

TECHNICAL WHITE PAPER

# FlashArray File Services Protection Using Veeam Backup & Replication

A best practices guide

# Contents

<b>Introduction .....</b>	<b>3</b>
<b>How to Use This Guide .....</b>	<b>3</b>
<b>Solution Architecture.....</b>	<b>3</b>
Components .....	4
System Recommendations for Veeam Components .....	5
Scaling Considerations.....	5
<b>Understanding Backup and Recovery Performance with Veeam and FlashArray File Services .....</b>	<b>5</b>
File Access Protocol.....	6
Data Profile.....	6
Concurrent Tasks .....	7
Scaling Recommendations for Veeam Backup Proxies .....	7
<b>Basic Configuration .....</b>	<b>8</b>
Deploying File Share Backup Proxies.....	8
Adding Network Shares .....	10
Creating Backup Jobs .....	14
<b>Configuration Best Practices .....</b>	<b>20</b>
File Systems and Managed Directories .....	20
Veeam Backup Proxies.....	22
Veeam Network Shares.....	22
Veeam Repositories .....	22
Performance .....	22
<b>SMB-specific Best Practices .....</b>	<b>23</b>
Manage SMB Backup Access .....	23
<b>NFS Best Practices .....</b>	<b>25</b>
Restrict NFS export access .....	25
<b>Backup Best Practices .....</b>	<b>25</b>
Protect Open Files.....	25
FlashArray File Services Snapshot Behavior .....	25
Use Folder-Level ACL Handling .....	27
<b>General Snapshot Best Practices .....</b>	<b>27</b>
Create Snapshots on File Share Managed Directories.....	27
Use a Consistent Snapshot Name .....	27
Use Scripts to Create Snapshots for Backups .....	27
<b>Recovery Best Practices .....</b>	<b>28</b>
Recreate Managed Directories Before Restoring Data .....	28
Restore Large File Shares First.....	28
Avoid Overwriting Large Data Sets.....	28
<b>Conclusion .....</b>	<b>28</b>
<b>Additional Resources .....</b>	<b>29</b>
<b>About the Author .....</b>	<b>30</b>

## Introduction

Pure Storage® FlashArray™ File Services brings the reliability, data reduction, and simplicity of Purity//FA to network-attached storage (NAS). Where traditional NAS is often complicated to deploy, difficult to manage, and painful to refresh, unified all-flash storage from FlashArray delivers the same experience that's been delighting customers for years. Thanks to Evergreen Storage™, FlashArray stays forever young, so you'll never need to plan another disruptive migration. FlashArray forms the foundation of a truly Modern Data Experience™.

Veeam Backup & Replication delivers scalable, high-performance protection of FlashArray file data. Using a parallel backup and restore model that provides easy deployment and performance scaling, Veeam provides everything you need to back up and restore FlashArray file data.

---

## How to Use This Guide

We wrote this solution overview and best practices guide for backup administrators, storage administrators, and others. Our goal is to help you understand and then implement Veeam Backup & Replication (VBR) to protect and recover data on FlashArray File Services. The guide covers the solution architecture, key performance factors for protecting FlashArray File Services, basic configuration for simple environments and POCs, and best practices for achieving optimal performance and efficiency.

This guide assumes a working knowledge of the concepts and interfaces used with Veeam Backup & Replication and FlashArray File Services. You can learn more about Veeam interfaces from the [Veeam Help Center](#). Refer to the FlashArray user guide in your FlashArray web console for more information about managing FlashArray File Services.

## Solution Architecture

VBR uses a scale-out approach to managing NAS data. VBR breaks backup and restore sessions into tasks that read or write a subset of the files, which can be distributed across multiple systems. This parallelism increases backup and restore speeds, and it significantly simplifies configuration compared to single-threaded, single-reader solutions.



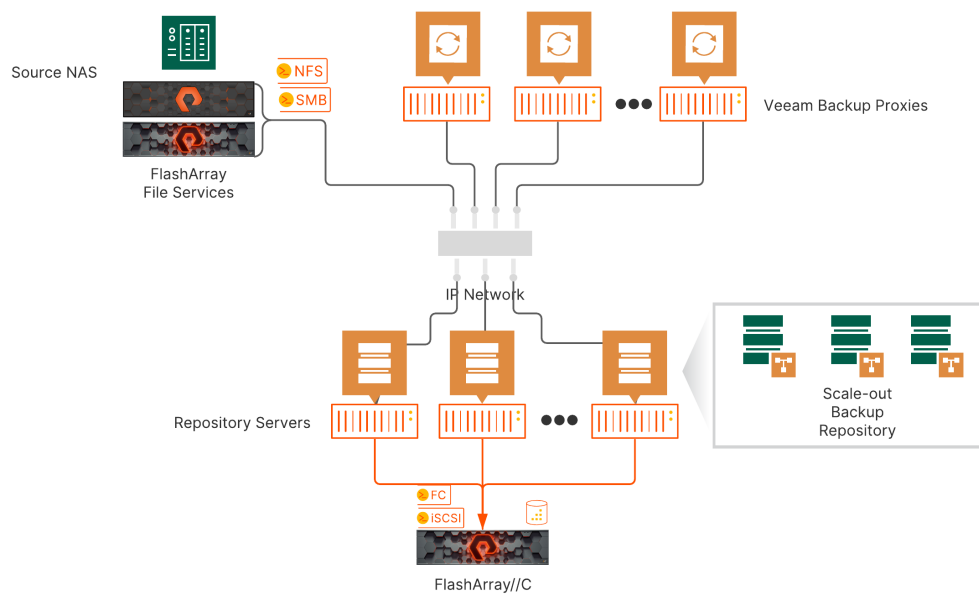


Figure 1. Solution architecture

During backup or restore sessions, a pool of Veeam backup proxies performs read and write operations against the NAS network shares using Server Message Block (SMB) or Network File Share (NFS) protocols, communicating with the scale-out backup repository (SOBR) for data storage. VBR can throttle the data traffic from the backup proxies to limit the load amount put on the FlashArray.

During backup and recovery sessions, VBR breaks the work into a set of tasks. The proxies dynamically divide the tasks.

## Components

**Pure Storage FlashArray File Services:** FlashArray File Services is the primary data tier, hosting unstructured data for users and applications to share. FlashArray File Services arranges the data in file systems and managed directories, which have policies applied to them. Users and applications access their data using the SMB or NFS protocol. VBR tracks the individual SMB or NFS file shares in its inventory of backup sources, and backup jobs define the frequency, target, and retention period for protecting data within the shares.

**Veeam Backup & Replication:** We developed this solution architecture using VBR V10. However, you can deploy it on VBR V11 with no changes. Consult the [Veeam Help Center](#) for details on VBR product changes between versions.

**Veeam Backup Proxies:** Veeam backup proxies are the data movers, reading and writing FlashArray data over SMB and NFS. During backups, the proxies compress the data, based on backup job settings, and send it to the repository. They also collect metadata about files and directories to track changes between backups.

Veeam designates proxies for particular workloads. In VBR V10, proxies can service VMware, Hyper-V, or in this case file shares. A single managed server can run multiple proxy types, so for example a server can back up both VMware VMs and file shares, but each proxy type on the same server is managed as a separate proxy.



**Backup Repository:** The Veeam Backup Repository provides the storage for backed up data and can be block- or file-based storage. Veeam supports a broad list of storage products, including FlashArray//C. For more information on using FlashArray//C as a repository platform, see the [Enhancing Veeam with FlashArray//C](#) white paper.

**Network Shares:** VBR tracks source file data in its inventory as network shares, which map to FlashArray managed directory exports. Network shares control parallelism, proxy affinity, and snapshot processing for any backup and restore sessions for that share.

**Backup Jobs:** Backup jobs are the policies that control how and when data will be protected. Jobs can include some or all data from any number of network shares, and network shares can be protected by more than one job. The job will define the backup storage, compression, retention, archiving, and other processing parameters. File share backups follow an incremental forever approach. The first time the backup job runs it will protect all the files and folders. All other backup sessions will be incremental, capturing only data changes.

## System Recommendations for Veeam Components

Table 1 lists the system configurations Pure Storage recommends for deploying Veeam with FlashArray File Services. Please refer to the [Veeam Help Center](#) for up-to-date minimum requirements.

Component	Operating System	CPU	RAM
<b>Backup Proxy</b>	Windows Server 2016 or newer	1 core per concurrent task	4GB minimum 4GB per additional concurrent task
<b>Repository Server</b>	Windows Server 2019 or newer Linux distribution <a href="#">supported by Veeam</a>	1 core per concurrent task 12 cores minimum	4GB per concurrent task 4GB for cache repository

Table 1. System recommendations

## Scaling Considerations

You can scale the solution vertically and horizontally at the backup proxies and repositories. Proxies and repositories deployed on larger servers can support more concurrent backup and restore tasks on fewer systems, but you can achieve similar results with more, smaller systems.

Task scaling is nonlinear. Each additional concurrent task will produce a smaller improvement until you reach maximum throughput. Each additional task beyond that point will decrease throughput. The actual maximum performance you can achieve will vary based on several factors. See the [Configuration Best Practices](#) section for more detailed guidance on scaling for different data profiles.

You can also scale repository performance to some degree. Faster repositories will be able to back up and restore more data. On some storage platforms, you can gain performance by using scale-out repositories across more servers and extents.

## Understanding Backup and Recovery Performance with Veeam and FlashArray File Services

When setting recovery service level agreements (SLAs) for file data, you need to consider several factors; file access protocol, data profile, data streams, and data access node resources all have measurable effects on backup and restore throughput.



Understanding these factors and their impacts will help you ensure you can achieve your recovery point objective (RPO) and recovery time objective (RTO).

### File Access Protocol

When reading or writing files on network storage, client systems use a protocol such as NFS or SMB. The protocol defines certain sequences of operations the client must perform to accomplish its specific task in a mixture of data and non-data operations. Protocol overheads are the non-data operations that the protocol forces. While protocol overhead exists with file systems on local or storage area network (SAN) block storage, it is generally more pronounced and visible with network file storage.

On any given NAS system, you can expect SMB to be slower than NFS by 10% or more, and the gap will grow as the average file size goes down. The next section explains how file size affects backup and recovery.

With FlashArray you should use the same protocol for backup and recovery as you use for primary client access. While you can export the same file system over NFS and SMB and use different protocols for client access and backup and recovery, this is not recommended. File system permissions are only approximated when you use a different protocol, and when permissions are not restored to the original state loss of access can result.

### Data Profile

Both NFS and SMB use a combination of data and non-data operations, such as file locking and queries, to transfer file data and metadata. The ratio of data to non-data operations affects how fast the storage and client or backup agent can exchange data. The ratio varies with file size because both NFS and SMB can transfer data in blocks of 1MiB or more. With small files under 1MiB in size, it is possible to send the entire file content in a single data request, but the same transaction requires multiple non-data requests; protocol overhead is relatively high as a result. With large files, protocol overhead is lower since the same file may need many data requests, with larger data blocks, but the same number of non-data requests per transaction. While performance may vary somewhat between SMB and NFS for the same data set due to protocol differences, both will see better performance with large files than with small ones.

The data profile, therefore, has a direct impact on the performance you can expect during backup and recovery. You can back up and restore data sets with large average file sizes faster than data sets with small average file sizes, and you can back up a handful of large files much faster than many small files.

File count and file system structure also play a large part in file scan performance. At scales of millions of files and directories, the processing time to identify changes can vary by minutes or even hours. This has a significant impact on overall backup times. The Veeam metadata cache reduces the impact of file counts compared to other backup products; tuning the backup I/O control for large shares will improve backup and restore times. However, you may not be able to reach the same performance on large shares as on smaller ones. You may also have to create two or more file shares in VBR to represent the same FlashArray managed directory, as was done for this paper.

Make sure your RPO and RTO calculations factor in average file sizes, file counts, and directory structure in addition to total data size. If you have multiple managed directories or file systems with large file counts, consider creating separate file shares for them in VBR so you can back them up in separate jobs and distribute the processing.



## Concurrent Tasks

VBR has a couple of ways it manages concurrent tasks. Repositories and proxies each have a defined maximum number of tasks. The total number of tasks across all active backup and restore sessions can't exceed the lowest maximum between the proxy and repository layers, so it's important to size them to avoid unplanned bottlenecks.

File shares also have a performance tuning option that controls concurrent tasks. Moving the Backup I/O Control slider to the left decreases the number of tasks and proxies a backup or restore session that can be used for the share, and moving it to the right increases the number. This setting applies across all backup sessions for that share. For example, if you leave the backup I/O control for a share set to the default, only two proxies can run backup tasks in parallel for that share, regardless of how many proxies you have and how many backup sessions you run.

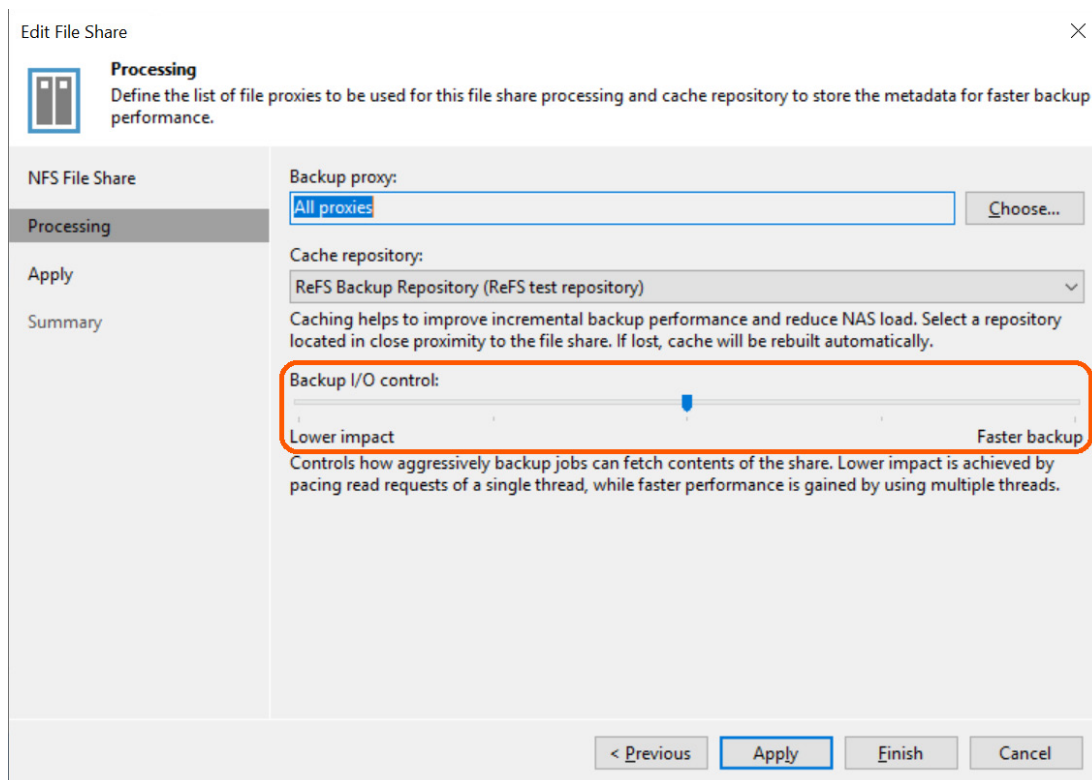


Figure 2. Veeam file share processing settings

## Scaling Recommendations for Veeam Backup Proxies

When deciding how many data access nodes to deploy, bear in mind that:

- You should start with at least two backup proxies for any environment size to provide redundancy. Two proxies should support over 3TiB/hour, depending on the networking and the data profile. Add proxies based on capacity and throughput needs.
- Backups are generally faster if you spread tasks across more proxies than if you run more tasks per proxy. Consider adding proxies before you increase the number of tasks per proxy. .
- Data profiles will affect how many tasks and proxies can effectively process a given data set. With very small average file sizes, additional tasks and proxies may not be able to increase throughput.



- Virtual proxies provide easy elasticity. You can quickly add resources or deploy new proxies if you need more throughput, and you can remove idle proxies. With virtual proxies, you will need to ensure the hypervisor hosts have sufficient available resources to support the number of proxies.

## Basic Configuration

The basic configuration should help you implement the solution as quickly and simply as possible. To add FlashArray File Services into your existing VBR environment, you need to perform three main steps: deploy file share backup proxies, create a network share client that represents the FlashArray, and configure one or more backup jobs.

**NOTE:** This document assumes you have a FlashArray joined to Active Directory and hosting data, and you have configured a Veeam repository.

## Deploying File Share Backup Proxies

Veeam file share backup proxies are simply managed Windows servers with a Veeam agent installed and designated for file share backups. See the [System Recommendations for Veeam Components](#) section for recommended system resources. See the [Veeam Backup Proxies](#) section for best practices on deploying backup proxies. To deploy a file share backup proxy,

1. Navigate to the Backup Infrastructure view. To open the New File Proxy wizard, from the Backup Proxy menu, click the Add Proxy button, then select File Share.

**NOTE:** On VBR V11, the menu item is “Backup proxy” instead of “File share.”

2. As Figure 3 shows, on the Server page, select the server that will act as a backup proxy. If the proxy is a new server not currently managed by VBR, click the Add New button, and follow the wizard to deploy the base Veeam software. Set the Max concurrent tasks value to the number of physical CPU cores or vCPUs in the server. Do not include Hyper-Threading in the CPU count. Accept the defaults on the other pages unless you need to modify the traffic rules to throttle backup speed. See the [Limiting Network Throughput](#) section for more details on network throttling.








Figure 3. New File Proxy wizard

Once the wizard is complete, you will see the proxy listed, with the “File share” type (Figure 4).











NAME	TYPE	HOST ↑	DESCRIPTION
 VMware Backup Pr...	VMware vSphere	sn1-r720-d06-01	Created by Veeam Backup & Replication
 File Backup Proxy	File share	sn1-r720-d06-01	Created by Veeam Backup & Replication
 sn1-r720-d06-03	File share	sn1-r720-d06-03	Created by SN1-R720-D06-01\Adminis
 sn1-r720-d06-05	File share	sn1-r720-d06-05	Created by SN1-R720-D06-01\Adminis
 veproxy1	File share	veproxy1	Created by SN1-R720-D06-01\roy at 6/
 veproxy2	File share	veproxy2	Created by SN1-R720-D06-01\roy at 6/
 veproxy3	File share	veproxy3	Created by SN1-R720-D06-01\roy at 6/
 veproxy4	File share	veproxy4	Created by SN1-R720-D06-01\roy at 6/
 veproxy5	File share	veproxy5	Created by SN1-R720-D06-01\roy at 6/
 veproxy6	File share	veproxy6	Created by SN1-R720-D06-01\roy at 6/

Figure 4. New File Proxy wizard

**NOTE:** On VBR V11, the proxy type is shown as “Agent” instead of “File share.”

Repeat the procedure for each backup proxy.



## Adding Network Shares

Before backing up a managed directory, you must add it to VBR as a network share. To add a network share:

1. Navigate to the Inventory view. Select the File Shares node. Click the Add File Share link in the inventory pane to open the Add File Share window (Figure 5). Click either the NFS share or SMB share link based on what you are adding.

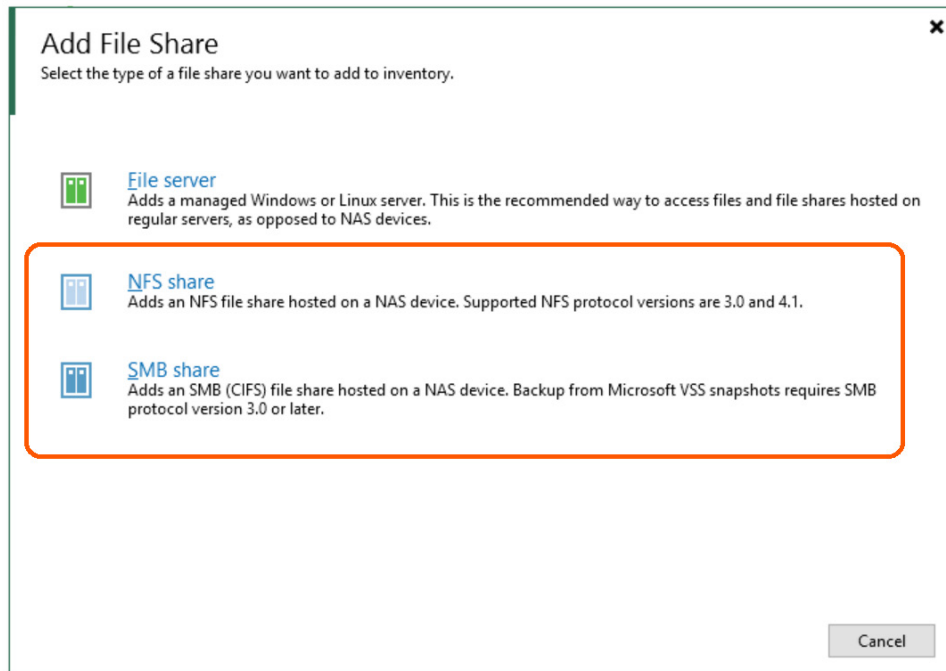


Figure 5. Add File Share window

2. Follow the below instructions for [NFS](#) or [SMB](#) shares.

### Adding an NFS File Share

To add an NFS file share to the VBR inventory:

1. On the NFS File Share page, enter the share path in NFS format, host:/share (Figure 6). You may use the FlashArray's hostname, virtual IP address, or fully qualified domain name (FQDN).

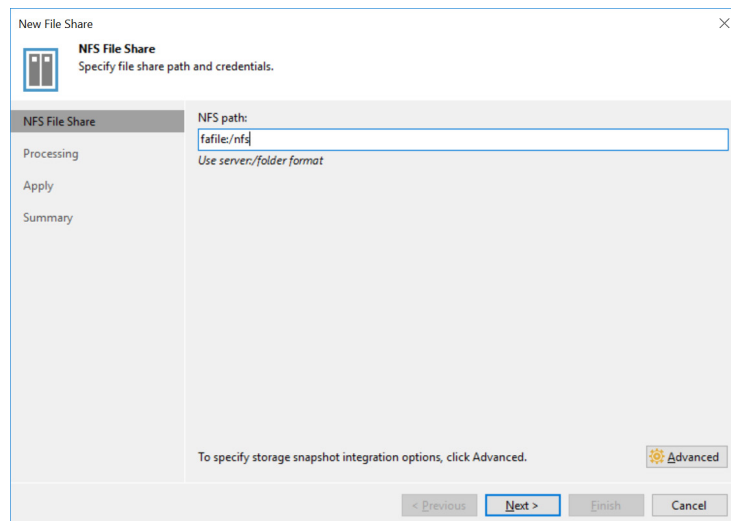


Figure 6. NFS File Share page



- a. If you want to back up from a snapshot instead of from the live file system, click the Advanced button. Enter the path to the snapshot, in NFS format (Figure 7). We recommend creating a daily snapshot on the FlashArray, using a static client name and suffix to avoid frequent configuration changes due to variable snapshot names. See the [Protect Open Files](#) section for more information on using FlashArray file snapshots.




Figure 7. NFS File Share - advanced options

3. On the **Processing** page (Figure 8), make any desired changes to the file share configuration. You can limit the file proxies that can back up or restore the share, select the cache repository for share metadata, and adjust the **Backup I/O control** for more parallelism. See the [Concurrent Tasks](#) section for more information on backup I/O control. Click the **Apply** button to commit the settings and complete the configuration.

**NOTE:** The cache repository must be a standard backup repository. It cannot be a SOBR.

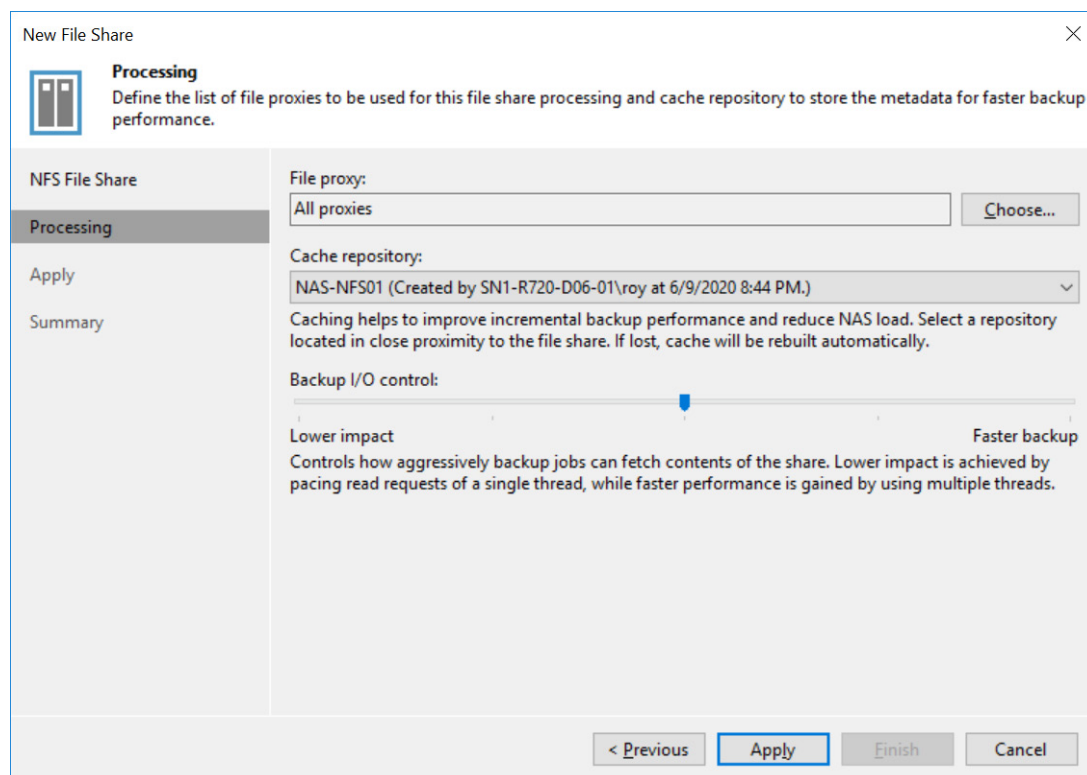


Figure 8. NFS File Share - Processing page



- On the **Apply** page (Figure 9), click the **Next** button to review the results summary, or click the **Finish** button to close the wizard.

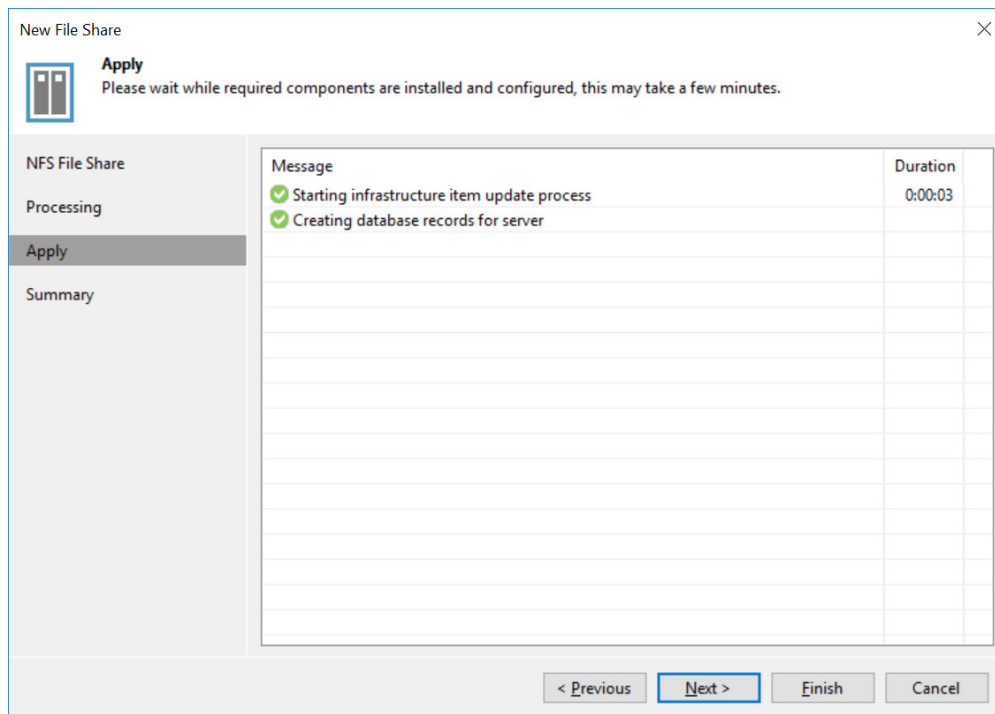


Figure 9. NFS File Share - Apply page

## Adding an SMB File Share

To add an SMB file share to the VBR inventory:

- On the SMB File Share page, enter the share path in Universal Naming Convention (UNC) format, \\host\share (Figure 10). You may use the FlashArray's hostname, virtual IP address, or FQDN.

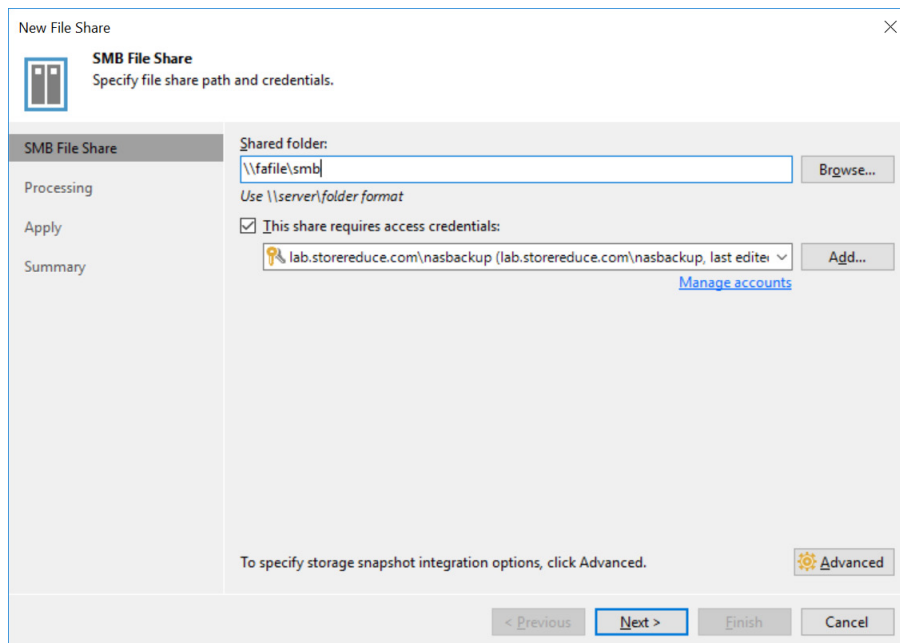


Figure 10. SMB File Share page



- a. If you want to back up from a snapshot instead of from the live file system, click the Advanced button. Enter the path to the snapshot, in UNC format (Figure 11). We recommend creating a daily snapshot on the FlashArray, using a static client name and suffix to avoid frequent configuration changes due to variable snapshot names. See the [Protect Open Files](#) section for more information on using FlashArray file snapshots.

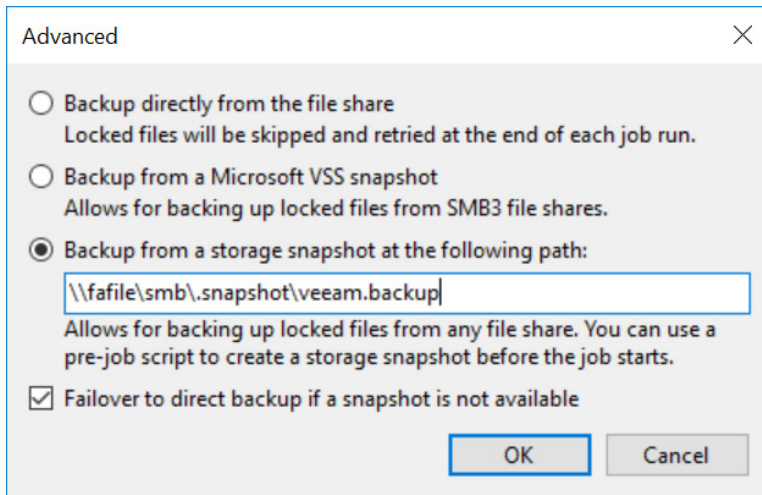


Figure 11. SMB File Share - advanced options

2. On the **Processing** page (Figure 12), you may make any desired changes to the file share configuration. You can limit the file proxies that can back up or restore the share, select the cache repository for share metadata, and adjust the Backup I/O control for more parallelism. See the [Concurrent Tasks](#) section for more information on backup I/O control. Click the **Apply** button to commit the settings and complete the configuration.

**NOTE:** The cache repository must be a standard backup repository. It cannot be a SOBR.

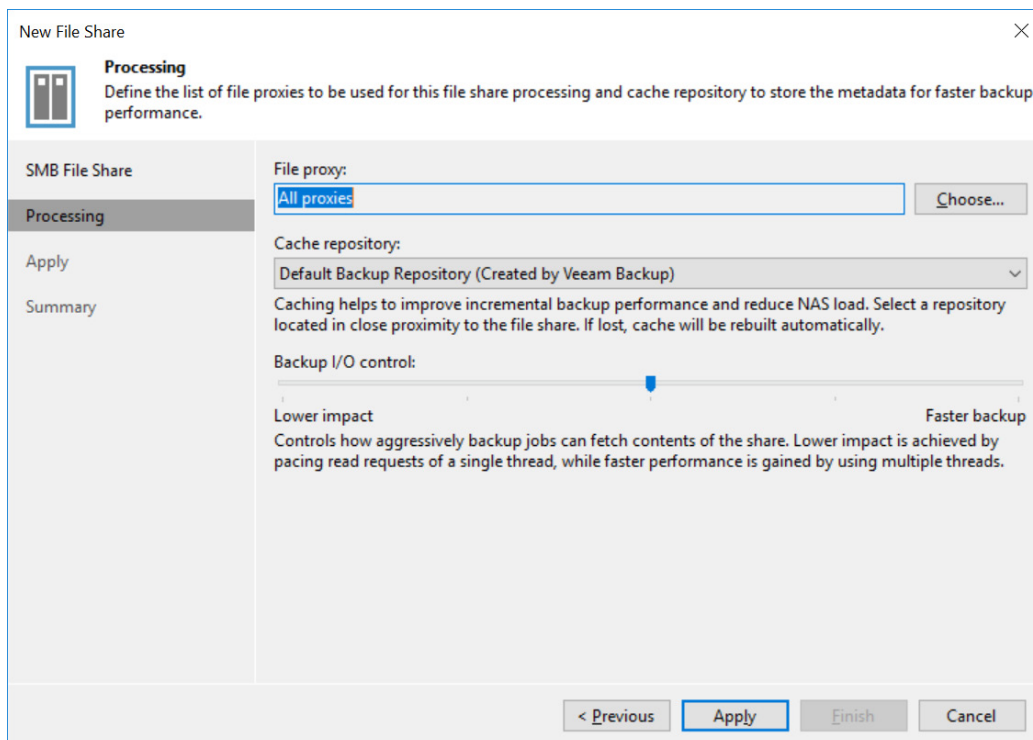


Figure 12. SMB File Share - Processing page



3. On the **Apply** page (Figure 13), click the **Next** button to review the results summary, or click the **Finish** button to close the wizard.

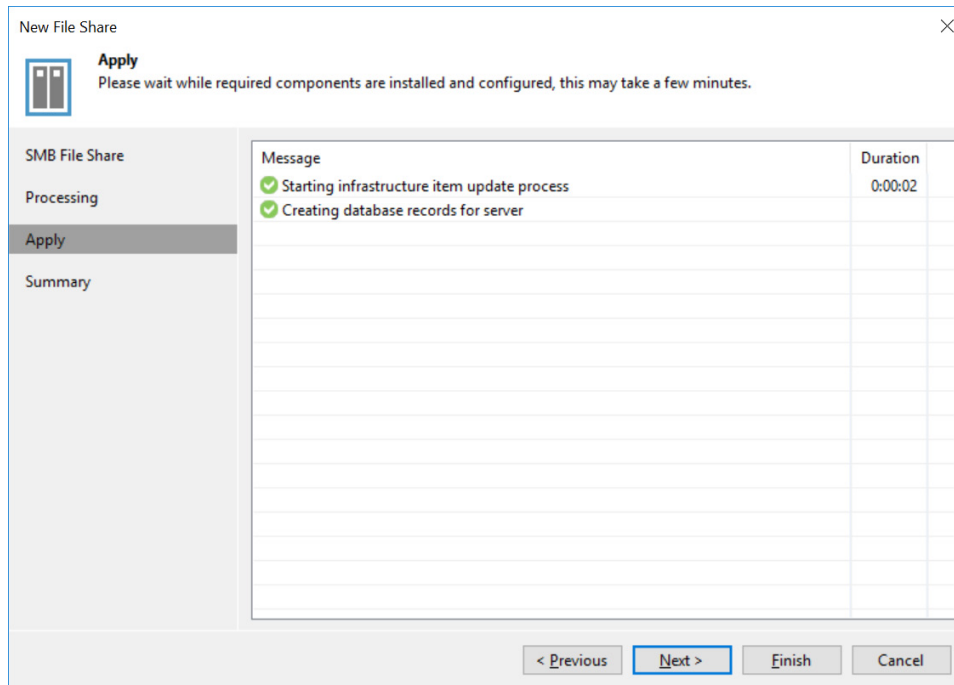


Figure 13. SMB File Share - Apply page

## Creating Backup Jobs

To protect file data, you must create one or more backup jobs. To create a backup job:

1. Backup Job > File Share Navigate to the Home view. From the Home menu, click the Backup Job button, then select the File share option. The New File Backup Job wizard appears.
2. On the Name page (Figure 14), enter a unique name for the backup job. You can also enter a description.


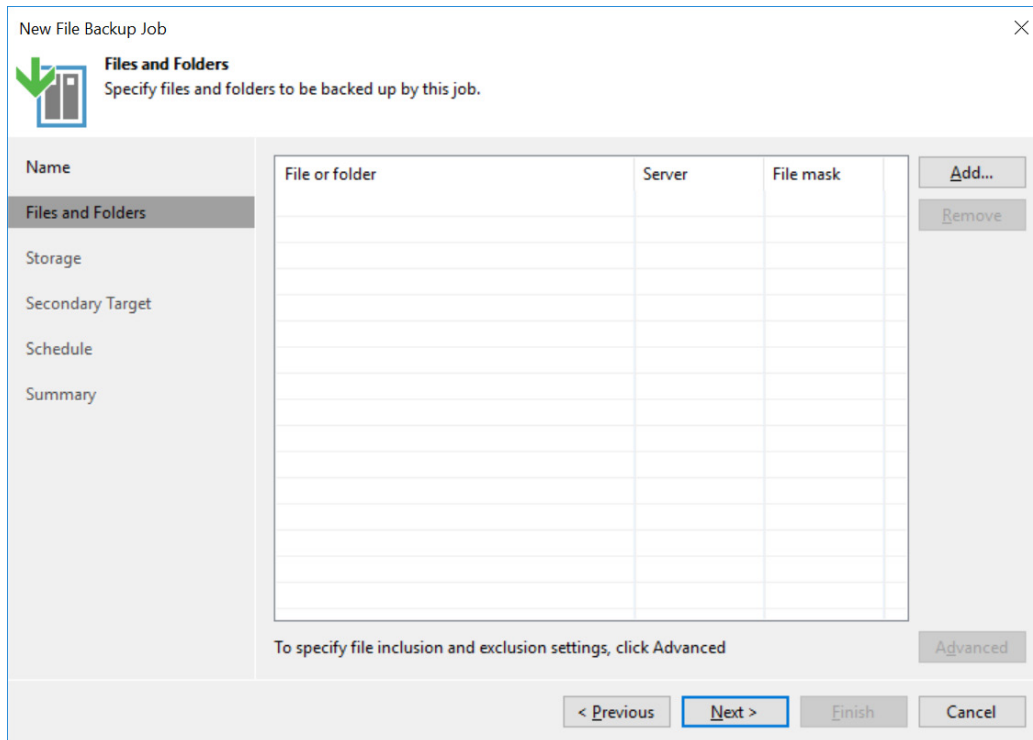


Figure 14. New File Backup Job - Name page



- On the Files and Folders page (Figure 15), click the **Add** button to open the **Select File or Folder** dialog and choose the source data to be backed up.



**Figure 15.** New File Backup Job - Files and Folders page

- As Figure 16 shows, from the **Server** dropdown, select the file share where the data resides. In the **Folders** pane, you can select the root folder to back up the entire share, or you can expand the file share contents in the **Folders** pane and select individual files and folders if you want to back up only part of the share.

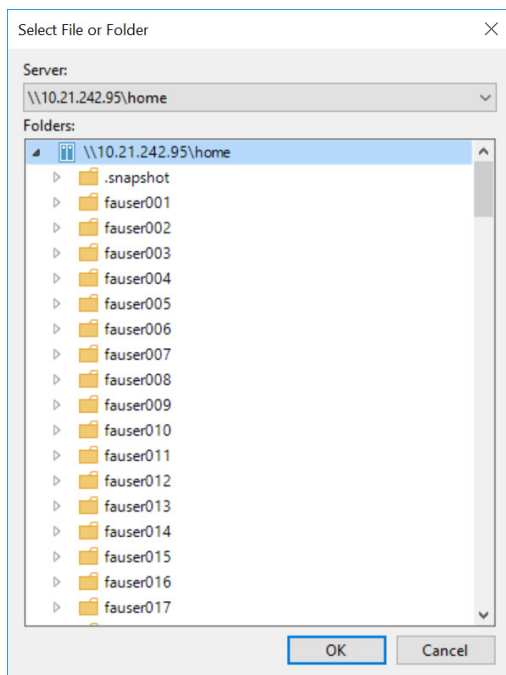


Figure 16. Select File or Folder dialog

5. As Figure 17 shows, you can add as many shares and paths to the backup job as you need. You can specify inclusion and exclusion filters for one or more paths by selecting them and clicking the **Advanced** button.

New File Backup Job

**Files and Folders**  
Specify files and folders to be backed up by this job.

Name	File or folder	Server	File mask
Files and Folders	\\fafile\smb	\\fafile\smb	Excluding \\...
Storage	fafile:/nfs	fafile:/nfs	Excluding fa...
Secondary Target			
Schedule			
Summary			

To specify file inclusion and exclusion settings, click Advanced

Advanced

< Previous Next > Finish Cancel

Figure 17. New File Backup Job - Files and Folders page

6. On the Storage page (Figure 18), select the backup repository to store the backup data. You can set the backup retention and, optionally, add archive retention settings. Click the **Advanced** button to access settings for handling ACLs, scripts, and other advanced options.

New File Backup Job

**Storage**  
Specify target backup repository and file retention policy for this job.

Backup repository:  
Scale-out Backup Repository 1 (Created by SN1-R720-D06-01\Administrator at 8/7/2020 10:15 PM.)

68.1 TB free of 149 TB

Keep all file versions for the last: 28 days

Retains recent versions of each file for the specified time period, allowing for restore of entire file shares to a point-in-time state, restore of deleted files, and restore of earlier file versions.

☐ Keep previous file versions for: 3 years

Archives older versions of active and permanently deleted files after they are no longer covered by the recent versions retention policy. For scalability reasons, we recommend using object storage.

Archive repository:  
Default Backup Repository (Created by Veeam Backup)

Files to archive:  
All

Advanced job settings include notification settings, automated post-job activity and other settings.

Advanced

< Previous Next > Finish Cancel

Figure 18. New File Backup Job - Storage page





- On the **ACL Handling** tab (Figure 19), you can choose whether to collect permissions from only folders or all folders and files. Collecting file permissions will slow backup performance significantly, so you should only select this option if you use complex permissions at the file level.

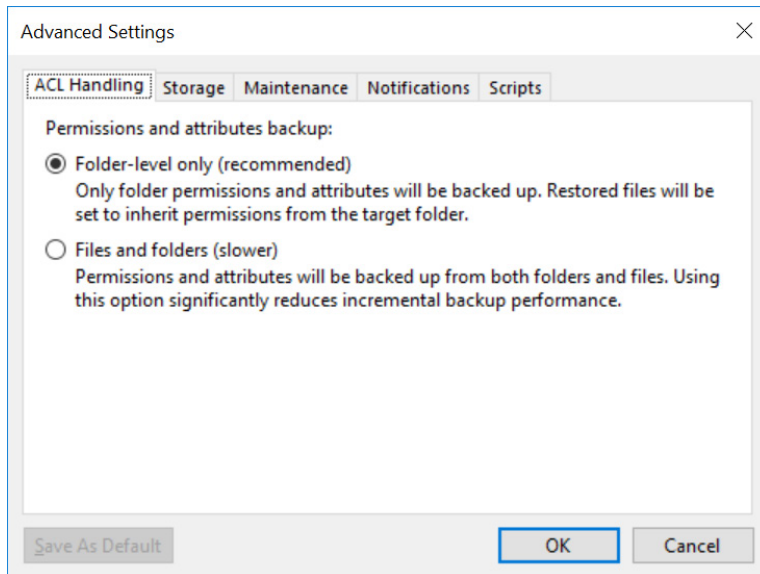


Figure 19. New File Backup Job - Advanced Settings - ACL Handling tab

- On the **Storage** tab (Figure 20), you can set the compression level. In most cases, you should choose the Optimal level. This will significantly reduce the amount of data sent between the backup proxies and repository servers. See the [Backup Job Compression](#) section for more information on compression settings.

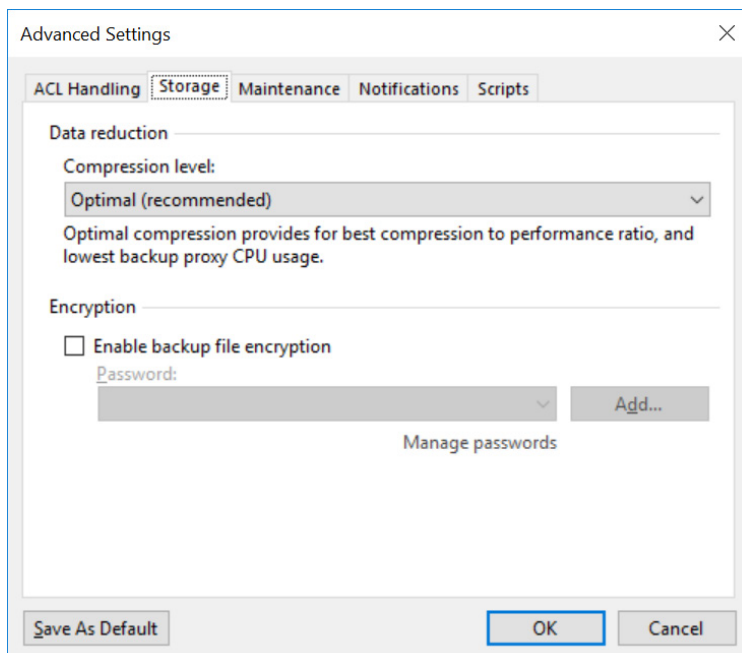


Figure 20. New File Backup Job - Advanced Settings - Storage tab



9. If you want to have the backup job run scripts before or after backup sessions, you can enter the script path and arguments on the **Scripts** tab (Figure 21). For example, you can use scripts to create snapshots for open file protection. For more information on protecting open files using snapshots, refer to the [Protect Open Files](#) section.

Advanced Settings

ACL Handling Storage Maintenance Notifications **Scripts**

**Job scripts**

☐ Run the following script before the job:

☐ Run the following script after the job:

☒ Run scripts every 1 backup session

☐ Run scripts on the selected days only

Saturday

Save As Default OK Cancel

Figure 21. New File Backup Job - Advanced Settings - Script tab

10. If you need additional data copies, you can configure them on the Secondary Target page (Figure 22).

New File Backup Job

**Secondary Target**

We can create additional copies of the short-term file store for redundancy, using the same or different retention policy. The data copy process will start automatically after each primary job run.

Secondary repositories:

Name	Capacity	Retention

Add... Edit... Remove

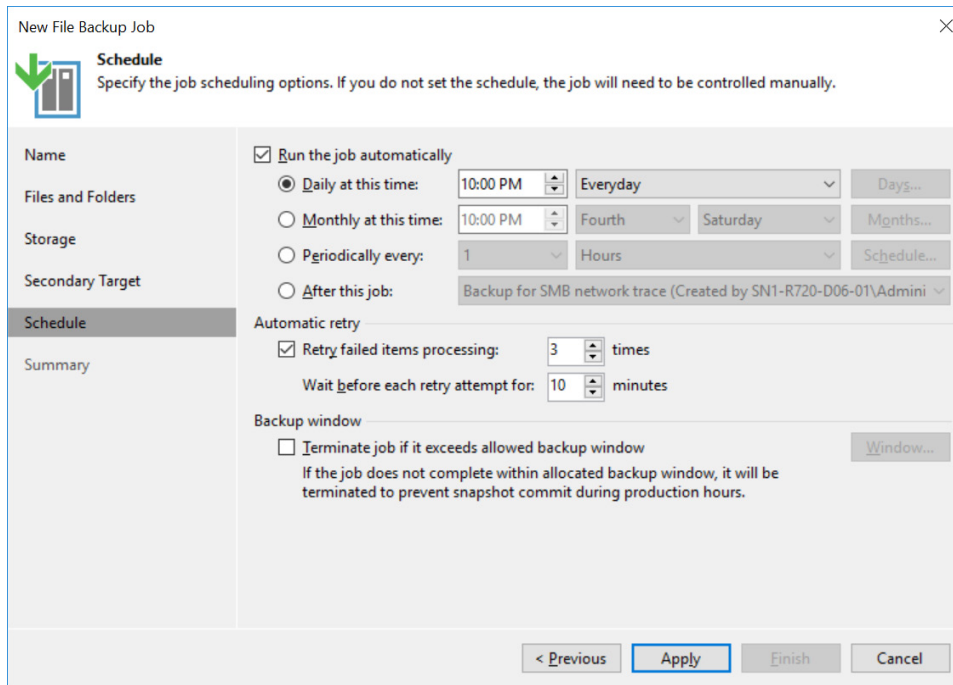
Data duplication to each secondary repository is performed automatically. You can customize retention, encryption and copy window settings by selecting the repository and clicking Edit.

< Previous Next > Finish Cancel

Figure 22. New File Backup Job - Secondary Target page



11. On the **Schedule** page (Figure 23), you can set a recurring schedule and related options for the backup job.



**New File Backup Job**

**Schedule**  
Specify the job scheduling options. If you do not set the schedule, the job will need to be controlled manually.

**Name**

**Files and Folders**

**Storage**

**Secondary Target**

**Schedule**

**Summary**

☒ **Run the job automatically**

☒ **Daily at this time:** 10:00 PM Everyday Days...

☐ **Monthly at this time:** 10:00 PM Fourth Saturday Months...

☐ **Periodically every:** 1 Hours Schedule...

☐ **After this job:** Backup for SMB network trace (Created by SN1-R720-D06-01\Admini)

**Automatic retry**

☒ **Retry failed items processing:** 3 times

Wait before each retry attempt for: 10 minutes

**Backup window**

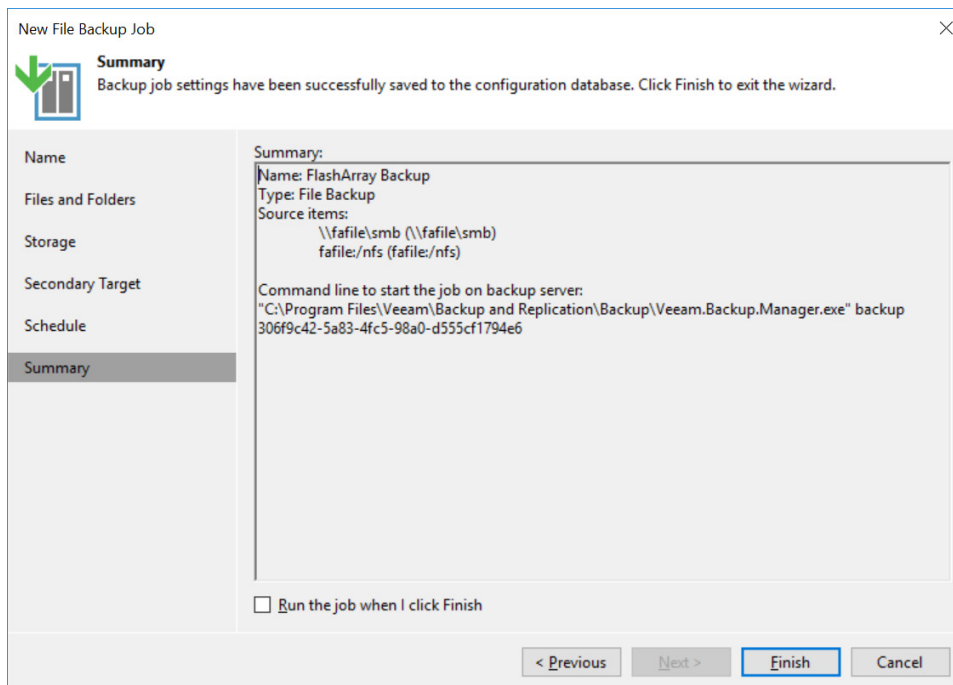
☐ **Terminate job if it exceeds allowed backup window** Window...

If the job does not complete within allocated backup window, it will be terminated to prevent snapshot commit during production hours.

< Previous Apply Finish Cancel

Figure 23. New File Backup Job - Schedule page

12. The **Summary** page (Figure 24) will list the options you selected. You can also choose to start the job immediately.



**New File Backup Job**

**Summary**  
Backup job settings have been successfully saved to the configuration database. Click Finish to exit the wizard.

**Name**

**Files and Folders**

**Storage**

**Secondary Target**

**Schedule**

**Summary**

**Summary:**

Name: FlashArray Backup  
Type: File Backup  
Source items:  
\\fatile\smb (\\fatile\smb)  
fatile/nfs (fatile/nfs)

Command line to start the job on backup server:  
"C:\Program Files\Veeam\Backup and Replication\Backup\Veeam.Backup.Manager.exe" backup  
306f9c42-5a83-4fc5-98a0-d555cf1794e6

☐ **Run the job when I click Finish**

< Previous Next > Finish Cancel

Figure 24. New File Backup Job - Summary page

You can repeat the process to create as many backup jobs as you need to meet your protection needs. The next section covers best practices to help you tune the environment for the best performance.



## Configuration Best Practices

### File Systems and Managed Directories

File systems act as top-level containers for FlashArray File Services data. Within a file system, FlashArray allows two types of directories: normal and managed. From the client's perspective, both appear identical, and you can apply permissions and other properties to them. However, there are several key differences between them: Normal directories are created through the file system from a client system as you would create directories on a local file system. The FlashArray administrative interfaces do not list normal directories and cannot manipulate them. Managed directories are special directories that allow you to apply policies, such as export and protection policies. You can only create managed directories through the FlashArray administrative interfaces, either with GUI, CLI, or API. When you create a managed directory, it adds a directory into the client view, with the path you specify. You can only take snapshots on managed directories. Managed directories can't nest within normal directories, although they can nest in other managed directories to a certain depth. The file system root is also a managed directory. There are several limitations on managed directories. You can't convert a normal directory to a managed directory, and vice versa. You also can't delete managed directories from client systems, and administrators cannot delete them unless they are empty. For more information on managed directories and how to use them, see the Pure Storage white paper [Introduction to Pure FlashArray File Services](#).

#### Group Associated Data Within File Systems

You should group data sets that you want to use, protect, and recover together into the same file system. For example, create user home directories as managed directories in a single file system. If different teams need shared SMB directories, create managed directories for each team in a single file system. Grouping directories will make protection and recovery simpler, especially as relates to snapshots and open file protection.

#### Protect Managed Directory and Policy Configurations

Managed directories are a key part of a FlashArray File Services environment. It is therefore important to protect their configurations so you can recover them if there's a catastrophic loss. You can create an export script to run regularly to capture the key configuration details into a recovery script, which you can then execute before full recovery to recreate the key structures. You will need the commands from Table 2, entered in order:

Command	Function
<code>purefs list -cli</code>	Outputs commands to recreate all file systems
<code>puredir list -cli</code>	Outputs commands to recreate all managed directories
<code>purepolicy nfs list --cli</code>	Outputs commands to recreate all NFS policies
<code>purepolicy nfs rule list --cli</code>	Outputs commands to recreate the rules in each NFS policy
<code>purepolicy smb list --cli</code>	Outputs commands to recreate all SMB policies
<code>purepolicy smb rule list --cli</code>	Outputs commands to recreate the rules in each SMB policy
<code>puredir export list --cli</code>	Outputs commands to recreate all managed directory exports



<code>purepolicy snapshot list --cli</code>	Outputs commands to recreate all protection policies
<code>purepolicy snapshot rule list --cli</code>	Outputs commands to recreate the rules for each protection policy
<code>purepolicy snapshot list --member --cli</code>	Outputs commands to recreate the memberships for each protection policy, effectively enabling snapshot schedules

Table 2. Commands to output key configuration elements

You must protect the scripts that are external to the FlashArray to ensure you can recover from a catastrophic loss. If you save them to a directory on a FlashArray file system, you can use VBR to back up the generated configuration scripts for extra protection.

Together these commands will generate a script that will recreate the entire set of file systems, managed directories, export policies, and protection policies. You can skip any part of the configuration that already exists when performing the restore. This is ideal for most major recoveries. If you want to capture configurations for only a subset, such as a single file system or managed directory, you can apply filters (which all support wildcards) to the commands as follows:

- To output only a specific file system, use:

```
purefs list <file system name> --cli
```

- To output only the directories in that file system, use:

```
puredir list --file-system <file system name> --cli
```

- To output all exports for that file system, use:

```
puredir export list --dir "<file system name>:*" --cli
```

- To output all snapshot policy memberships for that file system, use:

```
purepolicy snapshot list --member --filter "member.name='<file system name>:*'" --cli
```

We don't recommend capturing a subset of policies and rules, as it's not possible to filter the list of policies or rules based on members, only by policy name and type. It is simpler to capture all policies and use either the `puredir export list` or `purepolicy snapshot list --member` command to capture the specific directories' memberships.

**IMPORTANT:** Running the output of the `puredir export list --cli` command will make data accessible to end-users. If you prefer to recover data before creating exports, save the command's output to a separate file that you can run independently of the other script. You may also want to run a filtered capture of a subset of exports, such as all file system roots, by inserting a name or wildcard into the command. For example, to output the commands to recreate exports of file system roots, you would use the command `puredir export list --dir "*:root" --cli`.

### Create Managed Directories Before Restoring Data

When recovering an entire file system or managed directory after a catastrophic loss, you must recreate managed directories before restoring data. Backup software has no way of knowing whether a directory is normal or managed. On recovery, the backup software will create any required top-level directories as normal directories if they do not already exist. The software can't convert these to managed directories without a significant effort, which may include repeating data recovery.



If you are exporting configuration scripts regularly, you can simply run the latest version to recreate the managed directories and policies before you restore file data.

## Veeam Backup Proxies

**Placing Backup Proxies:** Veeam backup proxies need to read and write data as quickly as possible using the SMB and NFS protocols. Wherever possible, locate nodes in the same site as the FlashArray and on the same virtual local area network (VLAN).

**Joining Backup Proxies to Active Directory:** While proxies don't need to belong to Active Directory for backups to work, you can take advantage of Kerberos security for SMB file shares if they belong.

## Veeam Network Shares

You should align network shares in the Veeam inventory with FlashArray exported managed directories. Add a network share to VBR for each exported managed directory. This lets you tune performance for different directories that may have different data profiles or priorities.

## Veeam Repositories

You should follow Veeam's [guidance](#) on configuring repositories. If you use Pure FlashArray//C for repository storage, you should review the best practices in the [Enhancing Veeam with FlashArray//C](#) white paper.

You should balance the number of tasks available at the backup proxies and repositories, as having many more resources available on one layer than the other will lead to wasted resources and potentially lower-than-optimal performance.

## Performance

This section details the best practices for configuring VBR elements when protecting FlashArray File Services data.

### Backup I/O Control

For most shares, use the default backup I/O control setting. This gives a good balance of performance and storage impact. If you are not achieving the desired backup or restore throughput, move the slider to the right to increase available resources for that share. If you increase the backup I/O control all the way to the right and still don't see the performance you need, you will have to deploy more backup proxies or SOBR extents to increase the available task pool. Be aware, however, that some data profiles may never be able to reach the maximum throughput. Be aware that there is a point where setting too many parallel tasks on a share will actually decrease performance.

### Backup Job Compression

Backup jobs have several available compression settings. While VBR cannot compress traffic between the FlashArray shares and backup proxies, it does compress traffic from proxies to backup repositories. If you have a network bottleneck between the proxies and repositories, compression can significantly increase the amount of data the proxies can process. If you use the same servers as proxies and repository servers, enabling compression is unlikely to help throughput, and you should consider disabling it. If you separate proxies and repository servers, but you rely on backup storage such as Pure FlashArray//C to reduce the data size, you should enable compression in the backup job but set the repository to decompress data before writing.



### Antimalware Exclusions

Portions of the backup and restore processes involve heavy write activity on the backup proxies and repository servers. Real-time malware protection can significantly reduce performance. In lab tests, excluding the VeeamAgent.exe process from scans improved backup throughput by over 20%. VBR installs a 32-bit and 64-bit version, both of which should be excluded from scanning.

**NOTE:** To reduce the risk of spoofing, you should exclude processes using the full path or a file signature, if supported. Refer to your antimalware product documentation for instructions on setting exclusions.

### Limiting Network Throughput

There may be situations where you may want to limit backup traffic to reserve resources on the FlashArray. This will prevent backups from slowing down production workloads. There are several ways you can do this:

- Reduce the backup I/O control level on the file shares. Reducing the number of proxies and tasks will limit the number of connections to FlashArray, reducing the throughput and lowering the array load.
- Using VBR's network traffic rules to limit traffic between the backup proxies and repository servers. You can define a rule that restricts the combined backup throughput from all the proxies to all the repositories. All parallel backups will share the restricted bandwidth, but any single proxy has access to the full amount. The throttling setting does not take compression into account, so you should set the bandwidth cap lower than the limit you want to see at the FlashArray. Refer to [Enabling Traffic Throttling](#) in the Veeam Help Center for instructions on enabling network throttling.
- Using quality of service (QoS) on your network to ensure production workloads take priority over backup.
- Using the hypervisor's traffic shaping capabilities to control bandwidth if you are using virtual backup proxies.
- Combining some or all these options to minimize the chance of backups affecting production performance.

## SMB-specific Best Practices

### Manage SMB Backup Access

Since backups occur through network shares, you must configure a service account to access and back up the files. Create the service account in the Active Directory domain to the joined FlashArray File Services and add it to the Backup Operators group on the domain. You should not grant the account local login access or elevated privileges on any Windows system.

**IMPORTANT:** You must set the uidNumber and gidNumber attributes on the service account in Active Directory. You should not use this account for any other purpose in your environment.

You can add the account to VBR to easily reuse it across multiple file shares. You can also easily update the password in VBR when you change it on the domain and avoid having to update all the file share configurations. To add a credential:

1. From the VBR main menu, select **Manage Credentials**.
2. From the Manage Credentials window, click the **Add** button, then select **Standard account** from the menu.
3. In the Credentials dialog (Figure 25), enter the username and password in the appropriate fields, then click the OK button to create the credential.



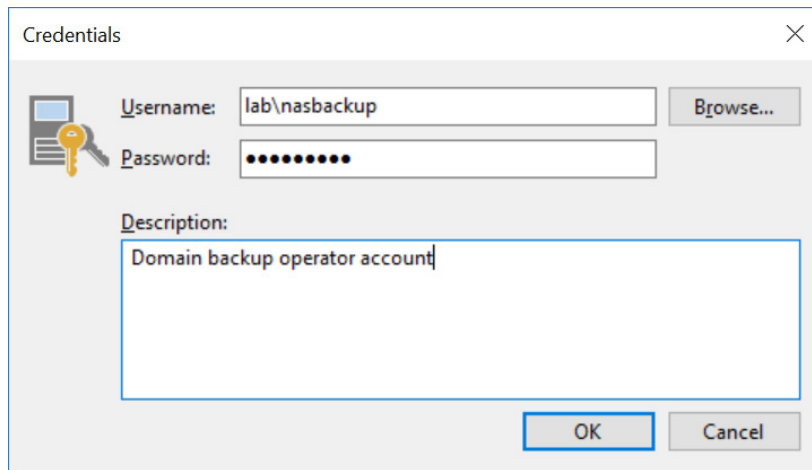


Figure 25. Credentials dialog

You can use the credential on SMB file shares by selecting it on the SMB File Share page of the configuration wizard (Figure 26).

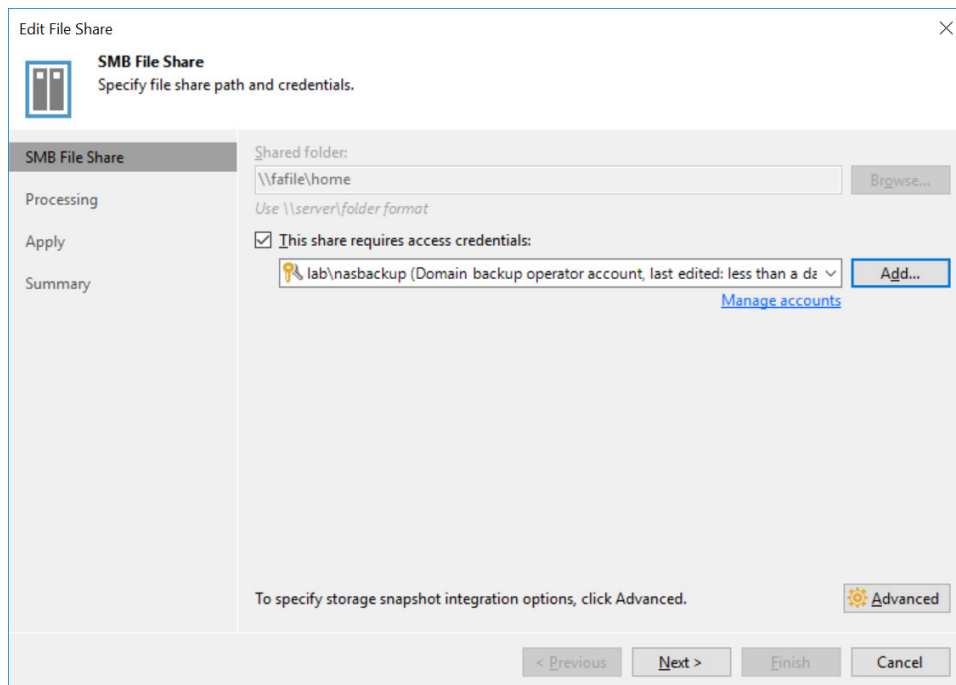


Figure 26. Edit File Share - SMB File Share page

You can change the password in VBR by selecting it in the Manage Credentials dialog and clicking the Edit button, then entering the new password in the Password field.





## NFS Best Practices

### Restrict NFS export access

You should use restrictive NFS export rules for your FlashArray file systems and managed directories unless you have a specific requirement. You will need to grant access to the Veeam backup proxies. Create your export rules as follows (Figure 27):

- If you have a range or name pattern that applies only to your backup proxies, you can use a single rule and enter the pattern in the **Client** field. Otherwise, create a separate rule for each proxy and enter its IP address.
- For the **Access** option, use the same root-squash or no-root-squash option as you set for your production clients. Using a different setting may lead to access issues.
- For the **Permission** option, select “rw.” You can choose the “ro” setting without affecting backups, but it will cause recovery failures.

Figure 27. FlashArray NFS export rule definition

## Backup Best Practices

### Protect Open Files

Ensuring that you back up critical files is an important part of an effective NAS protection strategy. There are several reasons why your backup software might not protect all your NAS data. Your backup software can't back up locked files or files in use by users and applications. Storage snapshots are a simple and effective way to provide that assurance. Integrating snapshots into VBR network share backups requires several other configuration changes, which vary based on the file protocol. This section details the procedures and best practices for using VBR with FlashArray File Services snapshots.

### FlashArray File Services Snapshot Behavior

FlashArray File Services creates snapshots of managed directories, including file system roots, capturing the state of the entire directory structure at the time of the snapshot. Snapshots will capture nested managed directories. You can manage snapshots manually, by script, or by policy, and access them through the .snapshot folder in the base of the managed directory where you took the snapshot. Snapshot names have two parts: a client name, which the administrator supplies, and a suffix, which the array increments automatically when managed through a protection policy. The array appends the client's name and suffix to the managed directory name to create the snapshot name. For example, for a FlashArray called “farray,” a



file system called “homedirs” might have a managed directory named “alan.” If the snapshot client name is “veeam” and the suffix has incremented to 348, the snapshot name would be “homedirs:alan.veeam.348.” Figure 28 shows the management view of the managed directory and snapshot.

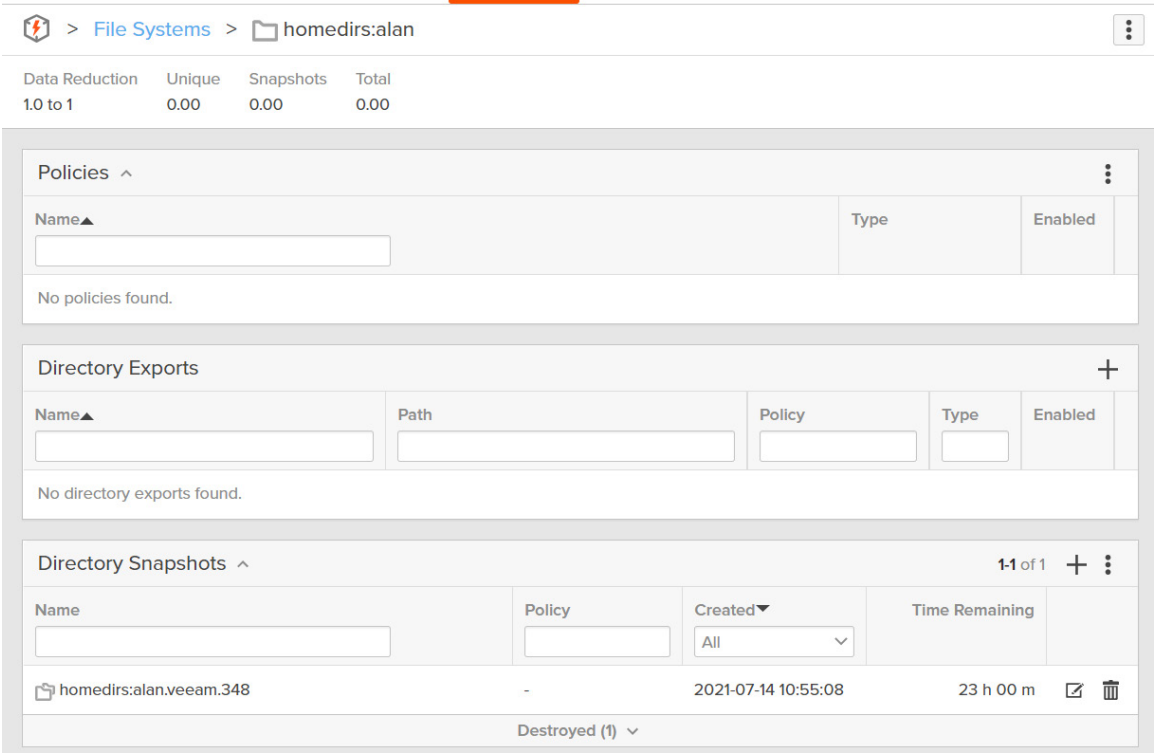


Figure 28. Managed directory snapshot, management view

Within the client view, you see only the client’s name and suffix of the snapshot, so for this example, an SMB client could access the snapshot at \\fatile\homedirs\alan\.snapshot\veeam.348. Figure 29 shows the client view.

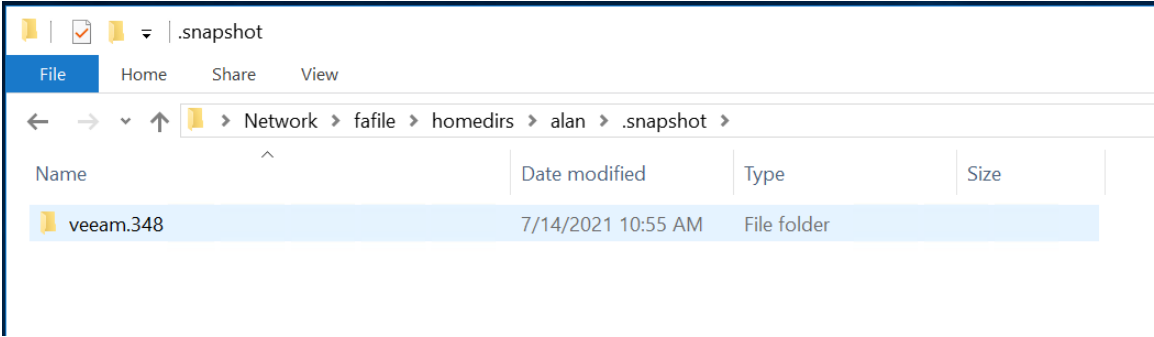


Figure 29. Managed directory snapshot, client view

You can also create snapshots manually using the FlashArray GUI or CLI commands. You can specify a suffix for these snapshots, which lets you define a consistent, predictable snapshot name.



## Use Folder-Level ACL Handling

Backing up permissions and attributes from all files can significantly reduce backup and restore performance. You should always use the **Folder-level only** ACL handling job option (Figure 30) unless you set explicit permissions on files. You should try to place directories that need file permissions captured into separate backup jobs to minimize the impact.

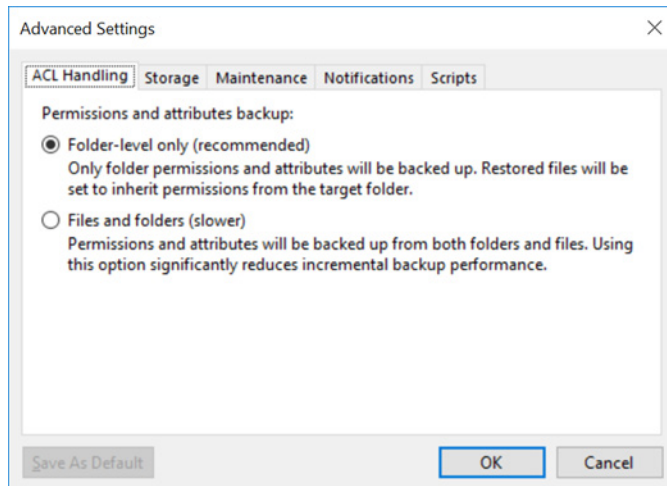


Figure 30. Restore overwrite selection

## General Snapshot Best Practices

### Create Snapshots on File Share Managed Directories

You can create snapshots at any managed directory level, including the file system root. However, you should align your snapshots for a VBR file share with the exported managed directory on FlashArray. For instance, if you have VBR configured to use a share coming from a file system root, create the snapshots on that root. If the file share is tied to a managed directory, create snapshots there and not at the root. Aligning snapshots to the file shares lets you more easily tie creating the snapshots to backup jobs and customize the timing and behavior for each file share.

### Use a Consistent Snapshot Name

You can simplify the VBR file share configuration by using the same client's name and suffix for all your snapshots. This lets you set the snapshot path for each file share to a consistent .snapshot subfolder. Following the earlier example, you could create snapshots at the file system root and name the snapshot veeam.backup. You could then statically configure the \\fafile\homedirs file share in VBR to use the snapshot path \\fafile\homedirs\.snapshot\veeam.backup.

**NOTE:** You will need to use scripts to achieve consistent snapshot naming.

### Use Scripts to Create Snapshots for Backups

Since scripts let you control the snapshot suffix, using them lets you reuse the same snapshot name and configure a static snapshot path in VBR file shares. FlashArray File Services does not refresh existing snapshots, so either your script will need to delete the existing snapshot and create a new one with the same name, or it will need to set the snapshot to expire before the next backup session so it doesn't exist.



## Recovery Best Practices

### Recreate Managed Directories Before Restoring Data

After a catastrophic loss that requires creating file system structures from scratch, you must recreate managed directories before you restore any data. File recovery will automatically create any missing directories, and since clients like backup software can only create normal directories, you would not be able to apply export or protection policies after restoring files. If you have protected your configuration, you can easily recover the managed directories before you begin the data recovery process. For more information on protecting the managed directory configuration, see the [Protect Managed Directory and Policy Configuration](#) section.

### Restore Large File Shares First

If you need to restore multiple file shares, you should start with the ones where you have increased backup I/O control. This ensures VBR has enough free tasks to allocate to those shares. You can start restoring other shares once the proxies have been allocated to the first sessions. VBR will allocate tasks to the rest of the restore sessions as they become available.

### Avoid Overwriting Large Data Sets

When you restore data, you have the option to keep or overwrite existing files (Figure 31). When restoring large data sets of thousands of files to their original folders, we recommend you use the **Keep** option to reduce the recovery time. You can also restore to a different location using the **Copy to** function.

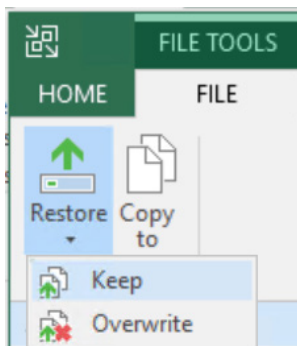


Figure 31. Restore overwrite selection

## Conclusion

Veeam Backup & Replication provides a simple yet powerful platform for data protection and recovery that you can use with confidence to ensure the availability of your FlashArray File Services data. Veeam can meet all your NAS backup and recovery needs, from individual files to entire arrays.

To learn more about FlashArray File Services or Veeam Backup & Replication, contact your Pure and Veeam account teams, or your reseller of choice.



## Additional Resources

- Learn more about [Veeam Backup & Replication](#).
- Visit the [Veeam Help Center](#).
- Learn about [Pure FlashArray File Services](#).
- Find out how to [enhance Veeam with FlashArray//C](#).



## About the Author



Roy Child is a Senior Solution Architect with Pure Storage, responsible for defining data recovery solutions and reference architectures for primary workloads such as Oracle, SQL, and VMware. Roy has worked in and with the data protection industry for more than 20 years, from end-user to IT architect with multiple backup and recovery products, followed by product management with Commvault. Roy joined Pure Storage in April 2019.

©2021 Pure Storage, the Pure P Logo, and the marks on the Pure Trademark List at <https://www.purestorage.com/legal/productenduserinfo.html> are trademarks of Pure Storage, Inc. Other names are trademarks of their respective owners. Use of Pure Storage Products and Programs are covered by End User Agreements, IP, and other terms, available at: <https://www.purestorage.com/legal/productenduserinfo.html> and <https://www.purestorage.com/patents>

The Pure Storage products and programs described in this documentation are distributed under a license agreement restricting the use, copying, distribution, and decompilation/reverse engineering of the products. No part of this documentation may be reproduced in any form by any means without prior written authorization from Pure Storage, Inc. and its licensors, if any. Pure Storage may make improvements and/or changes in the Pure Storage products and/or the programs described in this documentation at any time without notice.

THIS DOCUMENTATION IS PROVIDED "AS IS" AND ALL EXPRESS OR IMPLIED CONDITIONS, REPRESENTATIONS AND WARRANTIES, INCLUDING ANY IMPLIED WARRANTY OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE, OR NON-INFRINGEMENT, ARE DISCLAIMED, EXCEPT TO THE EXTENT THAT SUCH DISCLAIMERS ARE HELD TO BE LEGALLY INVALID. PURE STORAGE SHALL NOT BE LIABLE FOR INCIDENTAL OR CONSEQUENTIAL DAMAGES IN CONNECTION WITH THE FURNISHING, PERFORMANCE, OR USE OF THIS DOCUMENTATION. THE INFORMATION CONTAINED IN THIS DOCUMENTATION IS SUBJECT TO CHANGE WITHOUT NOTICE.

Pure Storage, Inc.  
650 Castro Street, #400  
Mountain View, CA 94041