



---

TECHNICAL WHITE PAPER

# Protecting Azure NetApp Files with Rubrik

Amrutha Naik, Ed Morgan, Pierre-François Guglielmi  
April 2020  
RWP-0502

---

# TABLE OF CONTENTS

## **3 ABOUT THIS WHITE PAPER**

## **3 EXECUTIVE SUMMARY**

## **3 CUSTOMER CHALLENGES**

## **4 HOW RUBRIK SOLVES THESE CHALLENGES**

## **4 RUBRIK OVERVIEW**

### 4 Fileset

### 5 SLA Domain Policy

### 6 Universal Fileset

What is a Rubrik Fileset?

Why Did Rubrik Choose To Use Filesets?

How Do Rubrik Filesets Protect Large Workloads?

## **8 NAS FILESET PROTECTION**

### 8 Sample NAS Fileset

### 10 NAS Share Backup Workflow

### 11 NAS Share Recovery Flow

## **12 PROTECT AZURE NETAPP FILES USING RUBRIK**

## **22 BACKUP AND RESTORE PERFORMANCE**

### 22 Test Environment

### 23 Backup Performance

First Full Backup

Incremental Backup

### 25 Restore Performance

## **25 CONCLUSION**

## **26 ABOUT THE AUTHORS**

## ABOUT THIS WHITE PAPER

The primary objective of this white paper is to study the different use cases that can be leveraged with the marriage of Rubrik & Azure NetApp Files. This white paper is intended to provide architects and data center administrators information about the implementation and benefits of Rubrik's integration with Azure NetApp Files, steps to configure Rubrik to protect Azure NetApp Files, and performance benchmarks of the solution.

## EXECUTIVE SUMMARY

As the amount of unstructured data continues to grow exponentially, enterprises face the daunting task of ensuring that critical data on Network Attached Storage (NAS) systems are fully protected, and the digitization of business requires enterprises to move faster and be more agile to survive. Applying new technologies to existing business activities (e.g., leveraging AI to increase customer satisfaction) will continue to fuel the cloud paradigm. According to IDC, enterprises will spend more than \$500 billion on cloud and cloud services by 2021, with 80% of application development on cloud platforms. For many enterprises, public cloud represents the ability to rapidly access resources for innovation while operating in a data-rich environment. With this exponential increase in cloud adoption, it is paramount to have a matching data protection and management solution.

For the first time, Microsoft will deliver an Azure native, first-party service for enterprise NFS/SMB file services based on NetApp ONTAP technology. This new development is driven by a strategic partnership between NetApp and Microsoft and further extends the reach of NetApp's world-class data services to Azure. This Azure cloud-native data service delivers high performance, reliability, and enterprise data management and security for customers who are moving enterprise NFS/SMB workloads to Azure.

The data service is Azure NetApp Files, the industry's first bare-metal cloud file storage and data management service. Azure NetApp Files (ANF) is a new Azure service, delivered by Microsoft and built on NetApp's industry leading ONTAP technology directly in the Azure data centers. The service allows customers to move their NFS and SMB file-share workloads into Azure – even legacy applications – without rearchitecting their applications. Users can reduce their migration from months or years, to days or weeks. Azure NetApp Files can be provisioned and managed seamlessly within the Azure portal. It has the same billing, CLI, and deployment paradigm as any other Azure service; the dedicated environment in Azure means that performance is both optimal and assured. Azure NetApp Files eliminates the need for time-consuming and expensive architectural change, delivering an easy way to seamlessly provision file-based workloads in Azure and helping organizations to meet their cloud mandate in record time. With Azure NetApp Files, any file-based workload can move to Azure, including NFS v3 and even SMB shares, with no change. To increase storage, simply choose from multiple levels of guaranteed performance seamlessly through the Azure portal – and change performance on the fly without moving data or creating any new volumes.

## CUSTOMER CHALLENGES

Some very common customer challenges with NAS protection include:

- Large file servers with billions of files are often particularly hard to handle as they may take hours to days for a full backup.
- Unsatisfactory backup performance or elongated recovery times are challenges that backup administrators typically encounter when protecting file data.
- Backup administrators often begin each day by remediating failed backups from the previous night, thus increasing operational management efforts.

## HOW RUBRIK SOLVES THESE CHALLENGES

Rubrik takes a modern incremental-forever approach, eliminating the need for periodic full backups. Performing incremental-forever backups reduces nightly backup windows, especially when there is a need to protect millions or billions of files. From a restore perspective, each backup looks like a full backup whether restoring a single file, a folder, or a full volume. This approach works not only for small enterprise NAS systems/file servers but even file services on cloud such as Azure NetApp Files. Operational management is minimized thus providing valuable production time for backup administrators to focus on more closely aligned business objectives.



## RUBRIK OVERVIEW

Let's understand a little more about the different building blocks of Rubrik in relation to NAS protection and get familiar with the common terminologies used in Rubrik NAS protection.

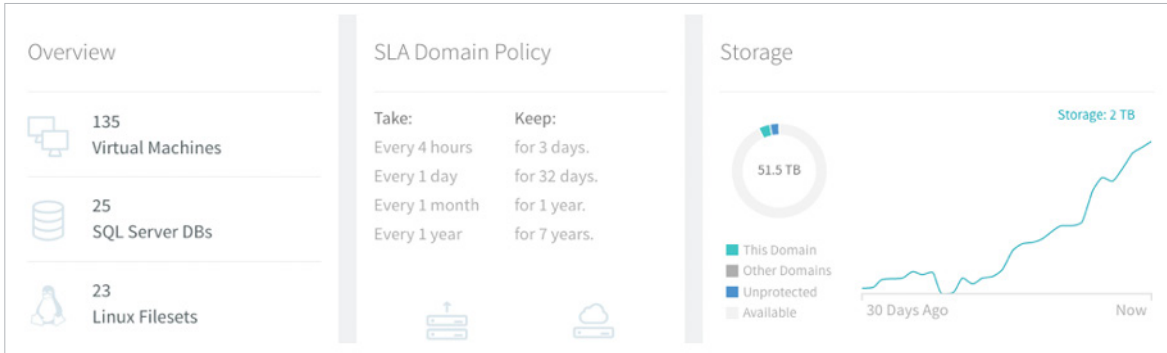
### FILESET

A fileset is a rule that defines the set of folders and files that should be backed up. The benefits of protecting NAS file data with filesets include:

- **Policy-Driven Management** — filesets are linked to an SLA Domain, which provides a configurable set of policies that can be applied to groups of virtual machines, applications, and hosts to achieve specific data protection objectives.
- **Centralized Management** — providing visibility through the Rubrik UI of all file data that has been protected and ready for granular recovery.
- **Incremental-forever backups** — via block mapping and intelligent metadata algorithms to dramatically reduce local storage requirements, provide much faster backups, and reduce network usage.
- **Multiple Recovery Options** — filesets provide granular file and folder-level protection with the ability to exclude sub-paths. Users can choose to restore files or folders to the original server, export to a different server, or download to the browsing system. There are many variations and granularity available to accomplish your recovery needs.

## SLA DOMAIN POLICY

As with all datasets, Rubrik's data management approach begins with SLA Domain Policies. A Rubrik SLA Domain Policy is a declarative policy encompassing the core items needed for backup and recovery, replacing the need to individually configure jobs, tasks, and other items. SLA Domain Policies are a core part of Rubrik's architecture and extend across all data types, as shown below:



Let's walk through the pieces needed to configure an SLA Domain Policy that can apply to all data types:

- **Backup Frequency:** Also known as the Recovery Point Objective (RPO), this describes the frequency in which backups are taken.
  - For file data, this determines how often a restore point is captured to provide recovery of a single file or multiple files. Due to the incremental-forever nature of Rubrik, some customers choose to protect more frequently than the traditional daily backup.
- **Availability Duration:** Also known as retention. Or, how long are backups retained?
  - For file data, retention can widely vary whether due to customer policy or regulatory and compliance requirements. For many customers, this often ranges from a week to 90 days but may be multiple years in order to meet regulatory compliance.

**Service Level Agreement**  
Choose how often we take Snapshots, and the length of time we keep them.

Backup Frequency	Take Snapshots:	Keep Snapshots:	Availability Duration
	Every (Hours) <b>4</b>	For (Days) <b>3</b>	
	Every (Days) <b>1</b>	For (Days) <b>32</b>	
	Every (Months) <b>1</b>	For (Years) <b>1</b>	
	Every (Years) <b>1</b>	For (Years) <b>2</b>	

- **Archival Policy:** Archive targets can be public cloud (AWS, Azure, or Google Cloud Platform) or on-premises (S3 compatible object stores, NFS, or tape). This dictates which archive target is used and when archives are maintained solely in the cloud and not on the local Rubrik cluster. If archives are maintained solely in the cloud (past 30 days for instance), RTO is longer due to the time required to retrieve data back to the Rubrik cluster.
  - For file data, archival can dramatically reduce long-term storage costs, especially when retention is required for years due to regulatory or compliance reasons. The amount of data retained on the local Rubrik cluster is most often determined by the time frame of restore requests. For many customers, this ranges from 30 days to 1 year depending on when most recoveries tend to occur.

Archival Policy

☒ Enable Archiving

☐ Enable Instant Archive ⓘ

Archival Location

Azure:se3demo

Move the slider to choose how long data is kept on the local Rubrik cluster before archiving.

66 days

Snapshots will be stored on the local Rubrik cluster for 66 days . Data will then be moved to your archival location and kept there for 6 years 299 days. Snapshots older than 7 years will no longer be available.

- **Replication Policy:** this relates to Disaster Recovery (DR). Simply put, how much replicated data should be maintained at a DR site?
  - For file data, this is often determined by how far back in time a customer would need to be able to restore files if there is a DR event. Some customers choose to keep the same amount of data on the remote Rubrik cluster as on their local Rubrik cluster, while others choose to have a shorter timeframe available on the remote Rubrik cluster to provide cost savings.

Replication Retention Policy

☒ Enable Replication

Move the slider to choose how long data is kept on the replication target. Leftmost means only the most recent replicated snapshot will be maintained.

30 days

Snapshots will be stored on the replication target for 30 days.

## UNIVERSAL FILESET

Rubrik filesets integrate with SLA policies and are universal in that various operating systems (Linux, AIX, and Windows) and NAS devices are protected utilizing the same fileset concept. There are both similarities and differences in how Rubrik manages data from Linux, AIX, Windows and NAS hosts. The similarities exist due to the way Rubrik utilizes filesets, and the differences are managed by Rubrik without requiring customer interaction after initial configuration.

## WHAT IS A RUBRIK FILESET?

As shown below, a fileset is the combination of full paths or filenames to define objects to include or exclude from backup. The Rubrik cluster interprets the syntax entered for the inclusion or exclusion fields of filesets by following a relative path rule. In the case of NAS filesets, after defining the share path, any associated filesets only need to include the desired path after the share.

The 'Add Fileset' dialog box contains the following elements:

- Fileset Name:** A text field containing 'NAS Home Dirs' with a file icon on the right.
- Share Type:** Radio buttons for 'NFS' (selected) and 'SMB'.
- Rules:** A section with a help icon and instructions: 'Use \*\* to include all files'. It includes three text areas:
  - Include:** Contains the example path `/usr/local/*.pdf`.
  - Exclude:** Contains the example path `/usr/local/temp/*.mov,*.mp3,*.mp4`.
  - Do Not Exclude:** Contains the example path `/company/*.mp4`.
- Enable Backup of Hidden Folders:** A checked checkbox.
- Buttons:** 'Cancel' and 'Add' buttons at the bottom.

## WHY DID RUBRIK CHOOSE TO USE FILESETS?

Filesets were designed to give customers more granular control while still providing a construct that scales in a simple way for large environments where operational overhead is challenging.

## HOW DO RUBRIK FILESETS PROTECT LARGE WORKLOADS?

If the data protected by a fileset exceeds 100 GB, Rubrik automatically subdivides incoming datasets into 100 GB partitions. Partitioning filesets provides increased ingest and restore performance as the workload can be distributed across all available nodes in the cluster. This distribution is done transparently from a customer perspective and allows Rubrik to protect many files as a single fileset. Further, filesets continue to scale horizontally as more nodes join the Rubrik CDM platform.

Filesets can be defined for Linux, AIX and Windows servers or NAS shares. For Linux, AIX and Windows (whether physical or virtual), Rubrik provides data protection for file systems through the pairing of the host with a fileset and an SLA policy. A single host can also be paired with multiple filesets, and each fileset can even be assigned to a different SLA Domain if desired.

For NAS, the Rubrik cluster pairs a fileset with NAS shares. The pairing of multiple filesets to a single share is similarly permitted and different SLA Domains for each share fileset are allowed. SLA assignment continues to match the core simplicity of the Rubrik architecture thus allowing the same SLA policy for VM's, Linux, AIX, Windows, or SQL also to be assigned to NAS datasets.

## NAS FILESET PROTECTION

NAS shares are protected similarly to Linux, AIX and Windows hosts and do not rely on Network Data Management Protocol (NDMP). Not using NDMP is a conscious design choice that further solidifies Rubrik as vendor agnostic, provides true incremental-forever for NAS, and stores data in a “native” format.

Most traditional backup applications still utilize NDMP. As a legacy control protocol created to move data between a backup server and the tape drives/libraries, NDMP has significant disadvantages. The disadvantages begin with the fact that NDMP was created decades ago when datasets were much smaller and physical tape was the primary backup media. In addition, NDMP does not dictate the format of the backup stream. As a result, each vendor sends the backup stream in a proprietary format that is usually not meant to be unpacked.


NAS file systems are exposed through either NFS or SMB (sometimes referred to as CIFS) protocols. Rubrik can backup file systems exposed through either protocol and preserve the corresponding metadata, including Access Control Lists (ACLs). Our approach is agnostic to the NFS protocol version and supports NFS 3.x and NFS 4.0. Similarly, for SMB, we support SMB 1.x, 2.x, and 3.x. This method includes storing data in a native format, backing up in an incremental-forever fashion, and instant access to all of the data management capabilities of the Rubrik platform including replication to DR, cloud archive, global search, erasure coding, and more. To further simplify datacenter operations, the Rubrik approach does not require NDMP accelerator nodes or “proxy VMs.”

### SAMPLE NAS FILESET

This section provides a few example use cases for defining NAS filesets:

- **Perform a one-level backup of all files in the root path**


```
Include: \*
Backup data:
\users.txt
```



- **Perform a recursive backup of everything under the root path**


```
Include: \
Backup data:
users.txt
Accounting
Userdata
Userdata\template.txt
Userdata\Sarah
Userdata\Sarah\file1.txt
Userdata\James
Userdata\James\file2.txt
Userdata\James\folder
Userdata\James\folder\file3.txt
```





- Perform an empty backup of directories that start with capital A (case sensitive) under \\Userdata

```
Include: \\Userdata\\A*
Backup data:
Userdata\\Alice
```

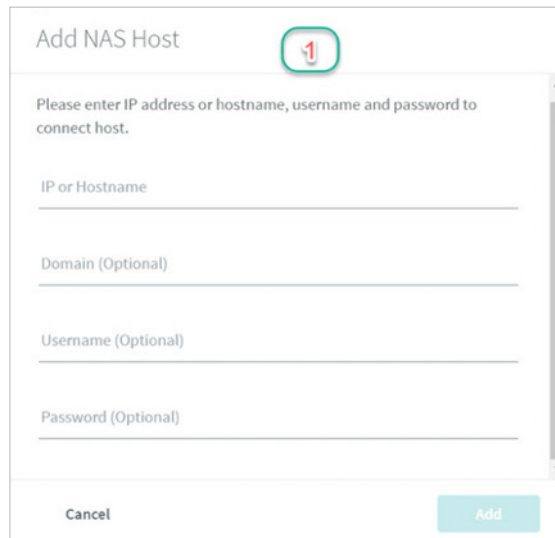


## NAS SHARE BACKUP WORKFLOW

As NAS appliances do not support the installation of any third-party agents, Rubrik manages and protects data in NAS shares by connecting directly to NAS appliances. The user can begin maintaining and protecting a NAS host by adding the host to the Rubrik cluster. Read permissions across the entire NAS share are required.

Rubrik provides a streamlined process for setting up NAS backups and having the Rubrik cluster connect directly to NAS servers. Setup is simple:

1. Add the IP address or FQDN and supply appropriate credentials:



**Add NAS Host** 1

Please enter IP address or hostname, username and password to connect host.

IP or Hostname

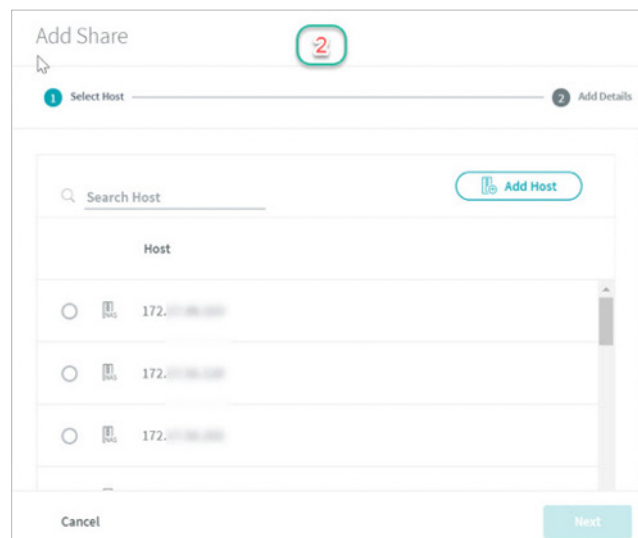
Domain (Optional)

Username (Optional)

Password (Optional)

Cancel Add

2. Select a NAS host or add a new one:



**Add Share** 2

1 Select Host 2 Add Details

Search Host Add Host

Host

172.17.0.100

172.17.0.101

172.17.0.102

Cancel Next

3. Add share details such as the protocol and the path:

3

✓ Select Host — Add Details

Specify share type, path and credentials to add a share

Share Type ☒ NFS ☐ SMB

NFS Path (/mnt/export)

Domain (Optional)

Username (Optional)

Password (Optional)

Cancel Back Finish

## NAS SHARE RECOVERY FLOW

NAS Share, Windows, and Linux filesets have similar recovery methods that allow users to search or browse and select individual files or folders to restore, export, or download.

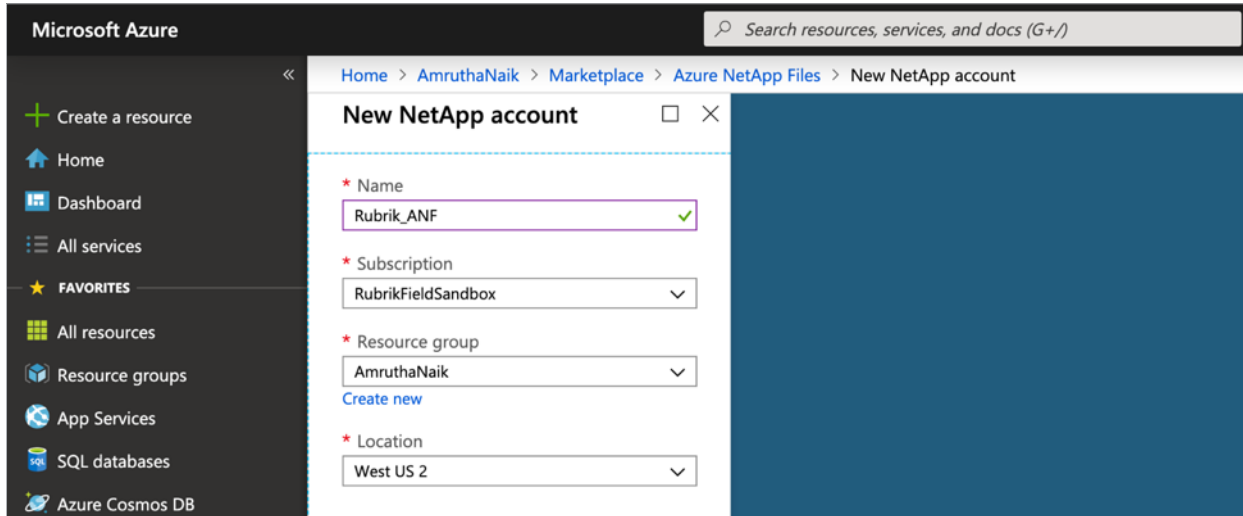
The steps below are handled automatically by Rubrik if applicable when a customer restores NAS data:

1. Based on the source path, determine which partitions need to be accessed whether a single partition or multiple partitions.
2. For each partition, assign it to a node to be executed. This allows restore scalability across many Rubrik nodes.
3. For each partition, create all directories that need to be built, including the ones in the partition boundaries.
4. Prepare writes: open all permissions and create new non-existing files.
5. Parallel writes: write all data in parallel.
6. Create all symbolic links.
7. Set all permissions and other metadata.

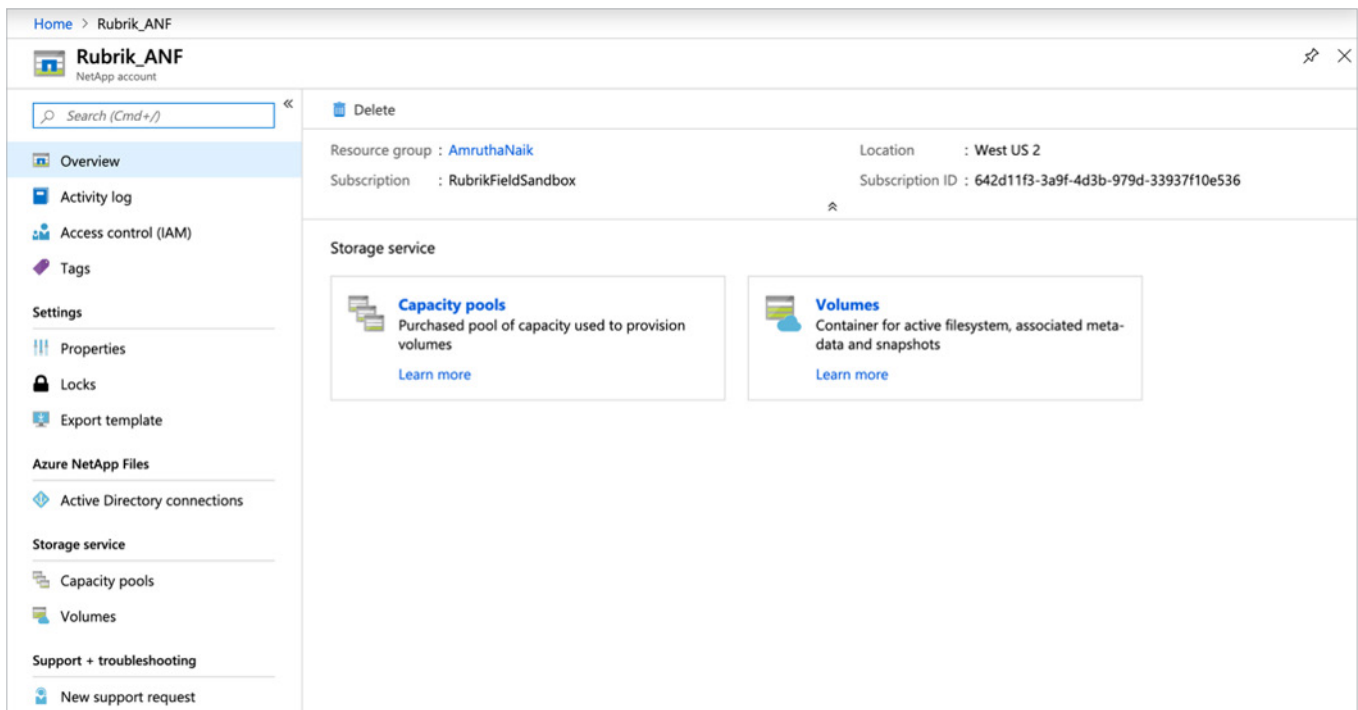
## PROTECT AZURE NETAPP FILES USING RUBRIK

The below section will cover the following two parts:

- Creating NetApp account in Azure portal, creating capacity pool, volumes and NFS/SMB shares
  - Configuring Rubrik to protect Azure NetApp Files
1. In the Azure portal search box, search for **Azure NetApp Files**. After launching the Azure NetApp Files wizard, create an account by pressing the **Add** button:



2. After the NetApp account is created, option for creating capacity pools and volumes is displayed as shown below:



3. Create a capacity pool by clicking on **Capacity pools** and click on **Add Pool**. Enter an appropriate name for the capacity pool, select the service level as per customer requirements and select the size of the pool and click on **OK**:

The screenshot shows the 'New capacity pool' dialog box in the Azure NetApp Files console. The dialog has three fields: 'Name' with the value 'Rubrik\_ANF\_Pool', 'Service level' set to 'Premium', and 'Size (TiB)' set to '4'. An 'OK' button is at the bottom right. In the background, the 'Capacity pools' page is visible, showing a table with columns 'NAME', 'CAPACITY', and 'SERV'. A message states: 'You don't have any capacity pools. Click Add pool to get started.'

4. Once the capacity pool is created, the next step is to create a volume. Click on **Add Volume** and provide the Volume name and quota. Under **Virtual network**, choose your current virtual network or click **Create new** to create a new Azure virtual network (Vnet):

The screenshot shows the 'Create a volume' page in the Azure NetApp Files console. The page has tabs for 'Basics', 'Protocol', 'Tags', and 'Review + create'. The 'Basics' tab is active. The page contains the following fields and values:

- Volume name:** Rubrik\_ANF\_Volume
- Available quota (GiB):** 4096 (4 TiB)
- Quota (GiB):** 100 (100 GiB)
- Virtual network:** (new) AmruthaNaik-vnet (192.168.70.0/25)
- Subnet:** an-anf-netapp (192.168.70.0/28)

At the bottom, there are three buttons: 'Review + create', 'Next: Protocol >>', and 'Download a template for automation'.

Then fill in the following information:

- Enter **myvnet1** as the Vnet name.
- Specify an address space for your setting, for example, **10.7.0.0/16**.
- Enter **myANFsubnet** as the subnet name.
- Specify the subnet address range, for example, **192.168.70.0/25**. Note that you cannot share the dedicated subnet with other resources.
- Select **Microsoft.NetApp/volumes** for subnet delegation.
- Click **OK** to create the Vnet.

**Create virtual network** [X]

\* Name  
AmruthaNaik-vnet

\* Address space  
192.168.70.0/25 ✓  
192.168.70.0 - 192.168.70.127 (128 addresses)

\* Subnet name  
an-anf-netapp ✓

\* Subnet address range ⓘ  
192.168.70.0/28 ✓  
192.168.70.0 - 192.168.70.15 (16 addresses)

\* Subnet delegation  
Microsoft.NetApp/volumes ▼

OK

Back to the **Create a volume** wizard, select the desired file share protocol to access the share, choose the file path, and select the file share protocol version. Click **Review + create**:

Home > Rubrik\_ANF\_Pool (Rubrik\_ANF/Rubrik\_ANF\_Pool) - Volumes > Create a volume

## Create a volume

Basics Protocol Tags Review + create

Configure access to your volume.

**Access**

Protocol type ☒ NFS ☐ SMB

**Configuration**

\* File path  ✓

\* Versions

**Export policy**

Configure the volume's export policy. This can be edited later. [Learn more](#)

<input type="checkbox"/>	INDEX	ALLOWED CLIENTS	ACCESS
<input type="checkbox"/>	1	0.0.0.0/0	Read & Write
		<input type="text"/>	<input type="text"/>

[Download a template for automation](#)

Your newly created volume will now appear in the Volumes section of the **Capacity pool** blade, as shown below:

Home > Rubrik\_ANF - Capacity pools > Rubrik\_ANF\_Pool (Rubrik\_ANF/Rubrik\_ANF\_Pool) - Volumes

### Rubrik\_ANF\_Pool (Rubrik\_ANF/Rubrik\_ANF\_Pool) - Volumes

Capacity pool

Search (Cmd+J)

+ Add volume Refresh

Search volumes

NAME	QUOTA	PROTOCO...	MOUNT PATH	SERVICE LEVEL
Rubrik_ANF_Volume	100 GiB	NFSv3	192.168.70.4/Rubrik-ANF-Volume	Premium

Overview

Activity log

Access control (IAM)

Tags

Settings

Properties

Locks

Export template

Storage service

Volumes

Monitoring

Metrics

Support + troubleshooting

New support request

5. Click the **Volumes** blade, and then select the volume for which you want to mount. Click **Mount instructions** from the selected volume, and then follow the instructions to mount the volume:

... > Rubrik\_ANF - Capacity pools > Rubrik\_ANF\_Pool (Rubrik\_ANF/Rubrik\_ANF\_Pool) - Volumes > Rubrik\_ANF\_Volume (Rubrik\_ANF/Rubrik\_ANF\_Pool/Rubrik\_ANF\_Volume) - Mount instructions

### Rubrik\_ANF\_Volume (Rubrik\_ANF/Rubrik\_ANF\_Pool/Rubrik\_ANF\_Volume) - Mount instructions

Volume

Search (Cmd+J)

- Overview
- Activity log
- Access control (IAM)
- Tags
- Settings
  - Properties
  - Locks
  - Export template
  - Storage service
  - Mount instructions**
  - Export policy
  - Snapshots
- Monitoring
  - Metrics
- Support + troubleshooting
  - New support request

#### Setting up your Azure instance

- Using the Azure Portal, associate your Azure instance with a subnet defined in the same VNet as the volume.
  - The subnet needs a network security group rule that allows traffic on the NFS ports (2049, 111), UDP and TCP.
- Open an SSH client and connect to your Azure instance.
  - [How to Use SSH keys with Windows on Azure.](#)
- Install the NFS client on your Azure instance:
  - On Red Hat Enterprise Linux or SuSE Linux instance:

```
sudo yum install -y nfs-utils
```
  - On an Ubuntu or Debian instance:

```
sudo apt-get install nfs-common
```

#### Mounting your file system

- Create a new directory on your Azure instance:

```
sudo mkdir Rubrik_ANF_Volume
```
- Select a mount target IP address:

192.168.70.4
- Mount your file system using the command/s below:
  - NFSv3

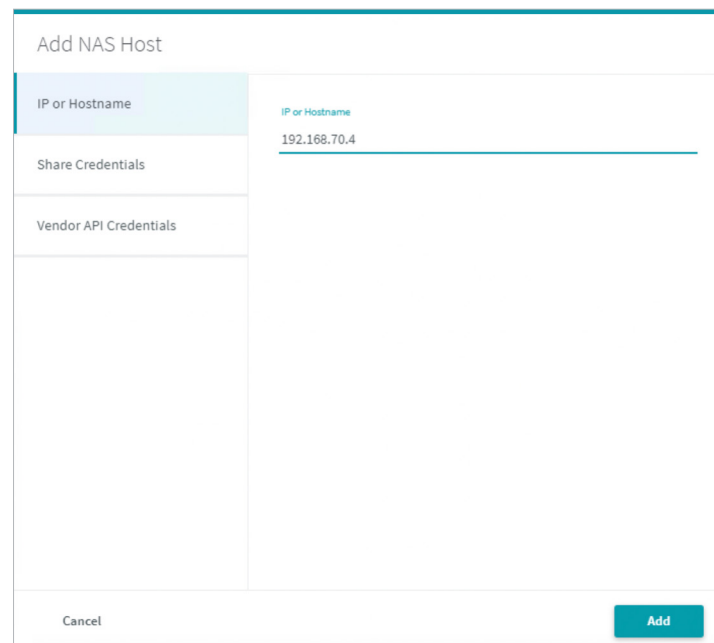
```
sudo mount -t nfs -o rw,hard,rsz=65536,wsz=65536,vers=3,tcp 192.168.70.4/Rubrik-ANF-Volume Rubrik_ANF_Volume
```

The screenshot below shows an example of mounting an Azure NetApp Files volume on a CentOS 7 client:

```
[amrutha@CentOS-7 ~]$ sudo mkdir Rubrik_ANF_Volume
[amrutha@CentOS-7 ~]$ sudo mount -t nfs -o rw,hard,rsz=65536,wsz=65536,vers=3,tcp 192.168.70.4:/Rubrik-ANF-Volume Rubrik_ANF_Volume
[amrutha@CentOS-7 ~]$ mount
sysfs on /sys type sysfs (rw,nosuid,nodev,noexec,relatime,seclabel)
proc on /proc type proc (rw,nosuid,nodev,noexec,relatime)
devtmpfs on /dev type devtmpfs (rw,nosuid,seclabel,size=4066508k,nr_inodes=1014627,mode=755)
securityfs on /sys/kernel/security type securityfs (rw,nosuid,nodev,noexec,relatime)
tmpfs on /dev/shm type tmpfs (rw,nosuid,nodev,seclabel)
devpts on /dev/pts type devpts (rw,nosuid,noexec,relatime,seclabel,gid=5,mode=620,ptmxmode=000)
tmpfs on /run type tmpfs (rw,nosuid,nodev,seclabel,mode=755)
tmpfs on /sys/fs/cgroup type tmpfs (ro,nosuid,nodev,noexec,seclabel,mode=755)
cgroup on /sys/fs/cgroup/systemd type cgroup (rw,nosuid,nodev,noexec,relatime,seclabel,xattr,release_agent=/usr/lib/systemd/systemd-cgroups-agent,name=systemd)
pstore on /sys/fs/pstore type pstore (rw,nosuid,nodev,noexec,relatime)
cgroup on /sys/fs/cgroup/perf_event type cgroup (rw,nosuid,nodev,noexec,relatime,seclabel,perf_event)
cgroup on /sys/fs/cgroup/hugetlb type cgroup (rw,nosuid,nodev,noexec,relatime,seclabel,hugetlb)
cgroup on /sys/fs/cgroup/net_cls,net_prio type cgroup (rw,nosuid,nodev,noexec,relatime,seclabel,net_prio,net_cls)
cgroup on /sys/fs/cgroup/pids type cgroup (rw,nosuid,nodev,noexec,relatime,seclabel,pids)
cgroup on /sys/fs/cgroup/cpuset type cgroup (rw,nosuid,nodev,noexec,relatime,seclabel,cpuset)
cgroup on /sys/fs/cgroup/cpu,cpuacct type cgroup (rw,nosuid,nodev,noexec,relatime,seclabel,cpuacct,cpu)
cgroup on /sys/fs/cgroup/memory type cgroup (rw,nosuid,nodev,noexec,relatime,seclabel,memory)
cgroup on /sys/fs/cgroup/devices type cgroup (rw,nosuid,nodev,noexec,relatime,seclabel,devices)
cgroup on /sys/fs/cgroup/bklcio type cgroup (rw,nosuid,nodev,noexec,relatime,seclabel,bklcio)
cgroup on /sys/fs/cgroup/freezer type cgroup (rw,nosuid,nodev,noexec,relatime,seclabel,freezer)
configfs on /sys/kernel/config type configfs (rw,relatime)
/dev/sda2 on / type xfs (rw,relatime,seclabel,attr2,inode64,noquota)
selinuxfs on /sys/fs/selinux type selinuxfs (rw,relatime)
systemd-1 on /proc/sys/fs/binfmt_misc type autofs (rw,relatime,fd=30,pgpr=1,timeout=0,minproto=5,maxproto=5,direct,pipe_ino=10913)
mqueue on /dev/mqueue type mqueue (rw,relatime,seclabel)
hugetlbfs on /dev/hugepages type hugetlbfs (rw,relatime,seclabel)
debugfs on /sys/kernel/debug type debugfs (rw,relatime)
/dev/sda1 on /boot type xfs (rw,relatime,seclabel,attr2,inode64,noquota)
/dev/sdb1 on /mnt/resource type ext4 (rw,relatime,seclabel,data=ordered)
tmpfs on /run/udev/1000-udev-limits type tmpfs (rw,nosuid,nodev,noexec,relatime,seclabel,size=8157272k,mode=700,uid=1000,gid=1000)
192.168.70.4:/Rubrik-ANF-Volume on /home/amrutha/Rubrik_ANF_Volume type nfs (rw,relatime,vers=3,rsz=65536,wsz=65536,namlen=255,hard,proto=tcp,timeo=600,retrans=2,sec=sys,mountaddr=192.168.70.4,mountvers=3,mountport=635,mountproto=tcp,local_lock=none,addr=192.168.70.4)
[amrutha@CentOS-7 ~]$
```

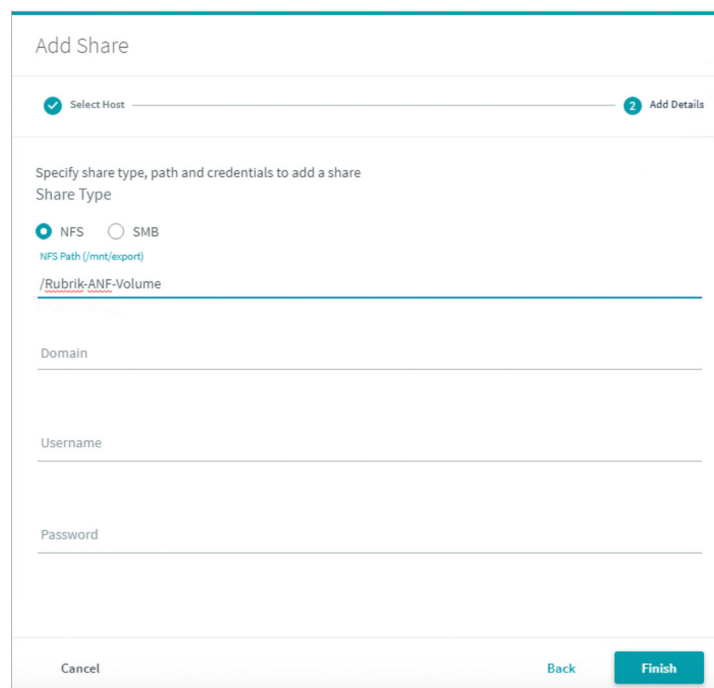


6. On Rubrik, click on **NAS Shares** under **Servers & Apps**. Click on **Add NAS Host** and provide the IP address or hostname of the NAS host retrieved from the above step (ANF mount target IP address). Click **Add**:



The screenshot shows the 'Add NAS Host' dialog box. It has a title bar 'Add NAS Host'. Below the title bar, there are three input fields: 'IP or Hostname' (with the value '192.168.70.4'), 'Share Credentials', and 'Vendor API Credentials'. At the bottom, there are 'Cancel' and 'Add' buttons.

7. Add the ANF share details. Select the protocol – NFS/SMB. Add the ANF mount path and credentials (if any). Click **Finish**:



The screenshot shows the 'Add Share' dialog box. It has a title bar 'Add Share'. Below the title bar, there is a progress bar with two steps: 'Select Host' (completed) and 'Add Details' (current step). The 'Add Details' section includes 'Specify share type, path and credentials to add a share'. Under 'Share Type', 'NFS' is selected. The 'NFS Path (/mnt/export)' field contains '/Rubrik-ANF-Volume'. There are also fields for 'Domain', 'Username', and 'Password'. At the bottom, there are 'Cancel', 'Back', and 'Finish' buttons.

8. In the **Add Fileset** page, provide a Fileset Name, select the protocol and choose the files in the share that have to be protected. In this example, we have chosen to select all the files in the fileset. Click **Add**:

### Add Fileset

Fileset Name

Rubrik\_ANF\_Fileset

Share Type

☒ NFS ☐ SMB

Rules ⓘ

Use \*\* to include all files

Include (/usr/local, \*.pdf)

\*\*

Exclude (/usr/local/temp, \*.mov, \*.mp3, \*.mp4)

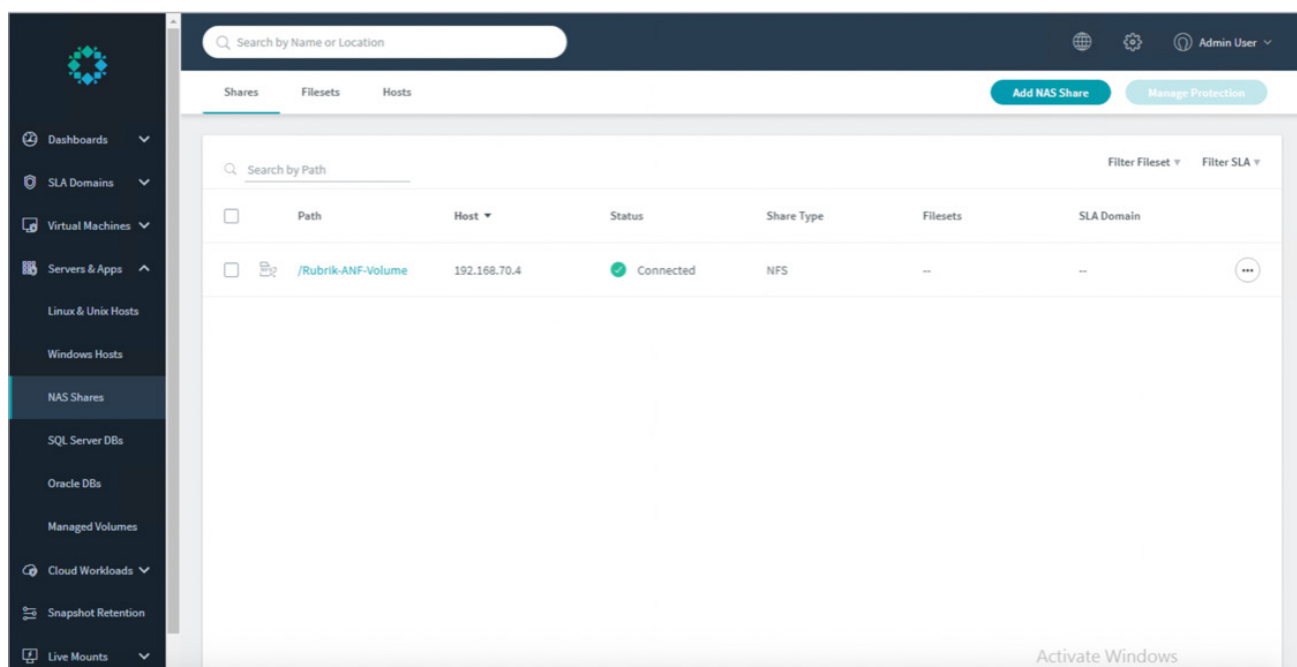
Do Not Exclude (/company, \*.mp4)

☐ Enable Backup of Hidden Folders

Cancel

Add

The newly added share is now listed in the Rubrik UI:



9. In the below step, we are adding an Azure container as an archival location for long-term retention. Provide Azure storage account details such as storage account name, access key, container name and click **Add**:

The screenshot shows the 'Add Archival Location' form. It includes the following fields and values:

- Archival Type:** Azure
- Storage Account Name:** anfestamruthanaik
- Access Key:** A masked field with 20 asterisks.
- Container:** anfarchival
- Archival Location Name:** Azure:anfarchival
- Instance Type:** Azure Default
- RSA Key:** A masked field with the text '-----BEGIN RSA PRIVATE KEY-----'.

At the bottom, there are 'Cancel' and 'Add' buttons.

10. Create an SLA Domain policy and set the backup frequency and retention period:

The screenshot shows the 'Create SLA Domain' form. It includes the following fields and values:

- SLA Domain Name:** ANF\_Rubrik
- Advanced Configuration:** A toggle switch that is currently turned off.
- Service Level Agreement:** Choose how often we take snapshots and the length of time we keep them.
- Take Snapshots:** A table with the following values:

Take Snapshots:	Keep Snapshots:
Every (Hours)	For (Days)
Every (Days)	For (Days)
1	7
Every (Months)	For (Months)
1	7
Every (Years)	For (Years)
1	7
- Local retention set to 7 years .**
- Snapshot Window:** Take snapshots from: : to : .

At the bottom, there are 'Cancel' and 'Create' buttons.

Before clicking **Create**, click the **Remote settings** button located at the bottom right-hand corner. Toggle on Archival, select the desired archival location and choose how long the backup data should be retained locally on the Rubrik cluster by using the **Retention On Brik** slider. Click **Create**:

Create SLA Domain

Remote Storage Configuration

Retention On Brik

07 years7 years

Archival

Azure:anfarchival

Enable Instant Archive

Archival starts immediately, and is retained on the archival location for 7 years .

Replication

A replication target has not been set up yet. Please [add a replication target](#) to configure retention.

SLA Domain Creation

Cancel

Create

Once your new SLA Domain policy is created, you can see a summary of its settings by selecting it in the list of existing SLA Domains:

ANF\_RubrikLocal

SLA Domain Policy

Take:

Every 1 day

Every 1 month on last day of the month

Every 1 year on last day of the year starting in January

Retain:

for 7 days

for 7 months

for 7 years

Snapshot Window

Not Configured

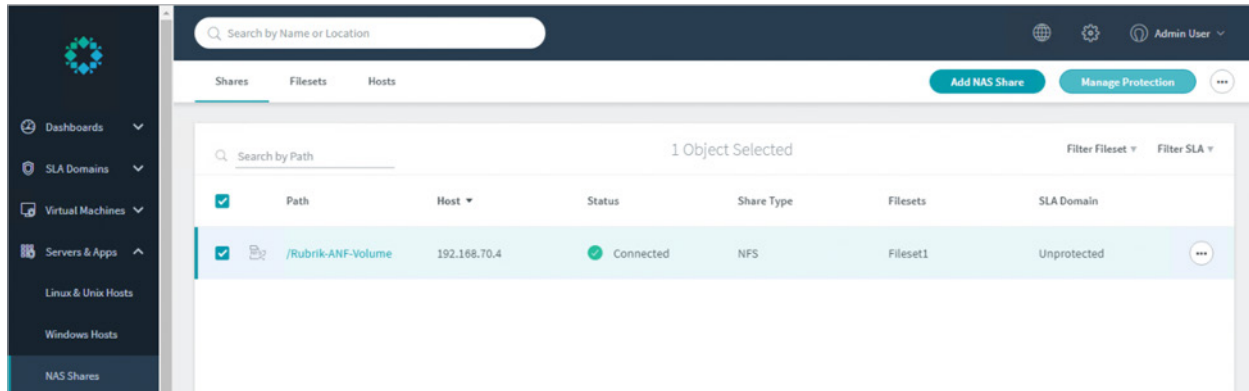
Replication Retention Policy

Not Configured

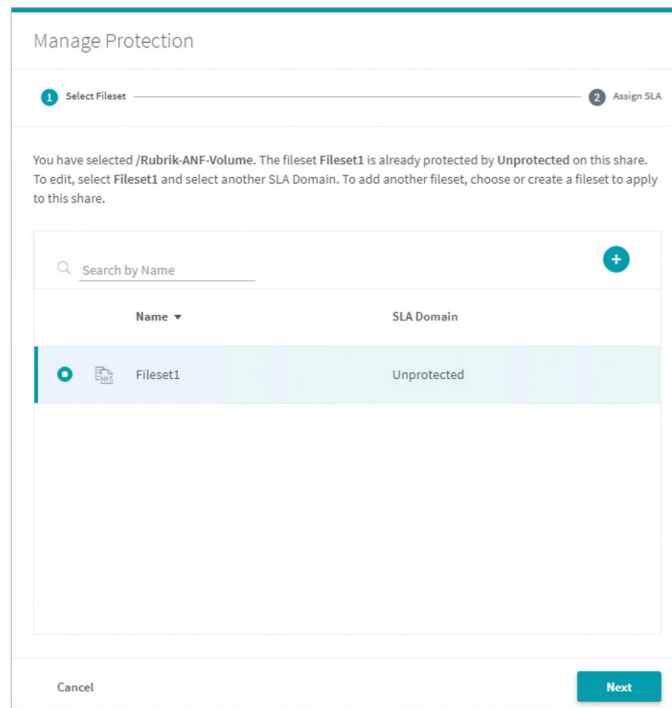
Archival Policy

Snapshots are stored on rubrik-anf-cc for 7 years .  
At the same time, data is instantly copied to Azure:anfarchival and will be stored there for 7 years .

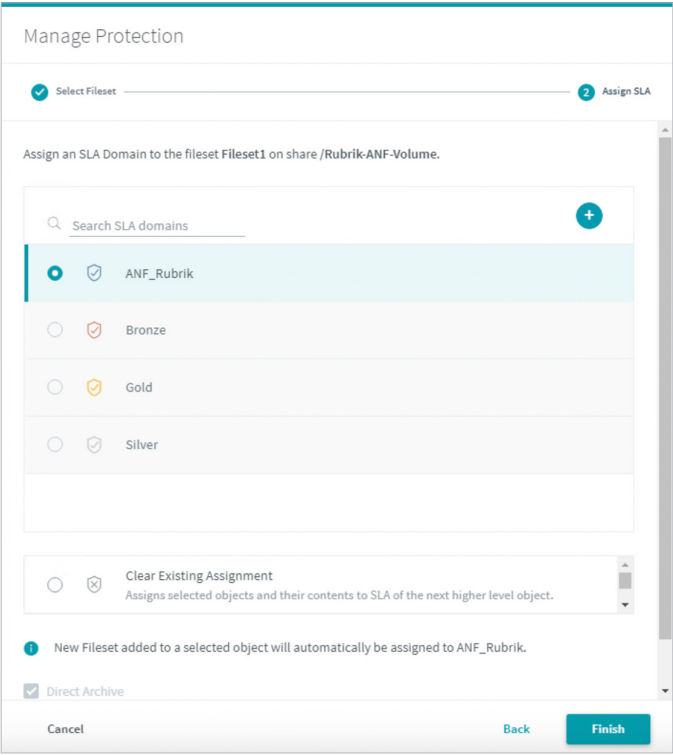
11. In the next step, assign the SLA Domain to the NAS fileset. Select the NAS fileset and click on **Manage Protection**:



Select the Fileset that needs to be protected and click **Next**:



Assign the SLA domain policy created in the previous step and click **Finish**:



The Rubrik UI now shows which filesets are protected for this ANF volume, and by which SLA Domain policy:

Shares

Filesets

Hosts

Add NAS Share

Manage Protection

Search by Path

Filter Fileset

Filter SLA

<div></div>	Path	Host	Status	Share Type	Filesets	SLA Domain	
<div><div></div><div><div></div><div></div></div></div>	<div>/Rubrik-ANF-Volume</div>	192.168.70.4	<div><div></div><div>Connected</div></div>	NFS	Fileset1	ANF_Rubrik	<div></div>

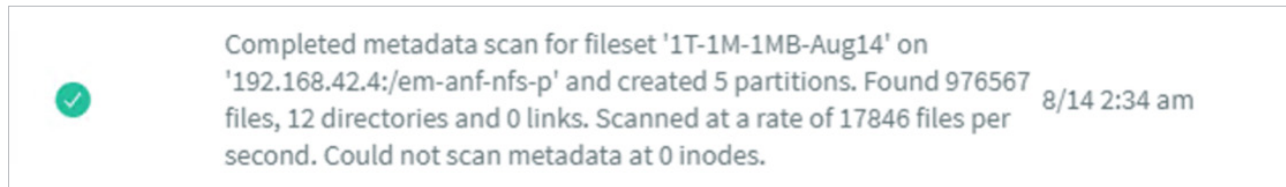
While Rubrik Cloud Edition can scale linearly to increase performance, compute and storage, these tests were performed with the default Cloud Edition installation.

## BACKUP PERFORMANCE

Let's take a look at the backup performance, starting with the first full backup, and followed by incremental backup.

### FIRST FULL BACKUP

During the scan phase of the backup job, it took around 1 minute to iterate through the file system and scan the files, scanning at a rate of ~18,000 files per second. This scan performance was found to be consistent across all 3 tiers of Azure NetApp Files.



On ingest the fileset was automatically split up into partitions and ingested in parallel across the cluster at a rate of ~90MB/s per node – close to the maximum throughput achievable by a DS3v2 instance.

rksupport@VRAZ76A5545D9: ~									
	23.8MB	47.7MB	71.5MB	95.4MB	119MB		23.8MB	47.7MB	71.5MB
VRAZ76A5545D9	=> 192.168.42.4			347KB 349KB 331KB	VRAZB4A64A6E0	=> 192.168.42.4			319KB 350KB 345KB
VRAZ76A5545D9	=> 192.168.42.42			88.2MB 88.3MB 76.2MB	VRAZB4A64A6E0	=> 192.168.42.40			75.4MB 77.9MB 76.0MB
VRAZ76A5545D9	=> 192.168.42.43			25.1KB 27.6KB 33.5KB	VRAZB4A64A6E0	=> 192.168.42.43			79.6KB 44.1KB 44.9KB
VRAZ76A5545D9	=> 192.168.42.41			68.1KB 52.5KB 75.8KB	VRAZB4A64A6E0	=> 192.168.42.42			47.7MB 20.8KB 35.3KB
VRAZ76A5545D9	=> 13.66.232.233			19.3KB 37.6KB 34.2KB	VRAZB4A64A6E0	=> 13.66.232.233			33.2MB 27.2KB 36.8KB
VRAZ76A5545D9	=> 168.63.129.16			5.63KB 34.6KB 46.9KB	VRAZB4A64A6E0	=> 168.63.129.16			23.0KB 33.3KB 33.8KB
VRAZ76A5545D9	=> 192.168.42.23			28.9KB 24.3KB 33.7KB	VRAZB4A64A6E0	=> 192.168.42.21			48.4KB 25.8KB 25.1KB
VRAZ76A5545D9	=> 192.168.42.22			61.1KB 39.8KB 68.7KB	VRAZB4A64A6E0	=> ec2-52-208-65-217.compute-1.amazonaws.com			24.0MB 20.8KB 20.5KB
VRAZ76A5545D9	=> 192.168.42.21			718B 425B 490B	VRAZB4A64A6E0	=> 192.168.42.5			732B 427B 407B
VRAZ76A5545D9	=> 169.254.169.254			5.28KB 3.15KB 3.62KB	VRAZB4A64A6E0	=> 192.168.42.5			5.26KB 3.15KB 3.08KB
VRAZ76A5545D9	=> 192.168.42.21			1.99KB 1.20KB 1.44KB	VRAZB4A64A6E0	=> 192.168.42.5			2.00KB 1.36KB 1.26KB
VRAZ76A5545D9	=> 192.168.42.21			2.54KB 1.52KB 1.79KB	VRAZB4A64A6E0	=> 192.168.42.21			2.54KB 1.63KB 1.52KB
VRAZ76A5545D9	=> 192.168.42.21			1.38KB 1.68KB 1.80KB	VRAZB4A64A6E0	=> 192.168.42.21			120B 72B 99B
VRAZ76A5545D9	=> 192.168.42.21			184B 178B 332B	VRAZB4A64A6E0	=> 192.168.42.21			238B 145B 194B
VRAZ76A5545D9	=> 192.168.42.21			1.58KB 752B 886B	VRAZB4A64A6E0	=> 192.168.42.21			0B 45B 32B
VRAZ76A5545D9	=> 192.168.42.21			634B 481B 415B	VRAZB4A64A6E0	=> 192.168.42.21			0B 96B 69B
VRAZ76A5545D9	=> 192.168.42.21			83B 54B 84B	VRAZB4A64A6E0	=> 192.168.42.21			392B 78B 56B
VRAZ76A5545D9	=> 192.168.42.21			174B 104B 166B	VRAZB4A64A6E0	=> 192.168.42.21			234B 47B 33B
VRAZ76A5545D9	=> 192.168.42.21			0B 67B 123B	VRAZB4A64A6E0	=> 192.168.42.21			0B 0B 0B
VRAZ76A5545D9	=> 192.168.42.21			0B 52B 38B	VRAZB4A64A6E0	=> 192.168.42.21			0B 8B 6B
VRAZ76A5545D9	=> 192.168.42.21			0B 0B 14B	VRAZB4A64A6E0	=> 192.168.42.21			0B 0B 0B
VRAZ76A5545D9	=> 192.168.42.21			0B 0B 38B	VRAZB4A64A6E0	=> 192.168.42.21			0B 8B 6B
TX:	cum: 13.7MB peak: 555KB			rates: 426KB 443KB 437KB	TX:	cum: 6.07MB peak: 495KB			rates: 476KB 449KB 444KB
XX:	2.33GB 84.9MB			88.2MB 88.3MB 76.2MB	XX:	1.02GB 84.9MB			75.4MB 77.9MB 76.0MB
TOTAL:	2.40GB 85.4MB			88.2MB 88.3MB 76.2MB	TOTAL:	1.05GB 84.9MB			76.0MB 78.5MB 77.3MB
rksupport@VRAZ08574A99F: ~									
	23.8MB	47.7MB	71.5MB	95.4MB	119MB		23.8MB	47.7MB	71.5MB
VRAZ08574A99F	=> 192.168.42.4			332KB 338KB 334KB	VRAZB4A64A6E0	=> 192.168.42.4			382KB 329KB 320KB
VRAZ08574A99F	=> 192.168.42.40			81.9MB 84.3MB 83.3MB	VRAZB4A64A6E0	=> 192.168.42.40			72.9MB 79.7MB 79.7MB
VRAZ08574A99F	=> 192.168.42.43			53.2KB 52.2KB 54.7KB	VRAZB4A64A6E0	=> 192.168.42.43			7.04KB 34.9KB 34.9KB
VRAZ08574A99F	=> 192.168.42.41			18.4KB 27.5KB 26.8KB	VRAZB4A64A6E0	=> 192.168.42.41			17.5KB 37.7KB 37.7KB
VRAZ08574A99F	=> 192.168.42.41			68.2KB 34.1KB 38.6KB	VRAZB4A64A6E0	=> 192.168.42.41			23.0KB 33.3KB 33.3KB
VRAZ08574A99F	=> 192.168.42.41			6.15KB 22.1KB 20.5KB	VRAZB4A64A6E0	=> 192.168.42.41			33.1KB 27.2KB 27.2KB
VRAZ08574A99F	=> 13.66.232.233			24.4KB 28.8KB 22.3KB	VRAZB4A64A6E0	=> 192.168.42.41			6.15KB 22.1KB 22.1KB
VRAZ08574A99F	=> 168.63.129.16			48.8KB 25.9KB 27.8KB	VRAZB4A64A6E0	=> 13.66.232.233			66.4KB 34.0KB 34.0KB
VRAZ08574A99F	=> 192.168.42.21			718B 427B 474B	VRAZB4A64A6E0	=> 168.63.129.16			0B 426B 426B
VRAZ08574A99F	=> 192.168.42.21			5.28KB 3.15KB 3.58KB	VRAZB4A64A6E0	=> 192.168.42.21			0B 3.15KB 3.15KB
VRAZ08574A99F	=> 168.63.129.16			1.99KB 1.20KB 1.33KB	VRAZB4A64A6E0	=> 192.168.42.21			1.54KB 1.26KB 1.26KB
VRAZ08574A99F	=> 192.168.42.21			2.54KB 1.52KB 1.69KB	VRAZB4A64A6E0	=> 192.168.42.21			338B 1.52KB 1.52KB
VRAZ08574A99F	=> 192.168.42.21			0B 454B 378B	VRAZB4A64A6E0	=> 192.168.42.21			0B 108B 108B
VRAZ08574A99F	=> 192.168.42.21			0B 146B 121B	VRAZB4A64A6E0	=> 192.168.42.21			0B 196B 196B
VRAZ08574A99F	=> 192.168.42.21			83B 57B 103B	VRAZB4A64A6E0	=> 192.168.42.21			0B 196B 196B
VRAZ08574A99F	=> 192.168.42.21			174B 117B 203B	VRAZB4A64A6E0	=> 192.168.42.21			0B 196B 196B
TX:	cum: 5.20MB peak: 502KB			rates: 473KB 447KB 444KB	TX:	cum: 4.12MB peak: 489KB			rates: 340KB 421KB 421KB
XX:	8.90GB 89.0MB			82.2MB 84.3MB 83.4MB	XX:	2.02GB 87.1MB			73.0MB 79.0MB 79.0MB
TOTAL:	8.98GB 90.1MB			82.2MB 84.3MB 83.4MB	TOTAL:	88.2MB 87.1MB			73.3MB 88.2MB 88.2MB

The first full backup of this dataset completed in 2hrs 59 mins, with an average end to end throughput of 93.6MB/s.

Activity Detail			
Event Type			
Backup			
Object Name			
1T-1M-1MB-Aug14			
Location			
192.168.42.4:/em-anf-nfs-p			
Status			
Success			
Start Time			
Aug 14, 2019 02:33:07 AM			
Duration			
2 hrs 59 mins 50 secs			
SLA			
ntap_nasda			
Data Transferred			
929.8 GB			
Logical Size			
1163.93 GB			
	Status	Activity	Date
	✓	Completed backup of Fileset '1T-1M-1MB-Aug14' from '192.168.42.4:/em-anf-nfs-p'	8/14 5:32 am
	✓	Successfully retrieved 976567 files in '1T-1M-1MB-Aug14' from '192.168.42.4:/em-anf-nfs-p'. A total of 998 GB were fetched in 2 hours, 57 minutes, 46 seconds. The average transfer rate was 93.6 MBps. Could not retrieve 0 files. The fetch phase took approximately 49 minutes, 53 seconds. The copy phase took approximately 2 hours, 7 minutes, 20 seconds. The verification phase took approximately 32 seconds.	8/14 5:32 am
	✓	Completed metadata scan for fileset '1T-1M-1MB-Aug14' on '192.168.42.4:/em-anf-nfs-p' and created 5 partitions. Found 976567 files, 12 directories and 0 links. Scanned at a rate of 17846 files per second. Could not scan metadata at 0 inodes.	8/14 2:34 am
	✓	Queued backup of '1T-1M-1MB-Aug14'	8/14 2:33 am
Download Server Logs			OK

## INCREMENTAL BACKUP

Another backup was then run on the same dataset as part of the defined SLA protecting the fileset.

Scan performance was again consistent, this time at just over ~21,000 files per second, with the backup being completed in 4 minutes 53 seconds.

Activity Detail			
Event Type			
Backup			
Object Name			
1T-1M-1MB-Aug14			
Location			
192.168.42.4:/em-anf-nfs-p			
Status			
Success			
Start Time			
Aug 23, 2019 02:34:17 AM			
Duration			
4 mins 53 secs			
SLA			
ntap_nasda			
Logical Size			
1163.93 GB			
	Status	Activity	Date
	✓	Completed backup of Fileset '1T-1M-1MB-Aug14' from '192.168.42.4:/em-anf-nfs-p'	8/23 2:39 am
	✓	Successfully retrieved 976567 files in '1T-1M-1MB-Aug14' from '192.168.42.4:/em-anf-nfs-p'. A total of 0 B were fetched in 3 minutes, 18 seconds. The average transfer rate was 0 Bps. Could not retrieve 0 files. The fetch phase took approximately 36 seconds. The copy phase took approximately 2 minutes, 16 seconds. The verification phase took approximately 24 seconds.	8/23 2:39 am
	✓	Completed metadata scan for fileset '1T-1M-1MB-Aug14' on '192.168.42.4:/em-anf-nfs-p' and created 5 partitions. Found 976567 files, 12 directories and 0 links. Scanned at a rate of 21700 files per second. Could not scan metadata at 0 inodes.	8/23 2:35 am
	✓	Queued backup of '1T-1M-1MB-Aug14'	8/23 2:34 am
	🕒	Scheduled backup of '1T-1M-1MB-Aug14' on '2019-08-23T02:34:15.420-07:00[America/Los_Angeles]'	8/22 2:39 pm
Download Server Logs			OK



## RESTORE PERFORMANCE

Restoration performance of this fileset was consistent with the results from the backup. A full restore of the share was performed with data being automatically multi-streamed back to the source NFS device by the Rubrik cluster at 90.29MB/s. A full 1TB restore completed in 3 hours 5 minutes.

Activity Detail

Event Type  
Recovery

Object Name  
4T-4M-1MB-July22

Location  
192.168.42.4;/em-anf-nfs-p

User Name  
admin

Status  
Success

Start Time  
Sep 11, 2019 09:18:55 AM

Duration  
3 hrs 4 mins 22 secs

Data Transferred  
929.82 GB

Throughput  
688.84 Mbs

Status	Activity	Date
✓	Successful restore of '/' to '192.168.42.4;/em-anf-nfs-p' based on snapshot taken at 'Sep 11, 2019 14:20:17 UTC'. Successfully restored 998 GB of data in 976566 files, 0 links, and 17 directories. Failed to restore 0 files and 0 directories.	9/11 12:23 pm
✓	Completed writing data to '192.168.42.4;/em-anf-nfs-p' from node 'VRAZ76A5545D9' for fileset '4T-4M-1MB-July22' in 3 hours, 4 minutes, 17 seconds. Restored at a rate of 88 files per second and 90.29 MB/s.	9/11 12:23 pm

Download Server Logs

OK

## CONCLUSION

Azure NetApp Files was introduced jointly by Microsoft and NetApp to allow organizations to migrate existing enterprise-grade applications that use high performance file shares from on-premises to the cloud without requiring major changes, or that need high database performance such as Oracle or SAP in the cloud for instance. Running such applications in the cloud requires to have a solid data protection strategy in which backup must be fast, reliable, secure and easy to use.

Rubrik chose from day 1 to backup file shares using industry standard protocols such as NFS and SMB to overcome many of the limitations that usually come with NDMP-based solutions. Because the Azure NetApp Files service uses the same protocols to access data, it was straight forward for Rubrik to support the new data service. With the help of Rubrik CDM Cloud Edition, a virtual version of Rubrik CDM specifically built for Microsoft Azure, backups of data born in the cloud stay in the cloud.

As described throughout this document, Azure NetApp Files can be backed up by Rubrik with the same policy-based approach as usual, with extreme ease of use via the HTML5 web interface and sustainable performance for both backup and recovery.

To learn more about Rubrik's general approach of NAS backup, read our [Protecting NAS at Scale with Rubrik](#) white paper.

## ABOUT THE AUTHORS

Amrutha Naik is a Senior Technical Alliances Manager at Rubrik. Prior to joining Rubrik, she worked as a Senior Product Manager at Tegile and also spent some years in Technical Marketing at EMC & NetApp gaining experience in virtualization, storage and data protection. Outside of work she enjoys painting, gardening, playing violin and using her husband as a lab rat for culinary experiments.

Ed Morgan is a Technical Product Manager at Rubrik. Before joining Rubrik he has worked in various consultancy and pre-sales roles for the past decade, with a particular focus on enterprise storage, virtualisation, and cloud, and spent 5 years serving in the British Army. Outside of work he enjoys rock music, CrossFit, endurance events, and writing about himself in the 3rd person. You can follow him on Twitter [@mo6020](#)

Pierre-François Guglielmi is the Technical Manager of Partner Solutions at Rubrik with many years of experience in virtualization and data protection. Pierre-François lives near Paris, France and enjoys CrossFit and playing music in his free time. You can follow him on Twitter [@pfguglielmi](#).



### Global HQ

1001 Page Mill Rd., Building 2  
Palo Alto, CA 94304  
United States

1-844-4RUBRIK  
[inquiries@rubrik.com](mailto:inquiries@rubrik.com)  
[www.rubrik.com](http://www.rubrik.com)

Rubrik, the Multi-Cloud Data Control™ Company, enables enterprises to maximize value from data that is increasingly fragmented across data centers and clouds. Rubrik delivers a single, policy-driven platform for data recovery, governance, compliance, and cloud mobility. For more information, visit [www.rubrik.com](http://www.rubrik.com) and follow [@rubrikInc](#) on Twitter. © 2020 Rubrik. Rubrik is a registered trademark of Rubrik, Inc. Other marks may be trademarks of their respective owners.