning file server and a hybrid synchronized solution, we need to ensure the data is protected. The plan is to protect the local file share using Windows Server Backup (on-premises) and to backup the Azure file share using Azure backup.

In addition, we can use the Azure backup agent to backup the file server to an Azure recovery services vault.

### Backup and recovery – on-premises

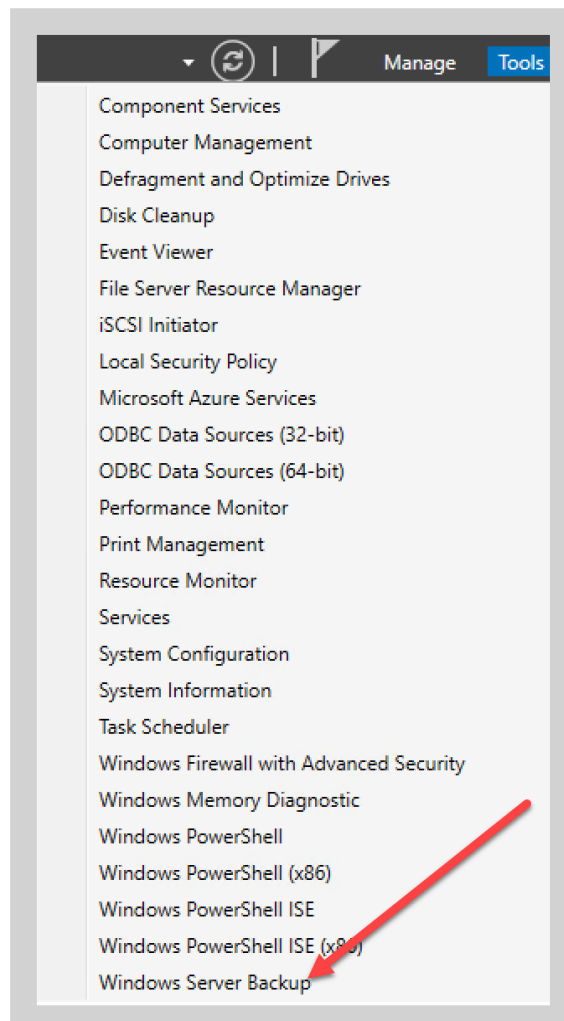
All servers should have a well-defined local backup plan. In this scenario, we are going to use Windows Server Backup. For the purposes of this exercise, we are only going to backup the File Server VM data and system state. It is a best practice to backup all the VMs and the Host server. You might consider using Microsoft System Center Data Protection Manager to provide a complete backup and recovery solution.

We have already installed the Windows Server Backup feature in the original setup. If you did not install it, then you can use Server Manager or PowerShell.

**For PowerShell, run this command in an administrative session:**

```
Install-WindowsFeature -Name windows-server-backup`  
-IncludeAllSubFeature -IncludeManagementTools
```

**Once the feature is installed, the Windows Server Backup application is available in the Tools menu of the Server Manager.**



Run the application.

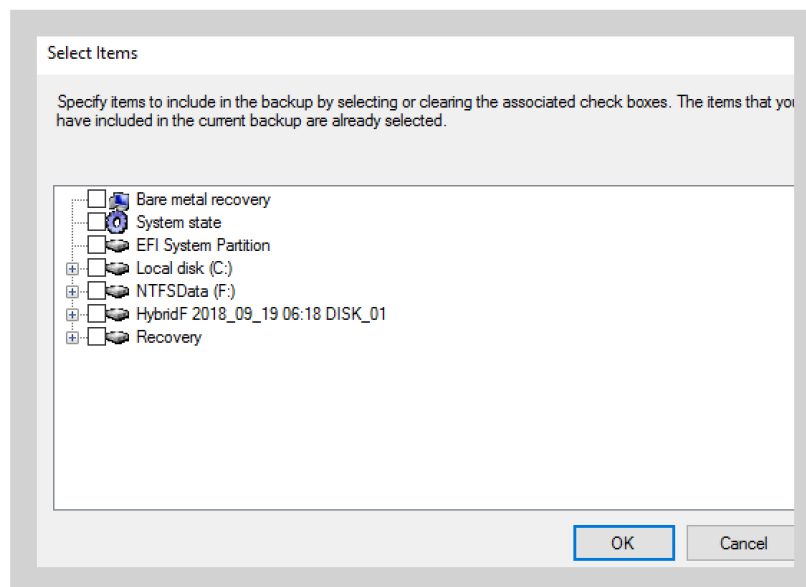
To be able to back up the file shares and the system state, you will need to have an additional external disk, a network file share on a different server, or a separate volume on a disk attached to the server. The choice is yours.

The best practice here for an SMB would be either an external disk or a network share. For ease of use here, we are backing up to a 150GB volume created on a disk in the VM. All backup disks used will be cleaned and held exclusively for the use of Windows Server Backup.

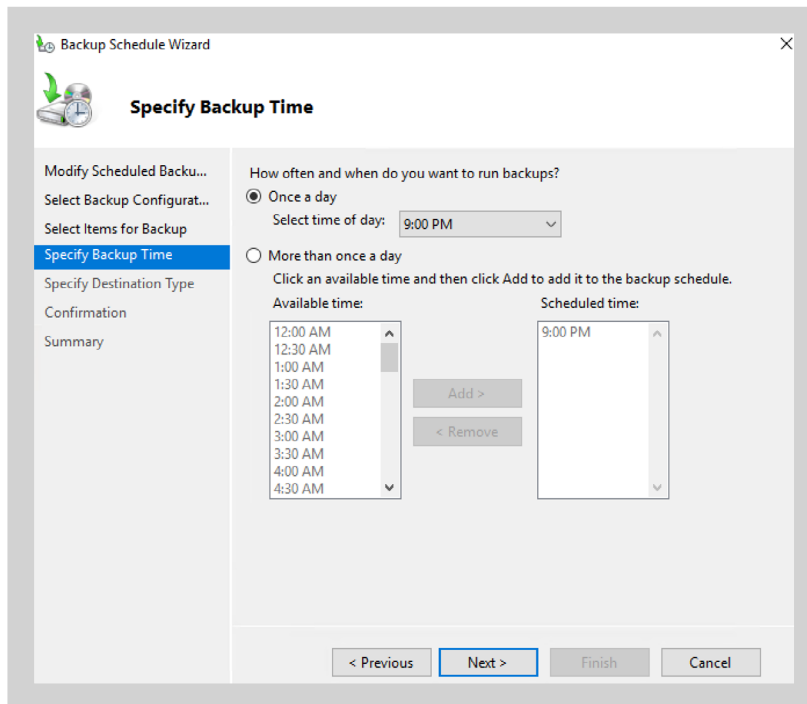
Within the WSB application, create a backup schedule to cover the file share folders on the F: drive. You can also add the system state. (WSB also provides a bare metal option, which is not included here.)

First step in the schedule backup wizard is to choose the backup type – full or custom. Choose custom.

**The next step is to select the files to backup.**

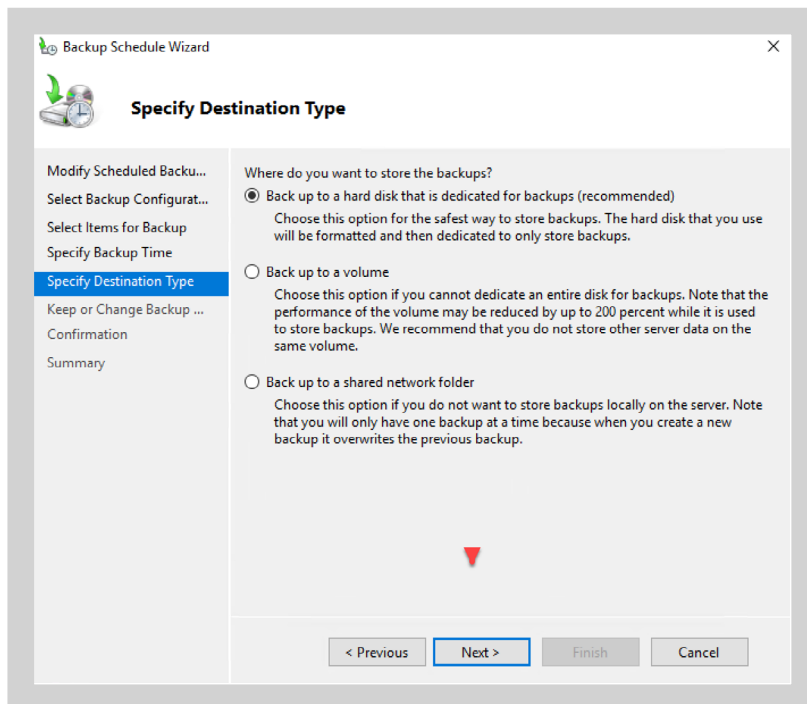


Next, choose the backup times. It's important to choose wisely for your desired recovery point requirements.

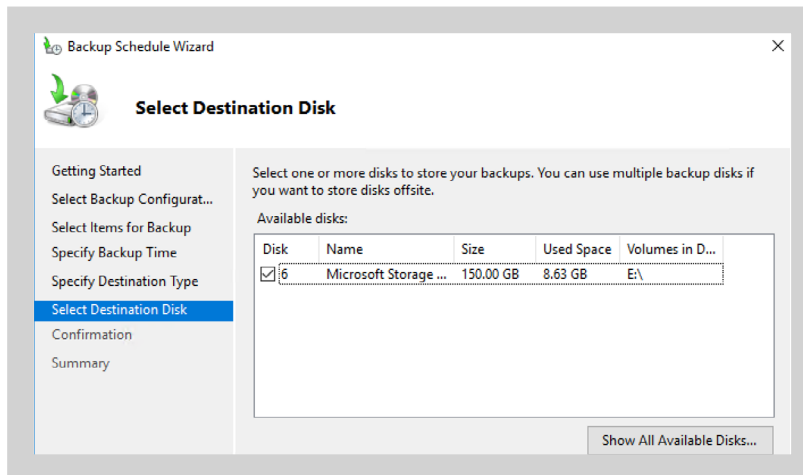


The screenshot shows the 'Specify Backup Time' step of the Backup Schedule Wizard. The left sidebar contains a list of steps: 'Modify Scheduled Backu...', 'Select Backup Configur...', 'Select Items for Backup', 'Specify Backup Time' (highlighted), 'Specify Destination Type', 'Confirmation', and 'Summary'. The main area is titled 'Specify Backup Time' and contains the question 'How often and when do you want to run backups?'. There are two radio button options: 'Once a day' (selected) and 'More than once a day'. Under 'Once a day', there is a 'Select time of day:' dropdown menu showing '9:00 PM'. Under 'More than once a day', there is a section titled 'Click an available time and then click Add to add it to the backup schedule.' This section has two columns: 'Available time:' with a list of times from 12:00 AM to 4:30 AM in 30-minute increments, and 'Scheduled time:' with a list showing '9:00 PM'. Between these columns are 'Add >' and '< Remove' buttons. At the bottom of the wizard are four buttons: '< Previous', 'Next >' (highlighted), 'Finish', and 'Cancel'.

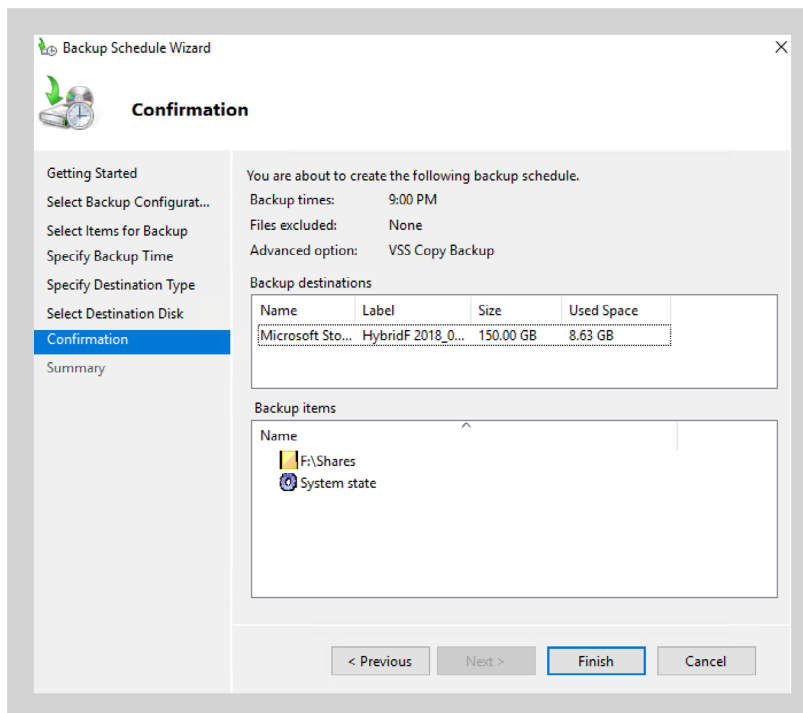
Pick the destination type and location.



The screenshot shows the 'Specify Destination Type' step of the Backup Schedule Wizard. The left sidebar is the same as the previous screen, with 'Specify Destination Type' highlighted. The main area is titled 'Specify Destination Type' and contains the question 'Where do you want to store the backups?'. There are three radio button options: 'Back up to a hard disk that is dedicated for backups (recommended)' (selected), 'Back up to a volume', and 'Back up to a shared network folder'. Each option has a brief description. The 'Back up to a hard disk...' option says 'Choose this option for the safest way to store backups. The hard disk that you use will be formatted and then dedicated to only store backups.' The 'Back up to a volume' option says 'Choose this option if you cannot dedicate an entire disk for backups. Note that the performance of the volume may be reduced by up to 200 percent while it is used to store backups. We recommend that you do not store other server data on the same volume.' The 'Back up to a shared network folder' option says 'Choose this option if you do not want to store backups locally on the server. Note that you will only have one backup at a time because when you create a new backup it overwrites the previous backup.' At the bottom of the wizard are four buttons: '< Previous', 'Next >' (highlighted), 'Finish', and 'Cancel'.



Finally, confirm the settings and create the schedule.



This completes the on-premises portion of the backup scenario.

The next stage is to back up the same files into Microsoft Azure to take advantage of the additional security and recovery options for cloud storage.

It is worth noting that both Windows Server Backup and Azure Backup provide many more features and options than we will be showing in this scenario. This is a starting point to help you begin to extend your on-premises solutions into the Azure cloud.

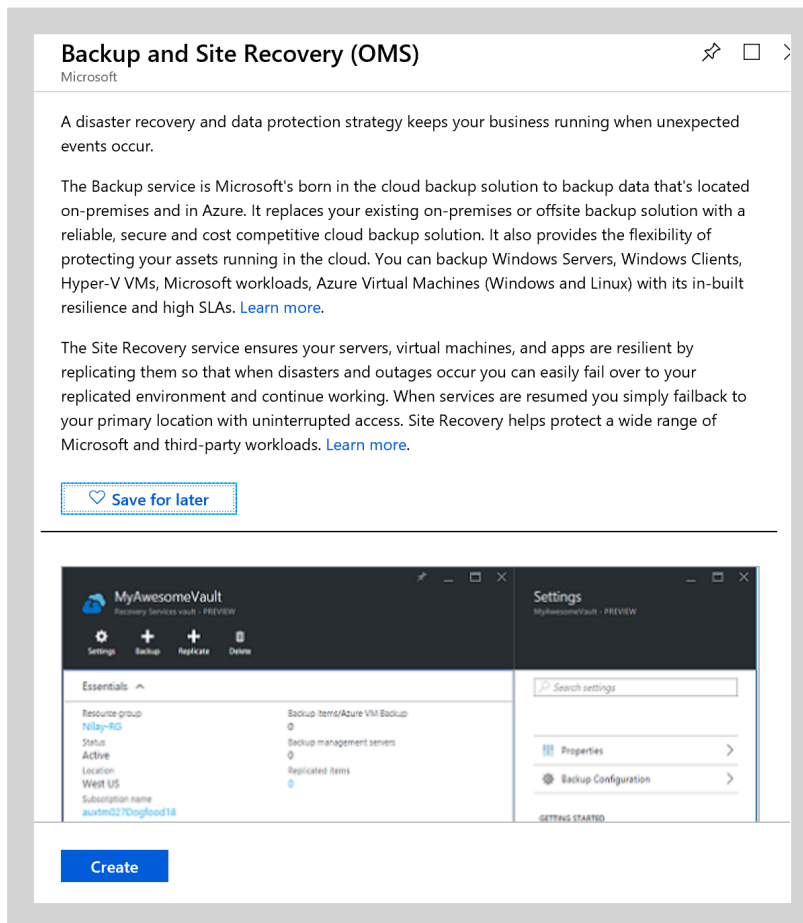
## Backup and recovery (files) – Azure

To backup the file share and system state into Azure requires some infrastructure configuration inside your Azure subscription.

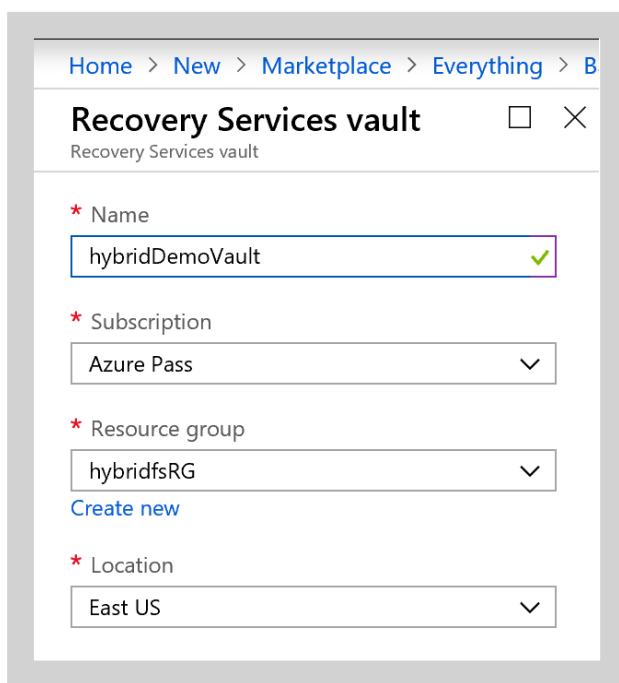
If you want to learn more about Azure Backup before deploying it for production, take a look [here](#). Once you are ready to get it working, then:

Navigate to Azure <https://portal.azure.com>, but do so INSIDE the file server VM.

We need to create a new resource, a Recovery Services vault. This is used to manage and host all your backups and data, as well as the site recovery service, if you choose to replicate your on-premises infrastructure into Azure.



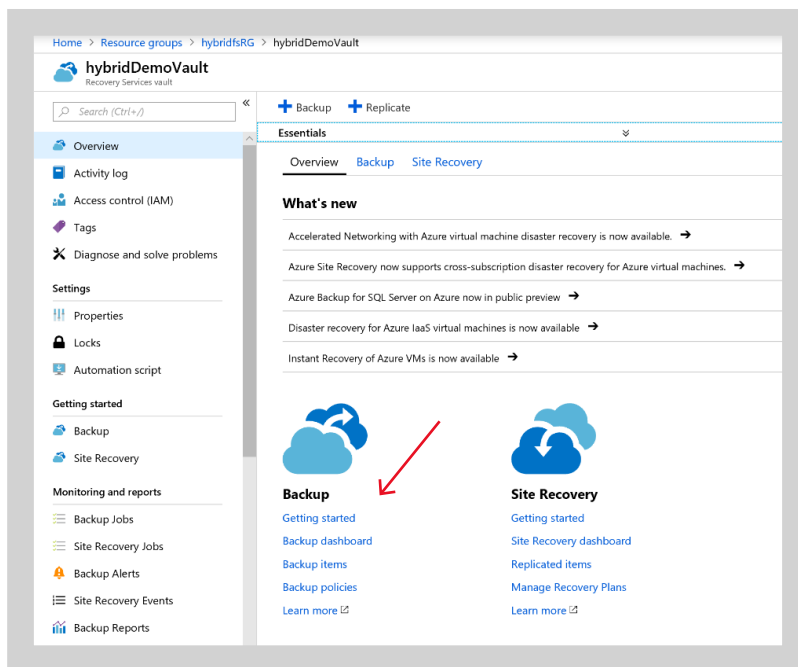
Complete the basic details.



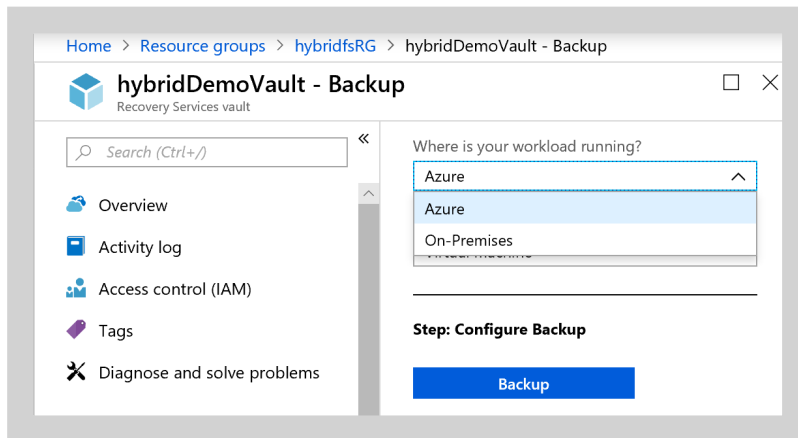
The screenshot shows the 'Recovery Services vault' creation page in the Azure portal. The breadcrumb navigation at the top reads 'Home > New > Marketplace > Everything > B'. The page title is 'Recovery Services vault' with a close button. Below the title, the form fields are as follows:

- Name:** 'hybridDemoVault' (with a green checkmark icon).
- Subscription:** 'Azure Pass' (with a dropdown arrow).
- Resource group:** 'hybridfsRG' (with a dropdown arrow and a 'Create new' link below it).
- Location:** 'East US' (with a dropdown arrow).

Once the vault is created, open the blade and select *Getting started* under Backup.

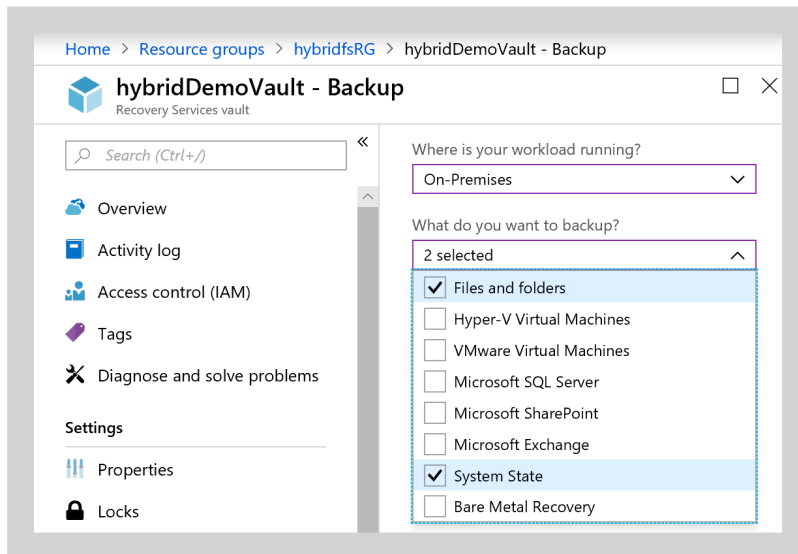


You can use Azure Backup to back up either Azure or on-premises servers:



You can also choose a whole selection of backup types.

Choose *System State* and *Files and folders*.





Choose *Prepare infrastructure* to pull up a selection of tasks to complete:

Prepare infrastructure

**Recovery Services Agent**  
Please follow the steps mentioned below.

1. Install Recovery Services agent  
[Download Agent for Windows Server or Windows Client](#)

2. Download vault credentials to register the server to the vault. Vault credentials will expire after 2 days.  

☐ Already downloaded or using the latest [Recovery Services Agent](#)

Download

3. Schedule backup using Recovery Services Agent UI. [Learn More](#)

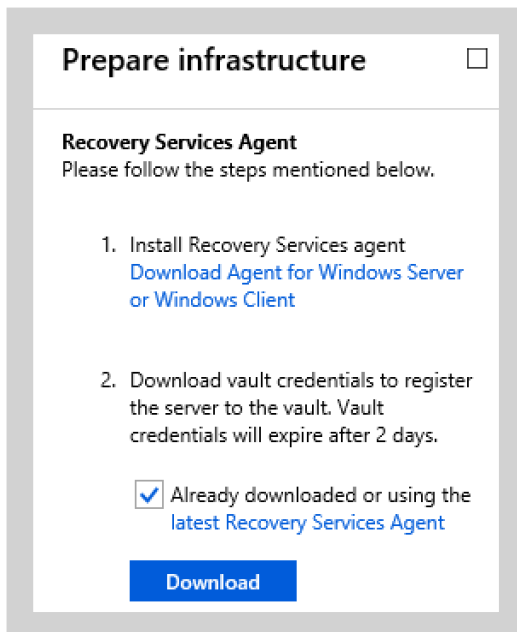
4. Once the backups are scheduled, you can use backup jobs page to monitor the backups. [Browse jobs page](#)

5. You can also Configure Notifications from alerts page to receive email alerts for backup failures. [Browse alerts page](#)

[Learn More](#)

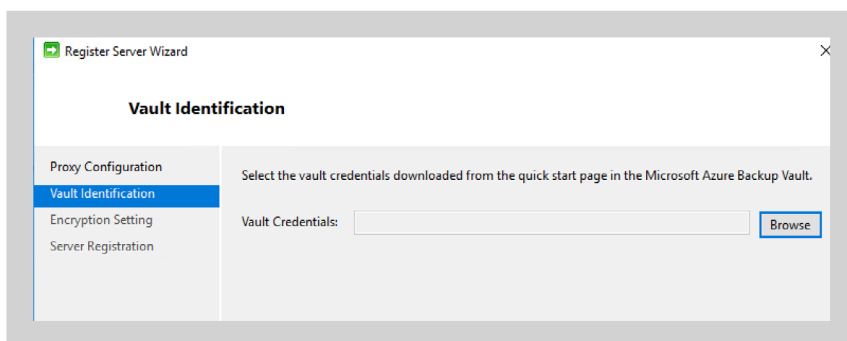

If you haven't opened the Azure portal from within the File Server VM, now is the time to go back and redo that. Then download the Windows Server Backup agent and install it.


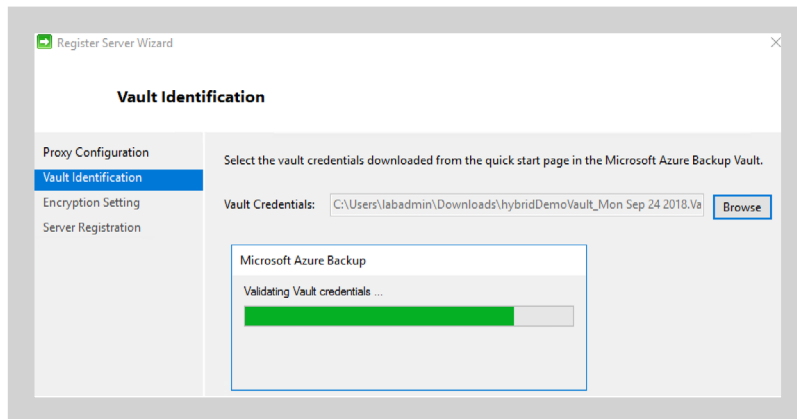
Once you have installed the agent, download the vault credentials:



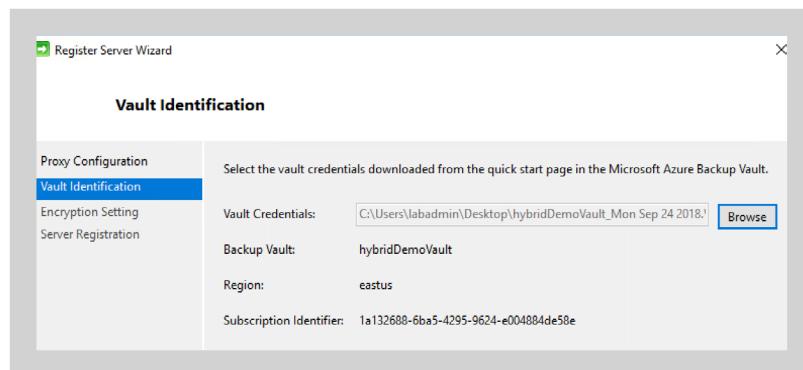
Once the agent is installed, the Azure Backup application (part of the Windows Server Backup application) will run the next step. Click *Register Server*.

This should take you to the registration wizard for vault identification:

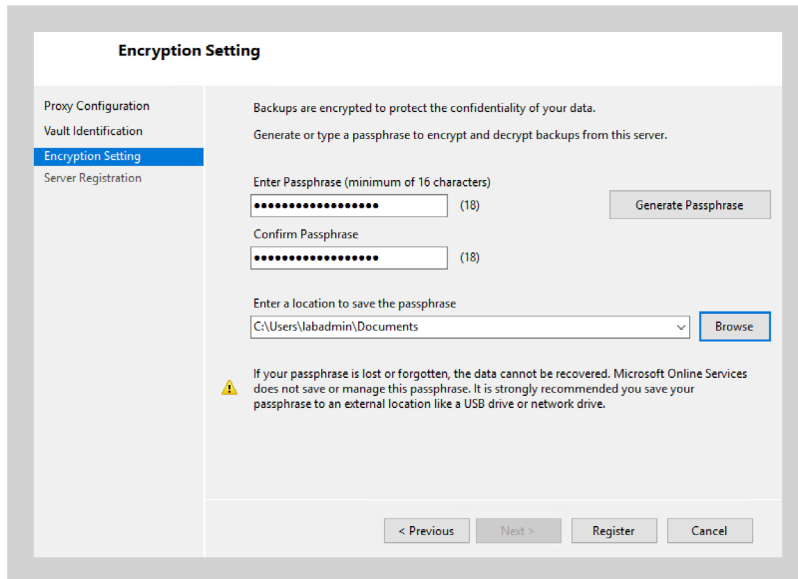




The server will then contact the Azure vault. Connect using the credentials you just downloaded.



In the next stage, choose the passphrase for the data encryption and choose *Register*.



**Encryption Setting**

Proxy Configuration  
Vault Identification  
**Encryption Setting**  
Server Registration

Backups are encrypted to protect the confidentiality of your data.  
Generate or type a passphrase to encrypt and decrypt backups from this server.

Enter Passphrase (minimum of 16 characters)  
 (18) Generate Passphrase

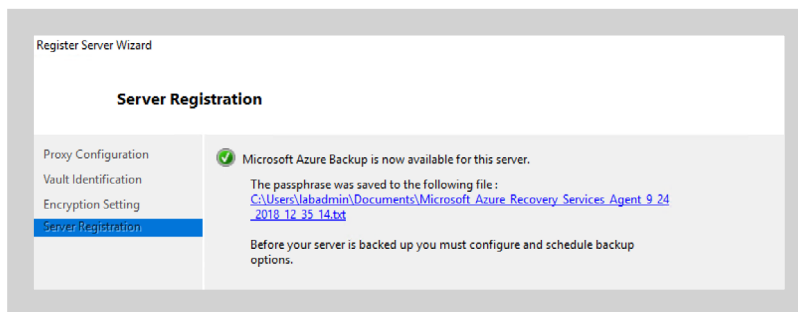
Confirm Passphrase  
 (18)

Enter a location to save the passphrase  
 Browse

**!** If your passphrase is lost or forgotten, the data cannot be recovered. Microsoft Online Services does not save or manage this passphrase. It is strongly recommended you save your passphrase to an external location like a USB drive or network drive.

< Previous   Next >   Register   Cancel

The server and passphrase are now registered.



Register Server Wizard

**Server Registration**

Proxy Configuration  
Vault Identification  
Encryption Setting  
**Server Registration**

✓ Microsoft Azure Backup is now available for this server.

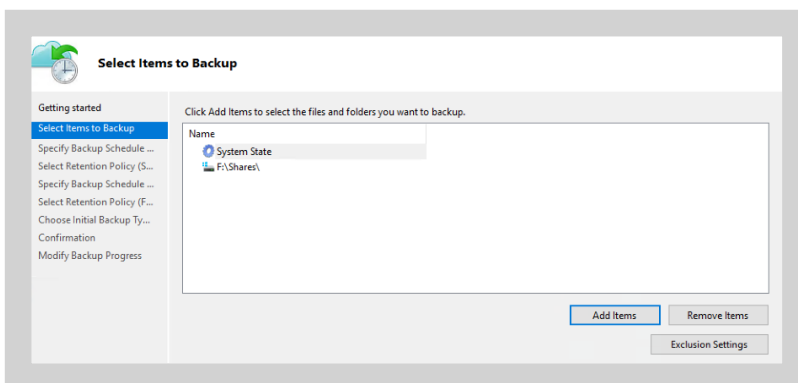
The passphrase was saved to the following file:  
[C:\Users\labadmin\Documents\Microsoft Azure Recovery Services Agent 9\\_24\\_2018\\_12\\_32\\_14.txt](C:\Users\labadmin\Documents\Microsoft Azure Recovery Services Agent 9_24_2018_12_32_14.txt)

Before your server is backed up you must configure and schedule backup options.

Once the agent is configured, having just registered the server, you can then create a new backup schedule for file share data and system state.

The Azure wizard is a little more detailed when both file and system state are selected.

First, select the items to back up:



**Select Items to Backup**

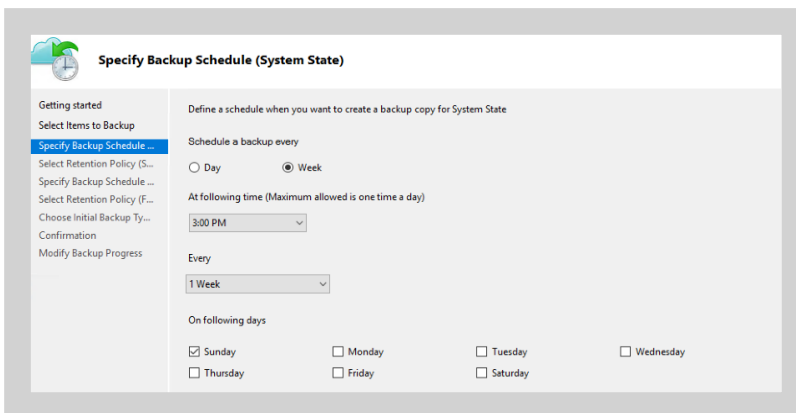
Getting started  
**Select Items to Backup**  
 Specify Backup Schedule ...  
 Select Retention Policy (S...  
 Specify Backup Schedule ...  
 Select Retention Policy (F...  
 Choose Initial Backup Ty...  
 Confirmation  
 Modify Backup Progress

Click Add Items to select the files and folders you want to backup.

Name
System State
F:\Shares\

Add Items Remove Items Exclusion Settings

Then select the system state schedule:



**Specify Backup Schedule (System State)**

Getting started  
Select Items to Backup  
**Specify Backup Schedule ...**  
Select Retention Policy (S...  
Specify Backup Schedule ...  
Select Retention Policy (F...  
Choose Initial Backup Ty...  
Confirmation  
Modify Backup Progress

Define a schedule when you want to create a backup copy for System State

Schedule a backup every

☐ Day ☒ Week

At following time (Maximum allowed is one time a day)

3:00 PM

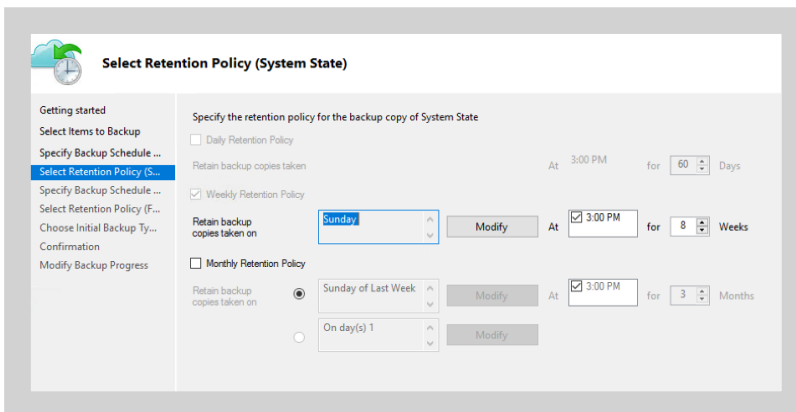
Every

1 Week

On following days

☒ Sunday ☐ Monday ☐ Tuesday ☐ Wednesday  
☐ Thursday ☐ Friday ☐ Saturday

Then select the system state retention policy:



**Select Retention Policy (System State)**

Getting started  
Select Items to Backup  
**Select Retention Policy (S...**  
Specify Backup Schedule ...  
Select Retention Policy (F...  
Choose Initial Backup Ty...  
Confirmation  
Modify Backup Progress

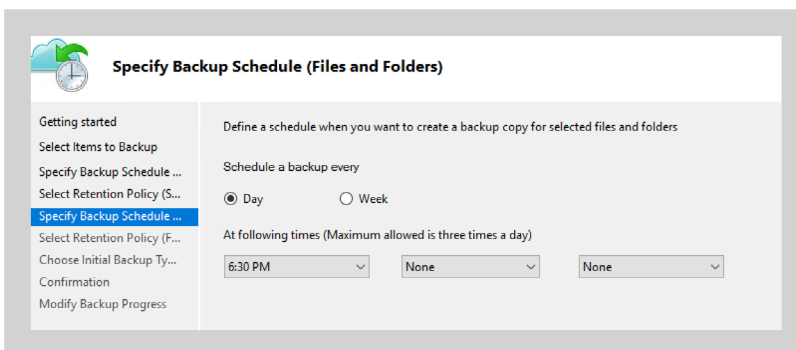
Specify the retention policy for the backup copy of System State

☐ Daily Retention Policy  
Retain backup copies taken At 3:00 PM for 60 Days

☒ Weekly Retention Policy  
Retain backup copies taken on Sunday At 3:00 PM for 8 Weeks

☐ Monthly Retention Policy  
Retain backup copies taken on Sunday of Last Week At 3:00 PM for 3 Months  
On day(s) 1

Next, select the files and folders schedule:



**Specify Backup Schedule (Files and Folders)**

Getting started  
Select Items to Backup  
Specify Backup Schedule ...  
Select Retention Policy (S...  
**Specify Backup Schedule ...**  
Select Retention Policy (F...  
Choose Initial Backup Ty...  
Confirmation  
Modify Backup Progress

Define a schedule when you want to create a backup copy for selected files and folders

Schedule a backup every

☒ Day ☐ Week

At following times (Maximum allowed is three times a day)

6:30 PM None None

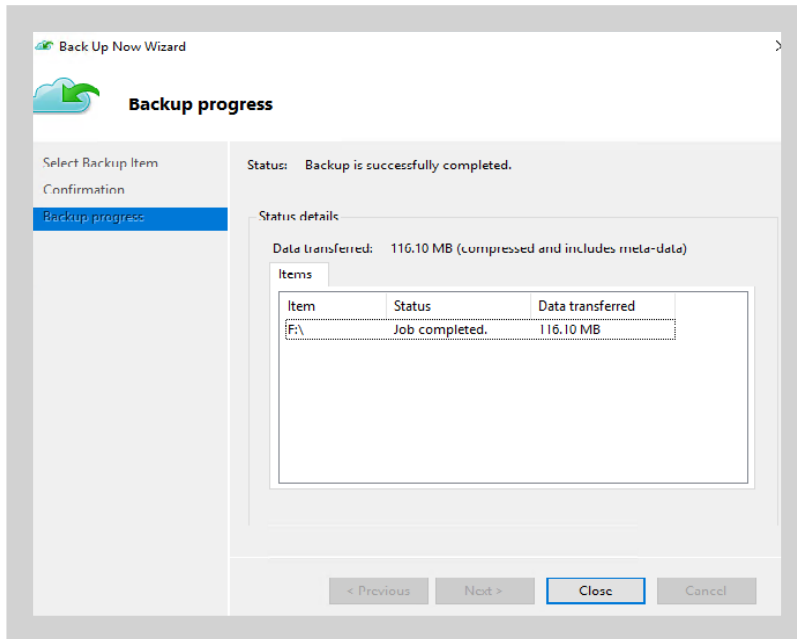
And then choose the files and folders retention policy:

The initial backup option assists if the initial backup is likely to be very large, in which case a different process can assist with importing the data to Azure. In our example case, it is a small backup and can be carried out over the network.

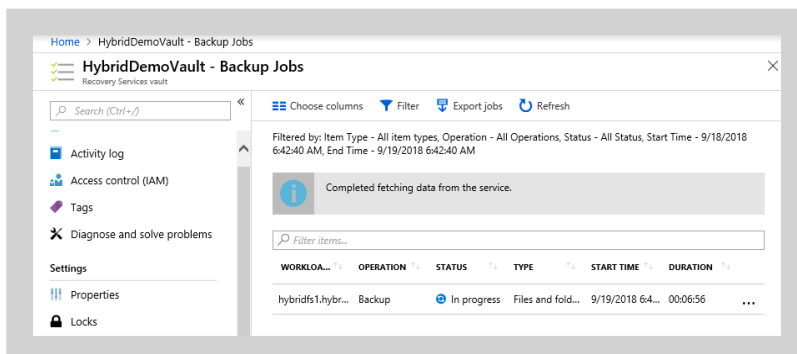
Finally, confirm the settings and create the schedule:

Policy Type	Selected Items	Excluded Items
System State	Selected Items (System State)	None
Files and Folders	Selected Items (F1)	None

When the schedule runs, you can see it in the Azure backup agent on-premises and also in Azure itself:



The added benefit is that you can track and manage the backup jobs and data from the Azure portal.



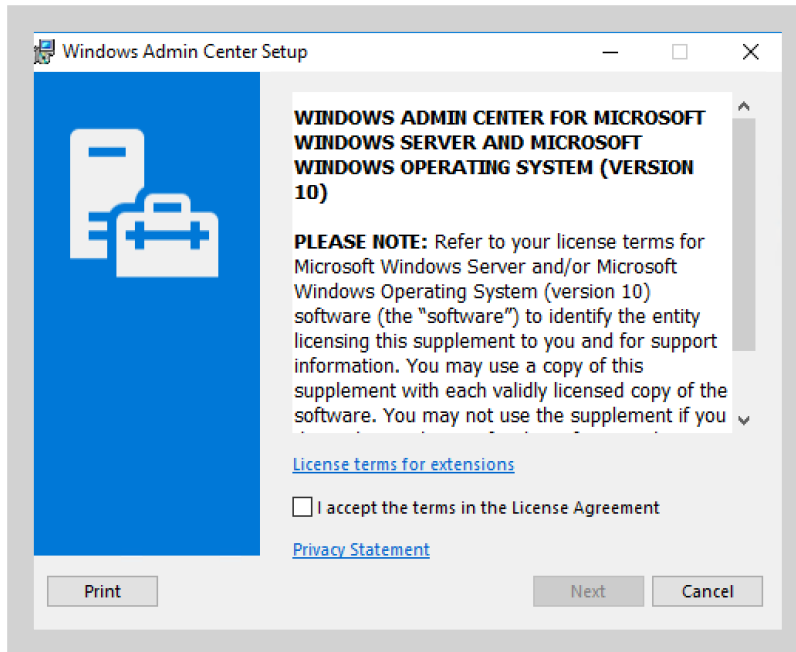
And this completes the setup of the full hybrid solution.

We have built the infrastructure on-premises and the file service in Azure and provided a file sync solution, as well as a combined on-premises and Azure backup solution.

## 3: Appendix

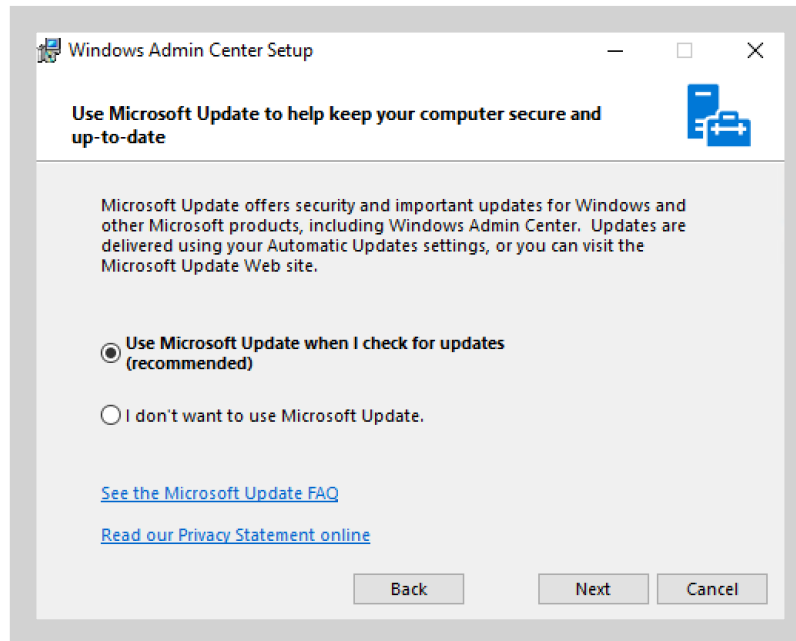
### Windows 10 client running Windows Admin Center (WAC)

Having installed the Windows 10 VM and joined it to the hybriddemo.cloud domain, navigate to and download [Windows Admin Center](#): (at least version 1809).



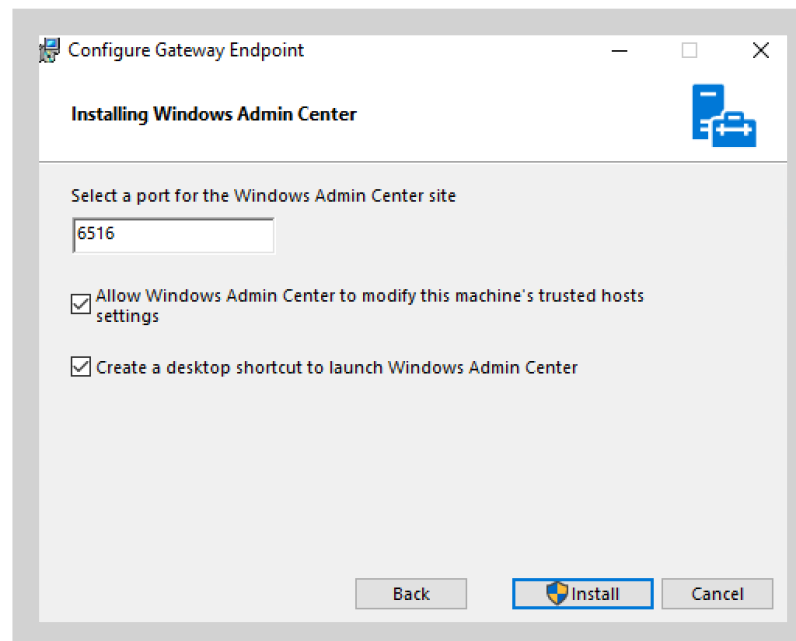


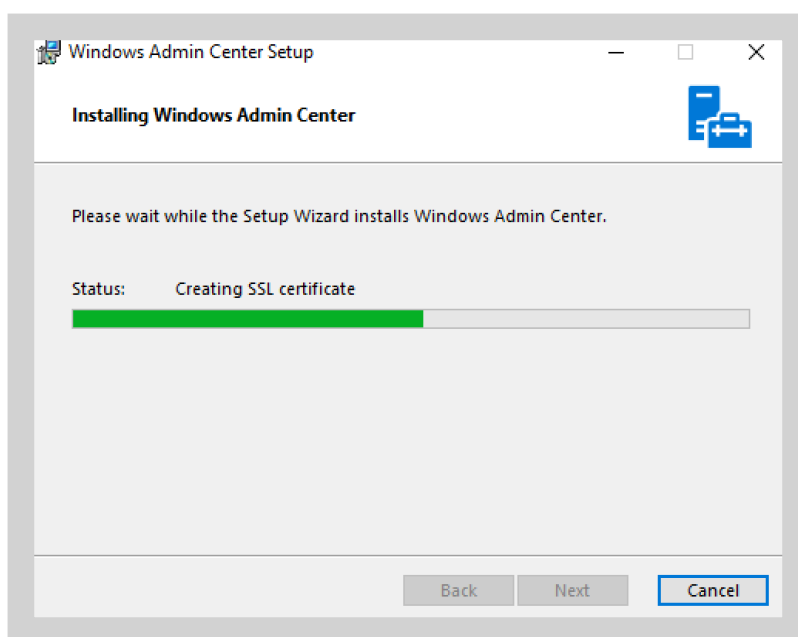

To run the application, accept the terms of the License Agreement and install WAC:

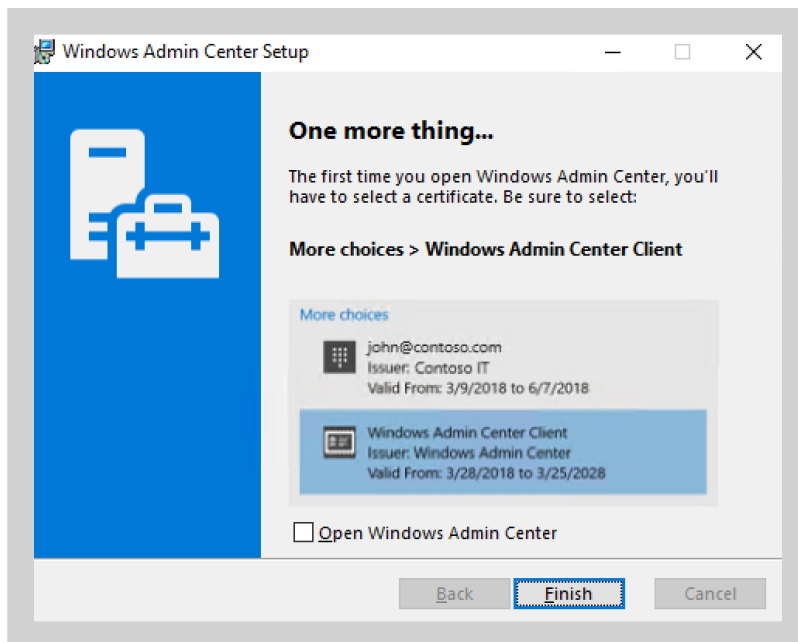


If you select use Microsoft update, WAC will be updated when Windows is updated.

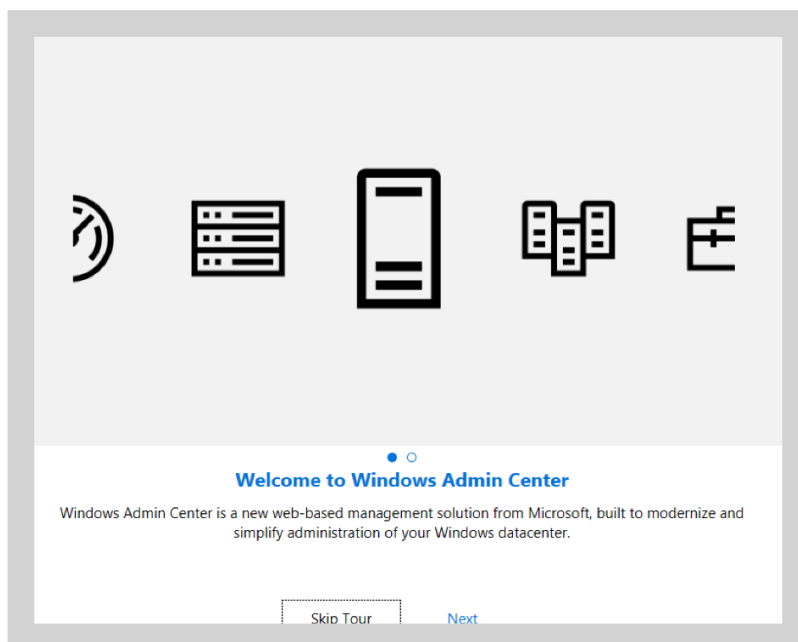
To start WAC installation process, choose a port number for the website (default is 6516). If you are not in a domain environment, you will need to allow WAC to amend the trusted hosts lists to allow connecting







to non-domain joined devices.



Once installed, WAC will open with a short tour of its features.

You need to add the DC and File server to the console by entering the DNS names into the *Add Server Connection* blade.

**Add Server Connection**

Connection Tags

+ Add Tags

Add one server Import a list of servers

☐ Don't attempt to resolve the server name.

Server name

hybridfs1.hybriddemo.cloud

✓ Found 'hybridfs1.hybriddemo.cloud'

Windows Admin Center

**Windows Admin Center**

All Connections

+ Add Connect Manage As Remove Edit Tags 3 items

Name ↑	Type	Last Connected	Managing As
hybridclient.hybriddemo.cloud	Windows PC	Never	HYBRIDEMO\labadmin
hybriddc.hybriddemo.cloud	Server	Never	HYBRIDEMO\labadmin
hybridfs1.hybriddemo.cloud	Server	Never	HYBRIDEMO\labadmin

If credentials are required, they will be requested.

Windows Admin Center Server Manager

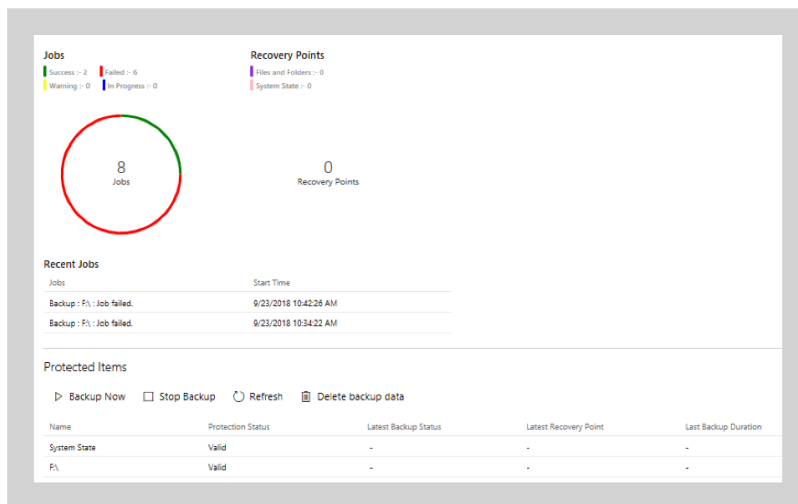
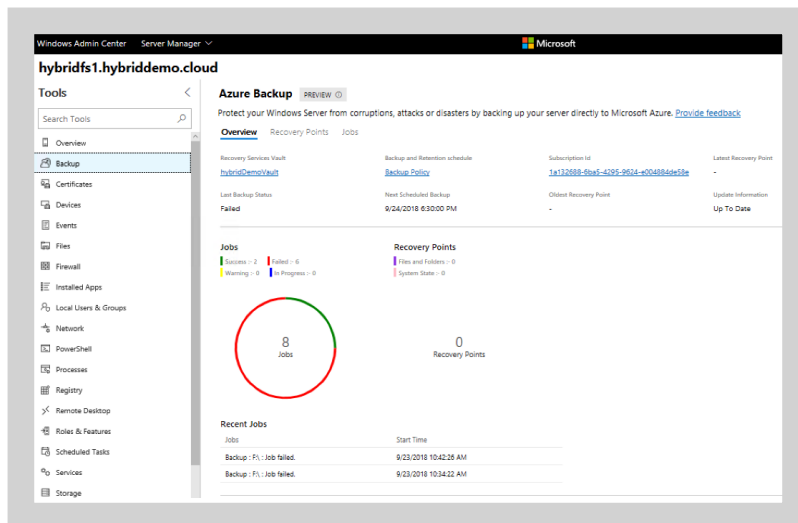
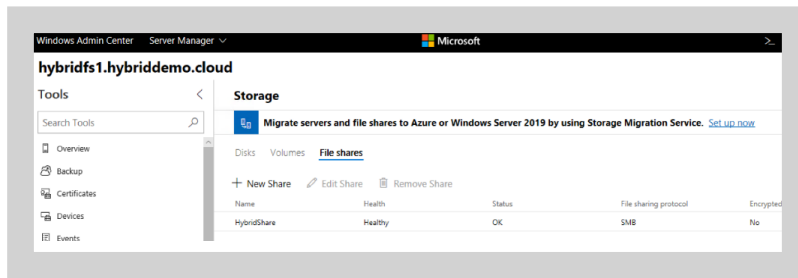
**hybridfs1.hybriddemo.cloud**

Tools Overview

Restart Shutdown Enable Disk Metrics Edit Computer ID Refresh

Computer Name	Domain	Operating System
hybridfs1	hybriddemo.cloud	Microsoft Windows Server 2016 Datacenter
Version	Installed Memory (RAM)	Disk Space (Free / Total)
10.0.14393	7.91 GB	252.61 GB / 281.91 GB
Processors	Manufacturer	Model
Intel(R) Xeon(R) CPU E7-8870 @ 2.40GHz	Microsoft Corporation	Virtual Machine
Logical Processors	Windows Defender	NIC(s)

Click on *hybridfs1* (or the name of your file server)



Explore the overview, the storage area for file shares and the backup. Notice the Azure integration for backup:

For more information about Windows Admin Center, check [here](#).

Feel free to explore the other sections of the Windows Admin Center.. You will find that almost all of the standard Windows Server management tools are included.