

WHITE PAPER

# FlashArray File Services Protection Using Veeam Backup & Replication

A Best Practices Guide

# Contents

<b>Introduction</b>	3
<b>How to Use This Guide</b>	3
<b>Solution Architecture</b>	4
Components	5
<b>System Recommendations for Veeam Components</b>	5
Scaling Considerations	6
<b>Understanding Backup and Recovery Performance with Veeam and FlashArray File Services</b>	6
File Access Protocol	6
Data Profile	7
Concurrent Tasks	8
Scaling Recommendations for Veeam Backup Proxies	9
<b>Basic Configuration</b>	9
Deploying Backup Proxies	10
Adding Network Shares	11
Creating Backup Jobs	16
<b>Configuration Best Practices</b>	21
File Systems and Managed Directories	21
Create Managed Directories Before Restoring Data	24
Veeam Backup Proxies	24
Veeam Network Shares	24
Veeam Backup Repositories	24
Performance	25
<b>SMB-specific Best Practices</b>	26
Manage SMB Backup Access	26
Restrict SMB Export Access	28
<b>NFS Best Practices</b>	29
Restrict NFS Export Access	29
<b>Backup Best Practices</b>	30
Protect Open Files	30
FlashArray File Services Snapshot Behavior	30
Use Folder-level ACL Handling	32
<b>General Snapshot Best Practices</b>	33
Create Storage Snapshots on File Share Managed Directories	33
Use Protection Policy for Storage Snapshots	33
<b>Recovery Best Practices</b>	35
Recreate Managed Directories Before Restoring Data	35
Restore Large File Shares First	35
<b>Conclusion</b>	35
<b>Additional Resources</b>	35



## Introduction

Pure Storage® FlashArray™ File Services brings the reliability, data reduction, and simplicity of Purity//FA to network attached storage (NAS). Real-time Enterprise File on the Pure Storage platform can meet demanding unstructured data needs with FlashArray File Services. FlashArray supports billions of files and thousands of concurrent users. Data services such as immutable snapshots and replication provide resilience against malicious or accidental data destruction and simplify disaster recovery.

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.

---

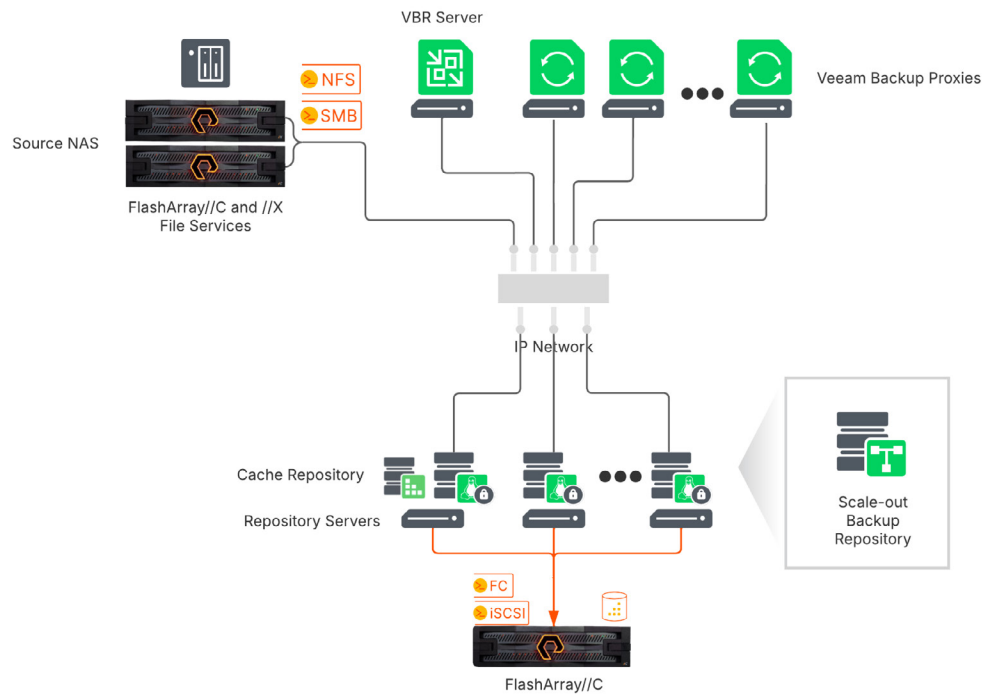
## How to Use This Guide

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 proofs of concept, 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 protecting 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 System (NFS) protocols, communicating with the Veeam 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, which are dynamically distributed across the backup proxies.

## Components

**Pure Storage FlashArray File Services:** Real-time Enterprise File on the Pure Storage platform can meet demanding unstructured data needs with FlashArray File Services. 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 V12.3, and it is applicable for other V12 releases as well.

**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 backup repository. They also collect metadata about files and directories to track changes between backups.

Veeam designates proxies for particular workloads. In VBR V12, proxies can service VMware, Hyper-V, or in this case file shares.

**Backup repository:** The Veeam Backup Repository provides the storage for backed up data and can reside on block or object 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 [Veeam and Pure Storage Security Blueprint](#) reference architecture.

**File shares:** VBR tracks source file data in its inventory as network shares, which map to FlashArray managed directory exports. Network shares define the 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 file 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
Backup Proxy	Windows Server 2016 or newer, or Linux	1 core per concurrent task	4GB minimum 4GB per additional concurrent task
Repository Server	Hardened Linux distribution supported by Veeam	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 in the same amount of time. On some storage platforms, you can gain performance by using scale-out repositories spread across multiple servers and repository 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 backup proxy 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, such as handshakes, file locks, and opening or closing a file. 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 protocol overhead to make SMB 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 both 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 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.

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 that a backup or restore session can use 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 at a time can run backup tasks for that share, regardless of how many proxies you have and how many backup sessions you run.

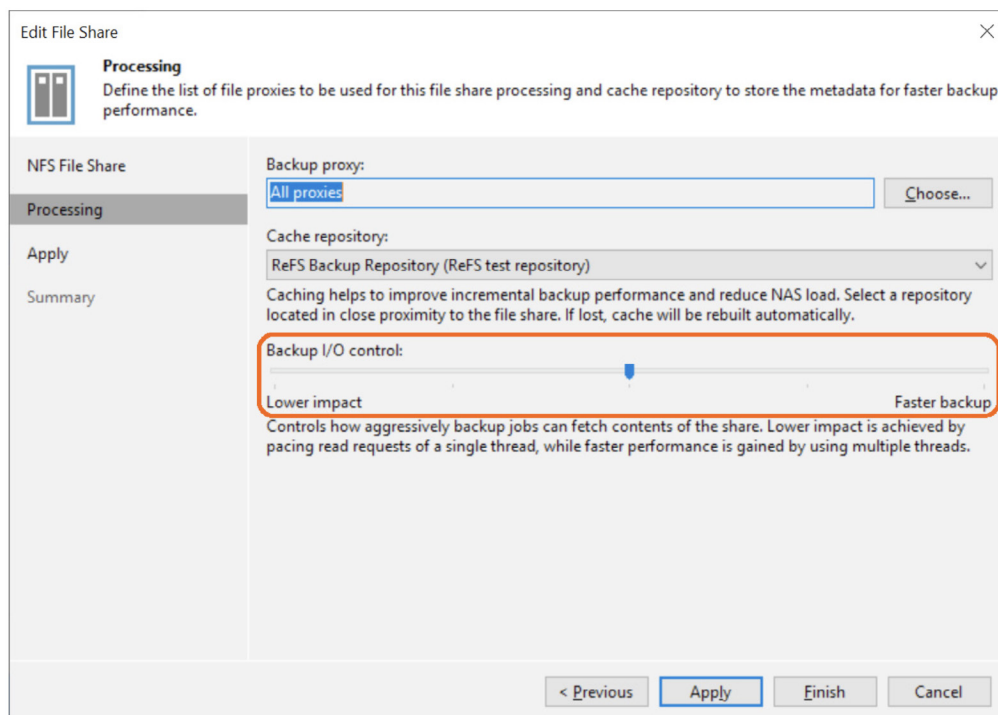


FIGURE 2 Veeam file share processing settings



## Scaling Recommendations for Veeam Backup Proxies

When deciding how many Veeam backup proxies to deploy, bear in mind that:

- To provide redundancy, you should start with at least two backup proxies for any environment size. Three proxies should support over 3.5TiB/hour, depending on the proxy networking and the data profile. Add proxies based on your 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 Real-time Enterprise File into your existing VBR environment, you need to perform three main steps: deploy backup proxies, create network shares that represent the FlashArray, and configure one or more backup jobs.

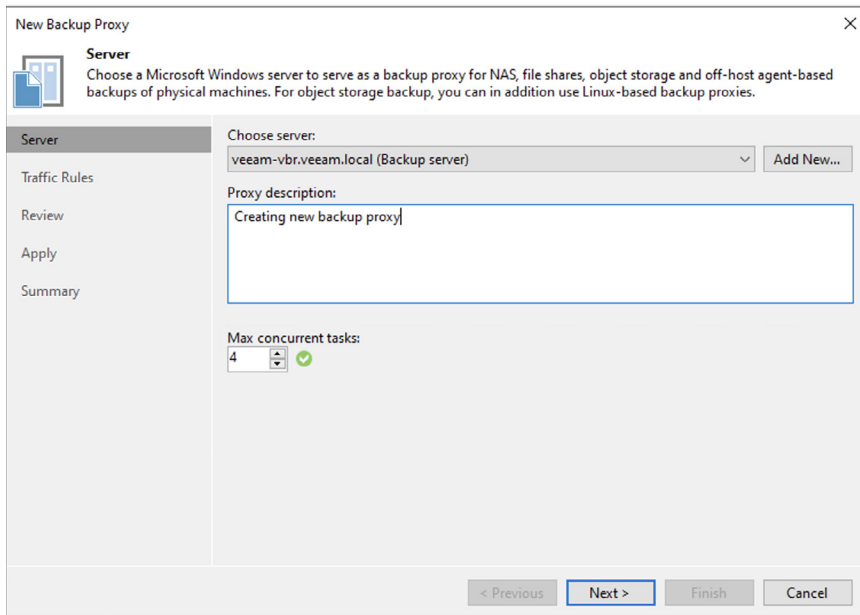
**NOTE:** This document assumes you have a FlashArray joined to Active Directory and hosting data and you have configured a Veeam Backup Repository. Alternatively, you can create an array local user for Veeam to use. Refer to the [FlashArray Admin Guide](#) for instructions on managing local users.



## Deploying Backup Proxies

Veeam backup proxies are simply managed servers with a Veeam proxy agent installed and designated for backups of NAS, file shares, and object storage. For this verification, Windows servers were used as backup proxies. 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 backup proxy for file share:

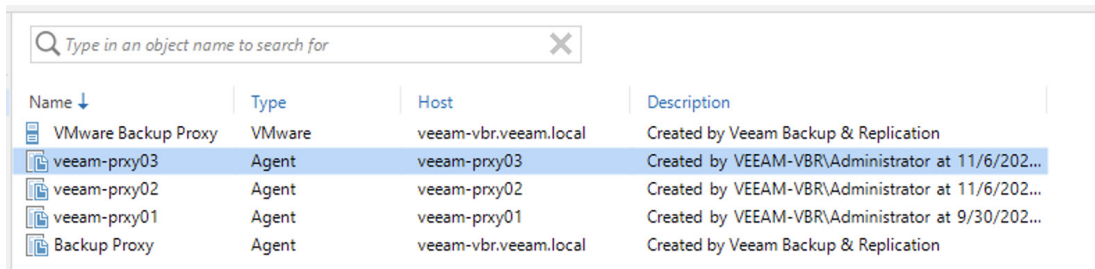
1. Navigate to the Backup Infrastructure view. To open the New Backup Proxy wizard, from the Backup Proxy menu, click the **Add Proxy** button.
2. As Figure 3 shows, on the Server page, select the server that will act as a backup proxy. If the proxy is a new server not currently managed by VBR, click 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.



The screenshot shows the 'New Backup Proxy' wizard, specifically the 'Server' page. The left sidebar contains links for 'Server', 'Traffic Rules', 'Review', 'Apply', and 'Summary'. The main area has a 'Choose server:' dropdown menu with 'veeam-vbr.veeam.local (Backup server)' selected. Below this is a 'Proxy description:' text box containing 'Creating new backup proxy'. At the bottom, the 'Max concurrent tasks:' is set to 4, indicated by a green checkmark. Navigation buttons at the bottom include '< Previous', 'Next >', 'Finish', and 'Cancel'.

FIGURE 3 New Backup Proxy wizard

Once the wizard is complete, you will see the proxy listed along with its Agent type (Figure 4).



Name ↓	Type	Host	Description
VMware Backup Proxy	VMware	veeam-vbr.veeam.local	Created by Veeam Backup & Replication
veeam-prxy03	Agent	veeam-prxy03	Created by VEEAM-VBR\Administrator at 11/6/202...
veeam-prxy02	Agent	veeam-prxy02	Created by VEEAM-VBR\Administrator at 11/6/202...
veeam-prxy01	Agent	veeam-prxy01	Created by VEEAM-VBR\Administrator at 9/30/202...
Backup Proxy	Agent	veeam-vbr.veeam.local	Created by Veeam Backup & Replication

FIGURE 4 Newly added backup proxy

Repeat the procedure for each backup proxy.

## Adding Network Shares

Before backing up a FlashArray 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 **Unstructured Data** option and select **File Shares**. 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.

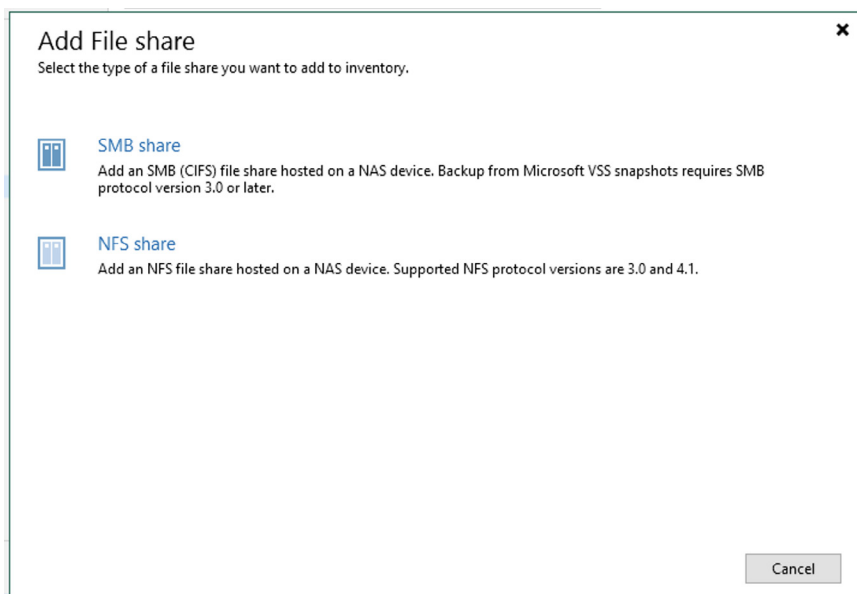


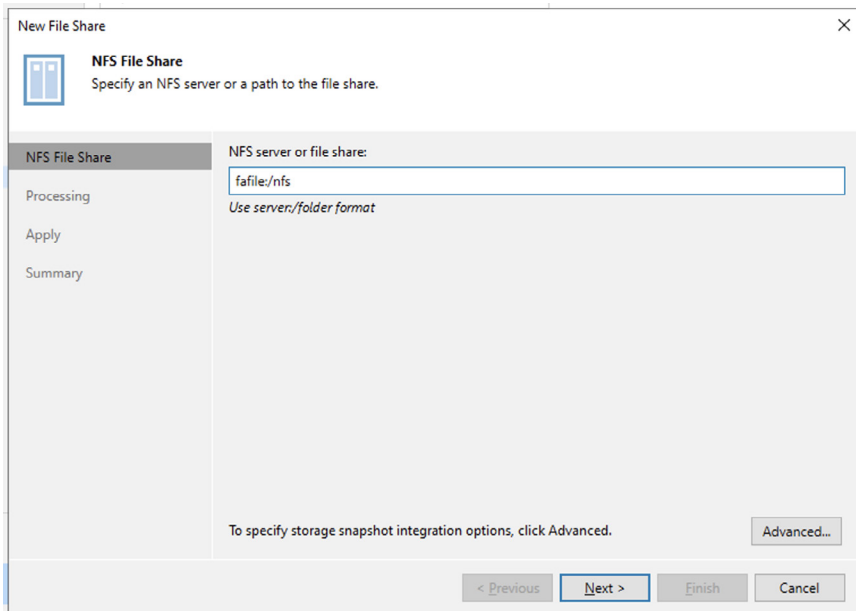
FIGURE 5 Add File share window

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 hostname, file virtual interface (VIF) IP address, or fully qualified domain name (FQDN).



New File Share

**NFS File Share**  
Specify an NFS server or a path to the file share.

NFS File Share

Processing

Apply

Summary

NFS server or file share:  
faffle:/nfs  
Use server/folder format

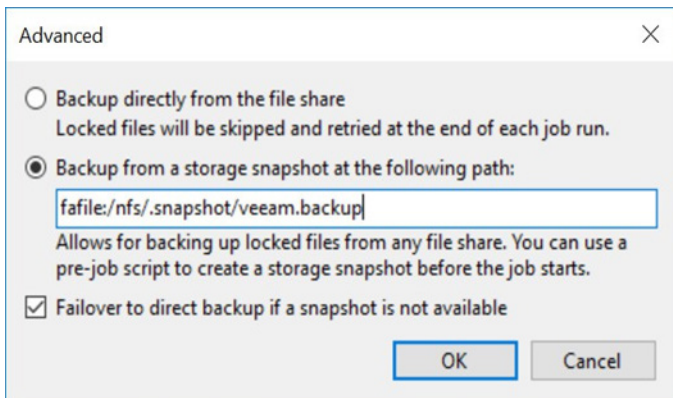
To specify storage snapshot integration options, click Advanced.

Advanced...

< Previous Next > Finish Cancel

FIGURE 6 NFS File Share page

If you want to back up from a storage snapshot instead of from the live file system, click the **Advanced** button. Enter the path to the storage snapshot in NFS format (Figure 7). We recommend configuring a protection policy specifically for file system backups that creates a daily snapshot just before your backup window begins, retains the snapshot for one day, and uses a static client name and suffix. This approach avoids the frequent configuration changes required when using variable snapshot names, ensuring a consistent, automated backup for your file systems. See the [Protect Open Files](#) section for more information on using FlashArray file snapshots.



Advanced

☐ Backup directly from the file share  
Locked files will be skipped and retried at the end of each job run.

☒ Backup from a storage snapshot at the following path:  
faffle:/nfs/.snapshot/veeam.backup  
Allows for backing up locked files from any file share. You can use a pre-job script to create a storage snapshot before the job starts.

☒ Failover to direct backup if a snapshot is not available

OK Cancel

FIGURE 7 NFS File Share—advanced options

- 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 shared 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.

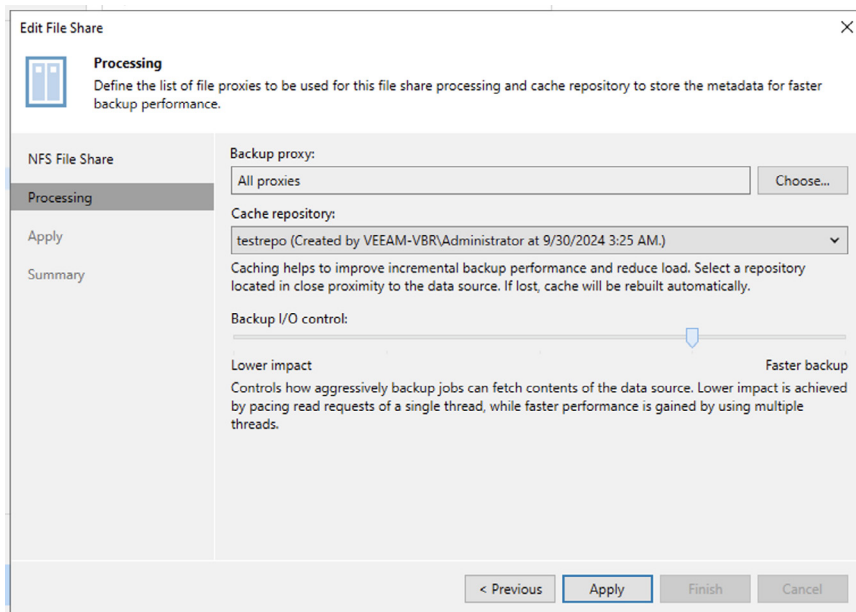


FIGURE 8 NFS File Share—Processing page

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

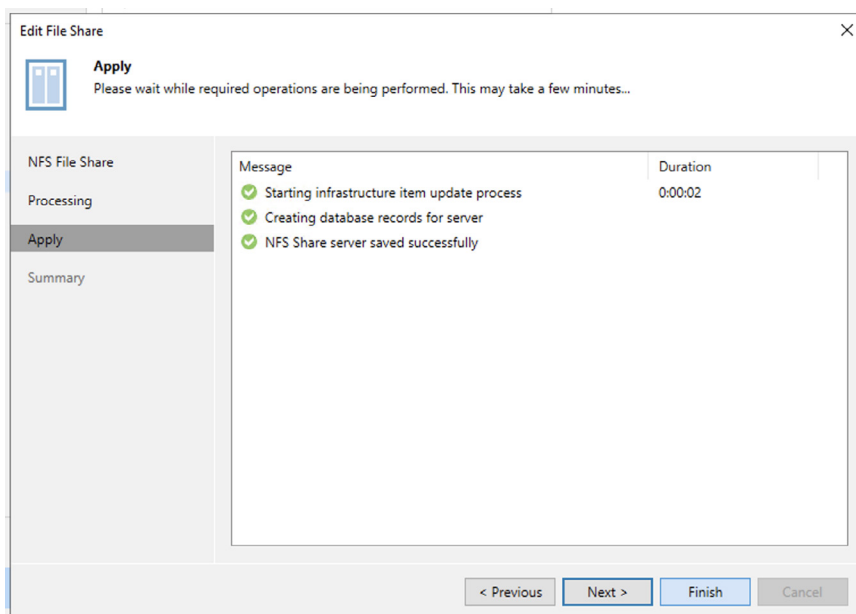


FIGURE 9 NFS File Share—Apply page

## Adding an SMB File Share

To add an SMB file share to the VBR inventory:

1. 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 hostname, virtual IP address, or FQDN.

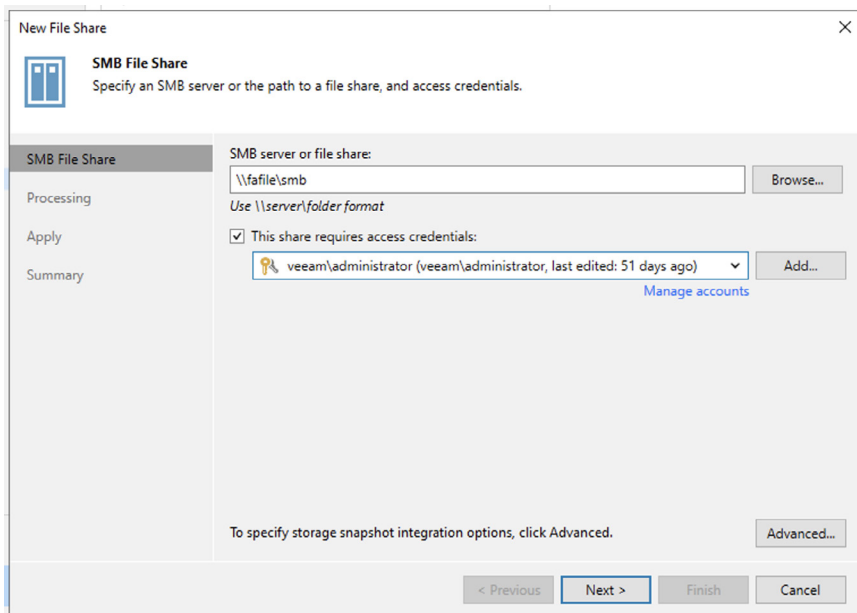


FIGURE 10 SMB File Share page

If you want to back up from a storage snapshot instead of from the live file system, click the **Advanced** button. Enter the path to the storage snapshot in UNC format (Figure 11). As with NFS file shares, we recommend creating a daily snapshot on the FlashArray that uses a static client name and suffix to simplify open file protection. See the [Protect Open Files](#) section for more information on using FlashArray file snapshots.

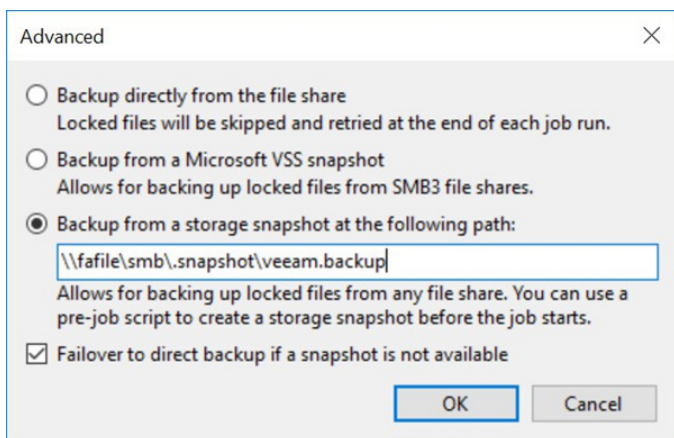


FIGURE 11 SMB File Share—advanced options

- 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 shared 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.

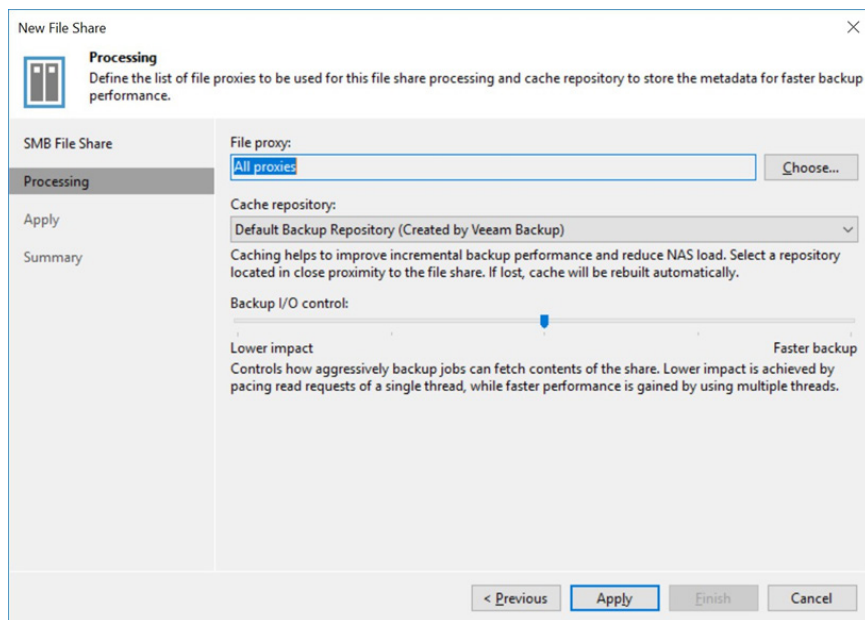


FIGURE 12 SMB File Share—Processing page

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

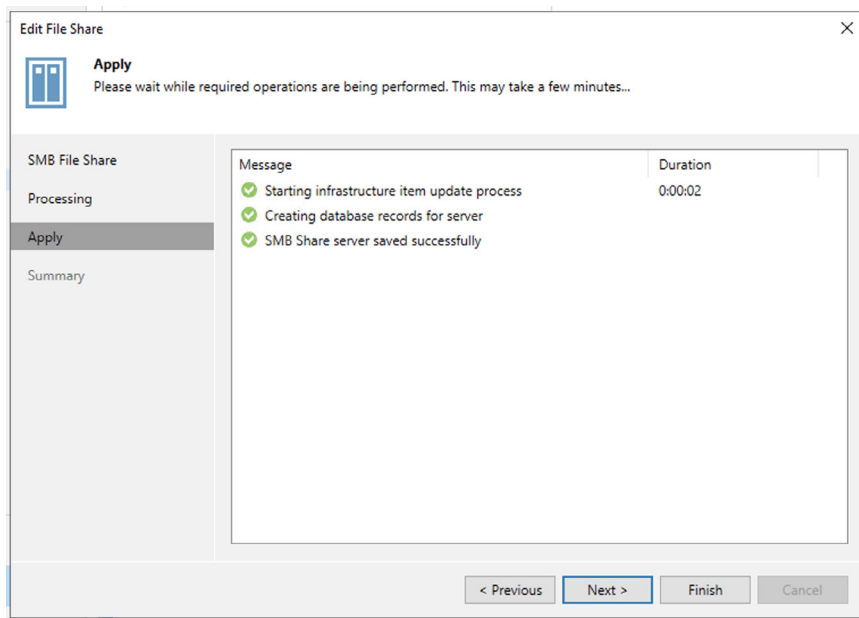


FIGURE 13 SMB File Share—Apply page

### Creating Backup Jobs

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

1. 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

3. On the Objects 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—Objects page





4. As Figure 16 shows, from the **Server** drop-down, 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.

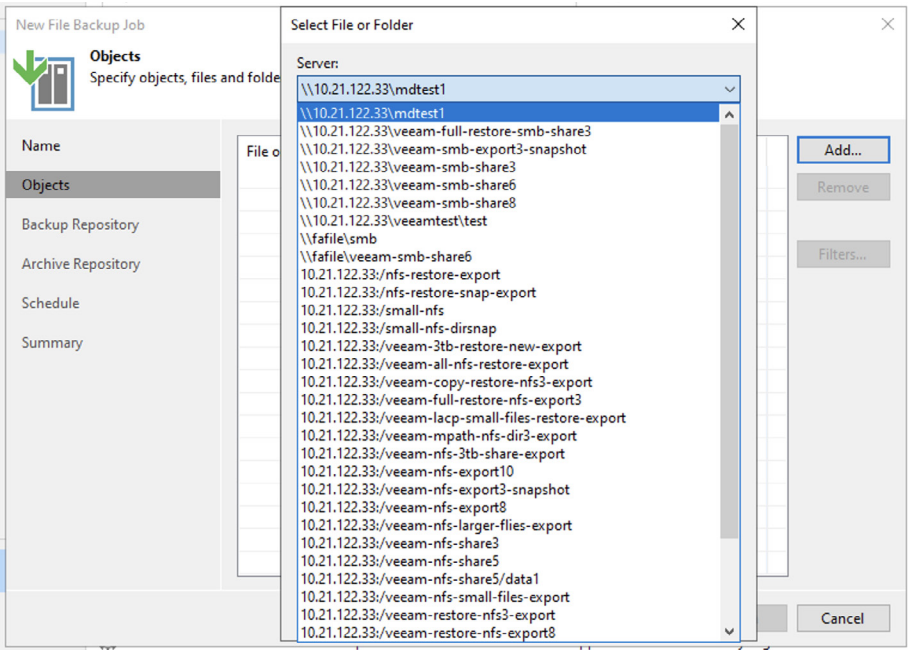


FIGURE 16 Select File or Folder dialog

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 **Filters** button.

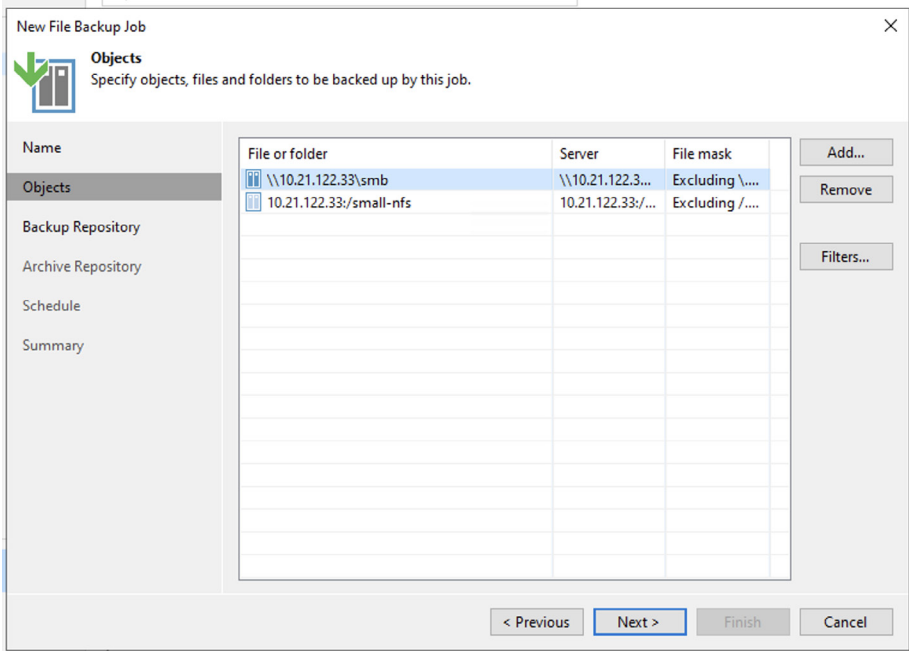
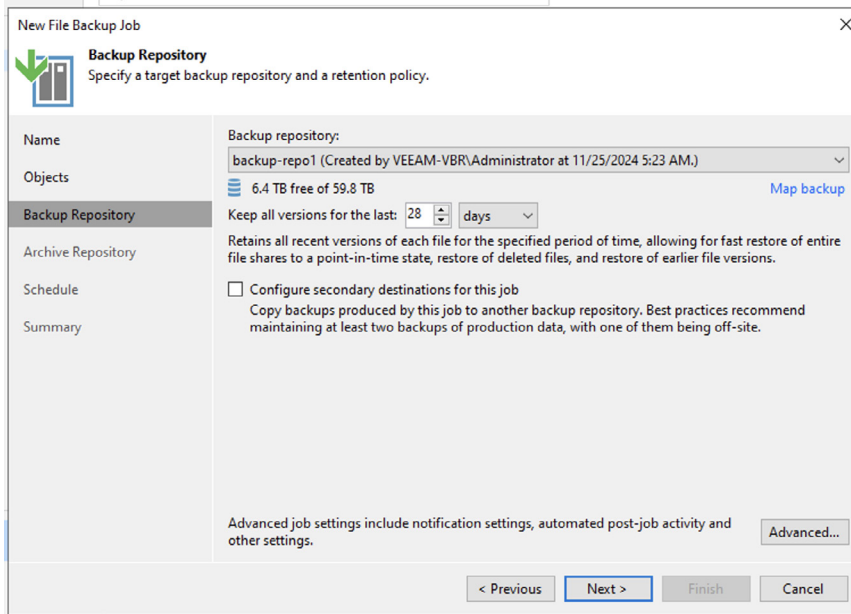


FIGURE 17 New File Backup Job—file or folders selection page

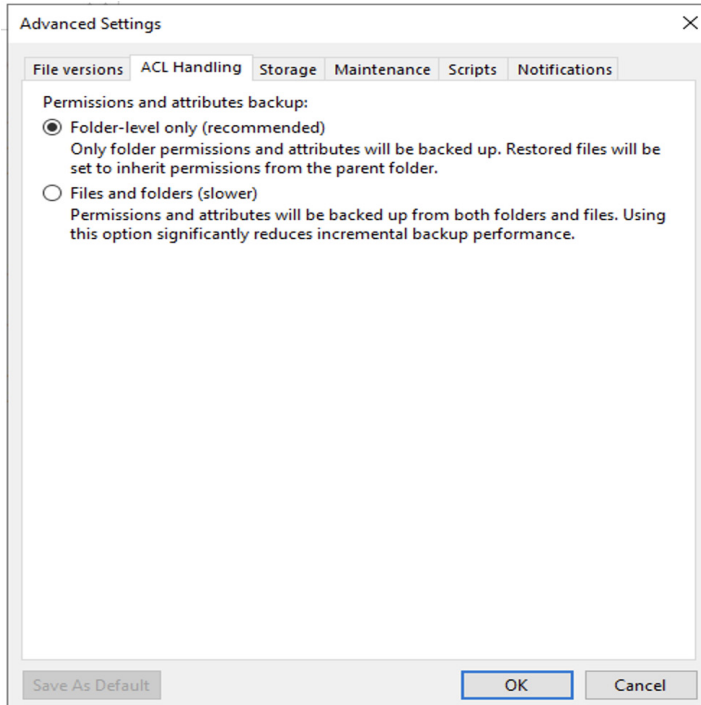
6. On the Backup Repository page (Figure 18), select the backup repository to store the backup data. You can set the backup retention and, optionally, add secondary destination settings. Click the **Advanced** button to access settings for handling ACLs, scripts, and other advanced options.



The screenshot shows the 'New File Backup Job' window, specifically the 'Backup Repository' tab. The window has a sidebar on the left with tabs: Name, Objects, Backup Repository (selected), Archive Repository, Schedule, and Summary. The main area is titled 'Backup Repository' with the subtitle 'Specify a target backup repository and a retention policy.' It contains a dropdown for 'Backup repository:' set to 'backup-repo1 (Created by VEEAM-VBR\Administrator at 11/25/2024 5:23 AM.)', showing '6.4 TB free of 59.8 TB' and a 'Map backup' link. Below this is a 'Keep all versions for the last:' field set to '28' days. A note states: 'Retains all recent versions of each file for the specified period of time, allowing for fast restore of entire file shares to a point-in-time state, restore of deleted files, and restore of earlier file versions.' There is an unchecked checkbox for 'Configure secondary destinations for this job' with a description: 'Copy backups produced by this job to another backup repository. Best practices recommend maintaining at least two backups of production data, with one of them being off-site.' At the bottom, there is an 'Advanced...' button and a note: 'Advanced job settings include notification settings, automated post-job activity and other settings.' Navigation buttons at the bottom are '< Previous', 'Next >' (highlighted), 'Finish', and 'Cancel'.

FIGURE 18 New File Backup Job—Backup Repository page

7. 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.



The screenshot shows the 'Advanced Settings' window, specifically the 'ACL Handling' tab. The window has a sidebar with tabs: File versions, ACL Handling (selected), Storage, Maintenance, Scripts, and Notifications. The main area is titled 'Permissions and attributes backup:' and contains two radio button options: 'Folder-level only (recommended)' (selected) with the description 'Only folder permissions and attributes will be backed up. Restored files will be set to inherit permissions from the parent folder.' and 'Files and folders (slower)' with the description 'Permissions and attributes will be backed up from both folders and files. Using this option significantly reduces incremental backup performance.' At the bottom, there are buttons for 'Save As Default', 'OK' (highlighted), and 'Cancel'.

FIGURE 19 New File Backup Job > Advanced Settings > ACL Handling tab

8. 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.

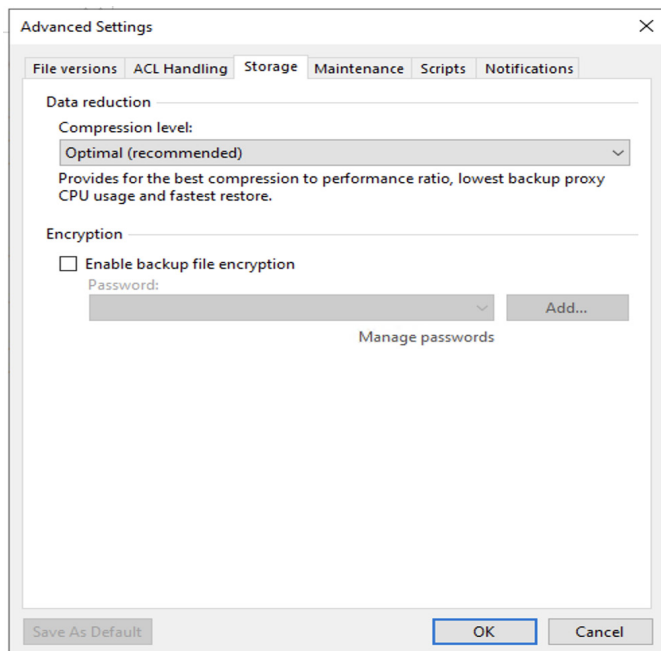


FIGURE 20 New File Backup Job > Advanced Settings > Storage tab

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 storage snapshots for open file protection. For more information on protecting open files using storage snapshots, refer to the [Protect Open Files](#) section.

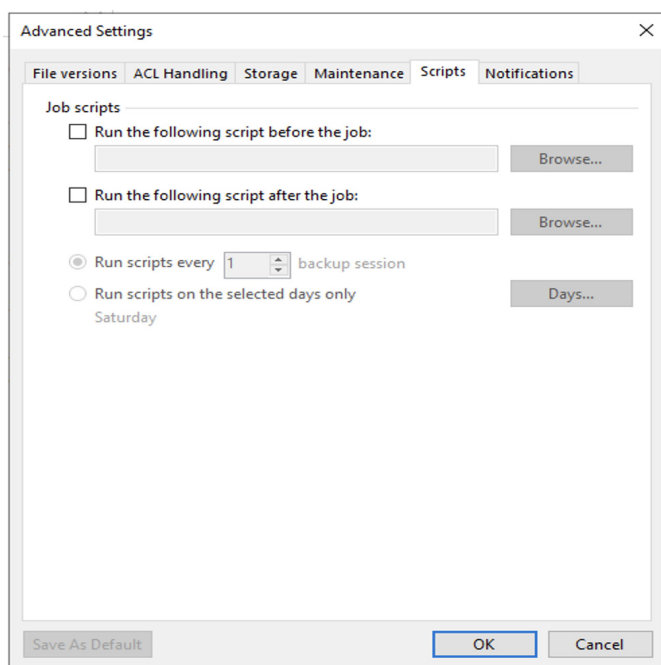


FIGURE 21 New File Backup Job > Advanced Settings > Scripts tab

10. You can optionally configure the archive repository (Figure 22).

FIGURE 22 New File Backup Job—Archive Repository page

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

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.

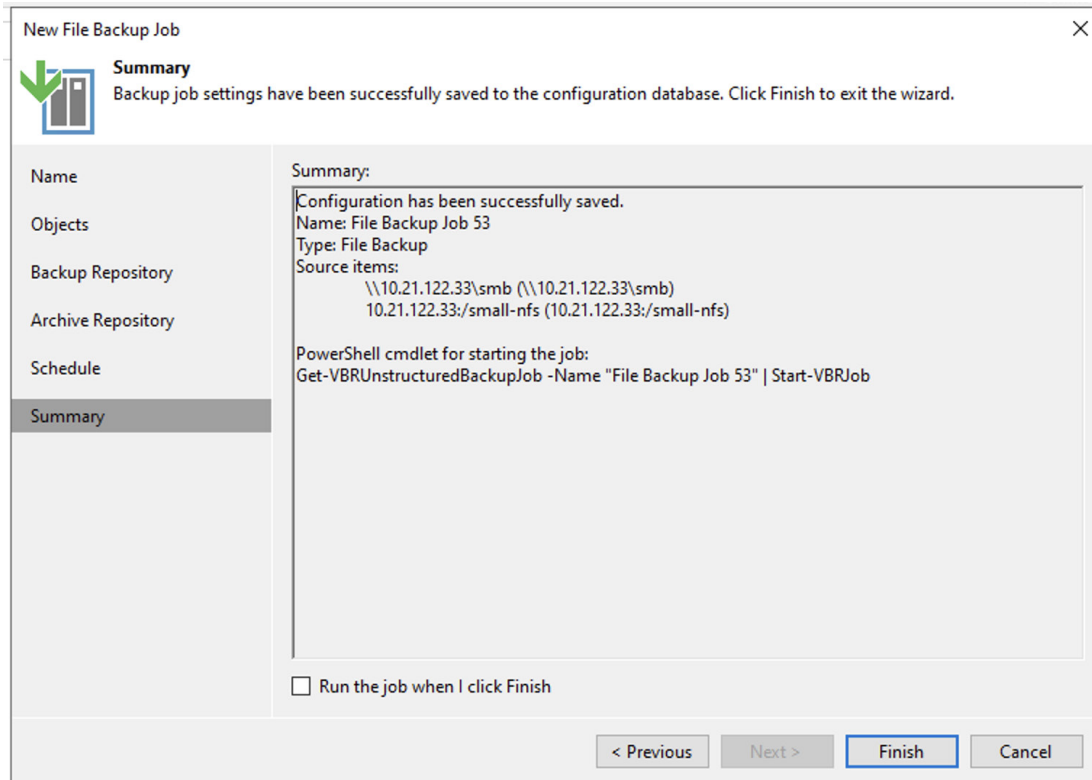


FIGURE 24 New File Backup Job—Summary page

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. Client systems can't delete managed directories, and administrators cannot delete them from the FlashArray management interface unless they are empty.

## 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 it 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>
```

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

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

To output all exports for that file system, use:

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

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 Backup Repositories

You should follow Veeam's [guidance](#) on configuring repositories. If you use Pure Storage FlashArray//C for repository storage, you should review the best practices in the [Veeam and Pure Storage Security Blueprint](#) reference architecture.

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 Storage FlashArray//C to reduce the data size, you should enable compression in the backup job but set the repository to decompress data before writing.

### Anti-malware 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 significantly. 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 anti-malware 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:

- Reducing 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 of 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.

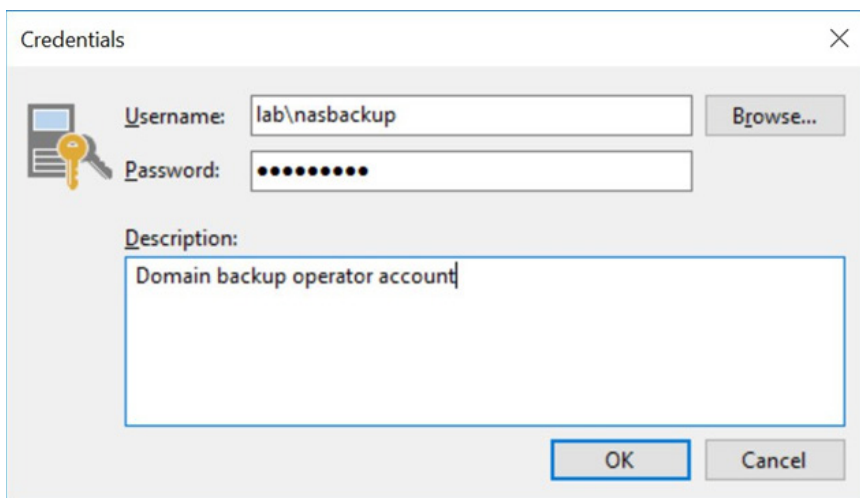


FIGURE 25 Credentials dialog

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

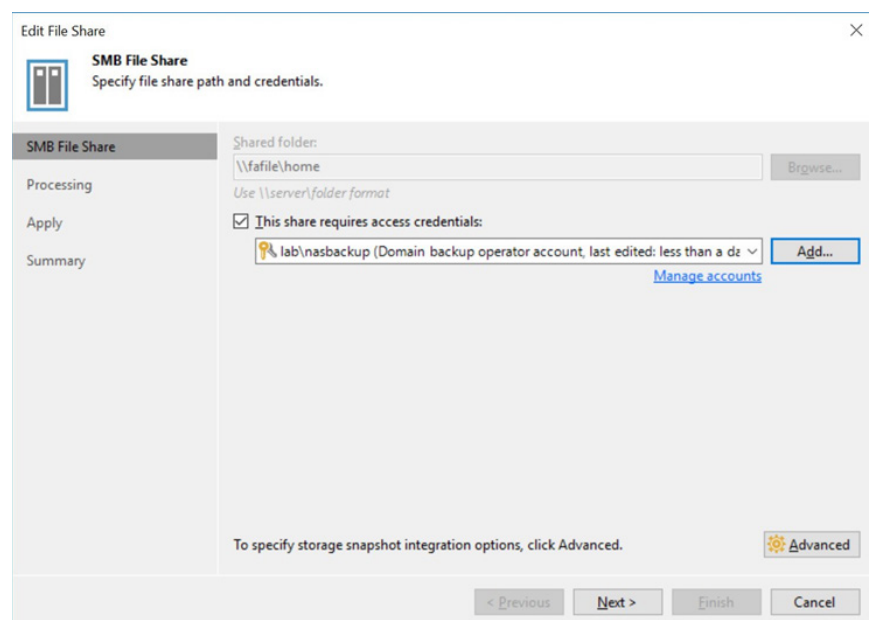


FIGURE 26 Edit File Share—SMB File Share page

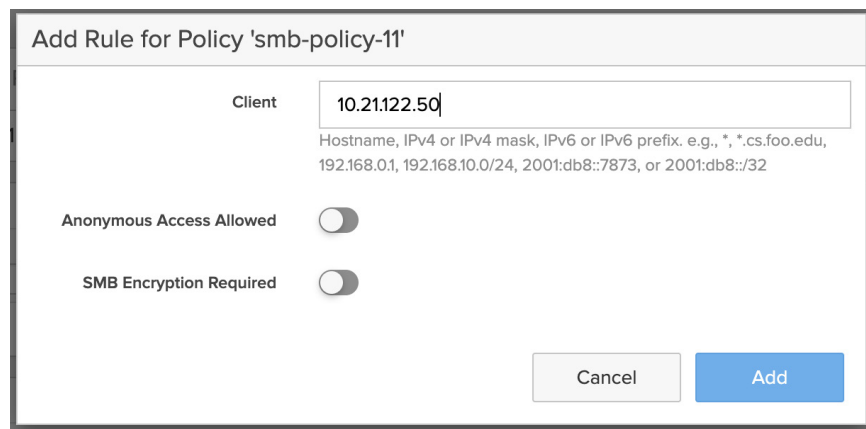
You can change the password in VBR by selecting it in the Manage Credentials dialog and clicking the **Edit** button, then entering the new password in the **Password** field.

**NOTE:** You can use an array local user instead of Active Directory if you prefer.

## Restrict SMB Export Access

You should use restrictive SMB 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.
- Based on need, the **Anonymous Access Allowed** option can be set.
- If encryption is needed, the **SMB Encryption Required** option can be enabled.



Add Rule for Policy 'smb-policy-11'

Client

Hostname, IPv4 or IPv4 mask, IPv6 or IPv6 prefix. e.g., \*, \*.cs.foo.edu, 192.168.0.1, 192.168.10.0/24, 2001:db8::7873, or 2001:db8::/32

Anonymous Access Allowed ☐

SMB Encryption Required ☐

Cancel Add

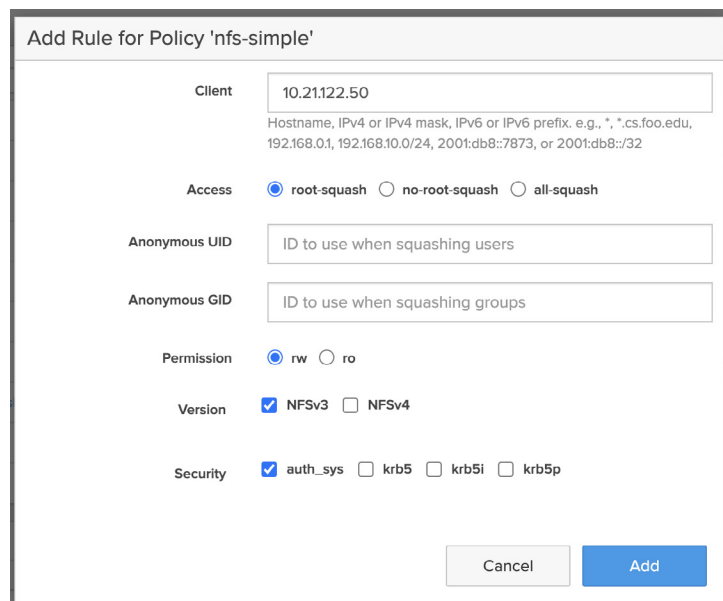
FIGURE 27 FlashArray SMB export rule definition

## 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 28):

- 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.
- Necessary NFS version and security configuration can be set as shown.



Add Rule for Policy 'nfs-simple'

Client: 10.21.122.50  
Hostname, IPv4 or IPv4 mask, IPv6 or IPv6 prefix. e.g., \*, \*.cs.foo.edu, 192.168.0.1, 192.168.10.0/24, 2001:db8::7873, or 2001:db8::/32

Access: ☒ root-squash ☐ no-root-squash ☐ all-squash

Anonymous UID: ID to use when squashing users

Anonymous GID: ID to use when squashing groups

Permission: ☒ rw ☐ ro

Version: ☒ NFSv3 ☐ NFSv4

Security: ☒ auth\_sys ☐ krb5 ☐ krb5i ☐ krb5p

Cancel Add

FIGURE 28 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 storage 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 storage snapshots of managed directories, including file system roots, capturing the state of the entire directory structure at the time of the snapshot. Storage 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 **fafile**, a file system called **veeam-validation-fs** might have a managed directory named **veeam-validation-dir**. If the snapshot client name is **veeam** and the suffix is set to **dp**, the snapshot name would be **veeam-validation-fs:veeam-validation-dir.veeam.dp**. Figure 29 shows the management view of the managed directory and snapshot.

ArrayHostsVolumesPodsFile SystemsPolicies

> File Systems > veeam-validation-fs:veeam-validation-dir

Size	Virtual	Data Reduction	Unique	Snapshots	Total
-	0.00	1.0 to 1	996.00 B	0.00	996.00 B

Policies ^1:1 of 1

Name ^	Type	Enabled
smb-simple	smb	true

Directory Exports1:1 of 1

Name ^	Status	Path	Policy	Type
sample-export	Disabled (Export Disabled)	/veeam-validation-dir	smb-simple	smb


Directory Snapshots ^1:1 of 1

Name	Policy	Created	Time Remaining
veeam-validation-fs:veeam-validation-dir.veeam.dp	-	2025-02-17 11:42:45	-

FIGURE 29 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 `\\fajfile\exportpath\.snapshot\veeam.dp`. Figure 30 shows the client view.



**FIGURE 30** Managed directory snapshot, client view

Besides the automated policies, you can also create snapshots on demand using the FlashArray GUI or CLI. When you create these manual snapshots, you can add a suffix to the name, which helps maintain a consistent, predictable naming convention for easier management.

While manual snapshots are useful for one-off needs, we recommend using protection policies to automate your snapshot strategy. For detailed best practices on how to set up and manage these policies, refer to the [General Snapshot Best Practices](#) section later in this document.

**NOTE:** When SafeMode™ Snapshots, a FlashArray feature that mitigates against ransomware attacks on file system data, is enabled, the FlashArray periodically creates immutable snapshots of all file systems and prevents manual eradication of file systems and snapshots.

## 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 31) 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.

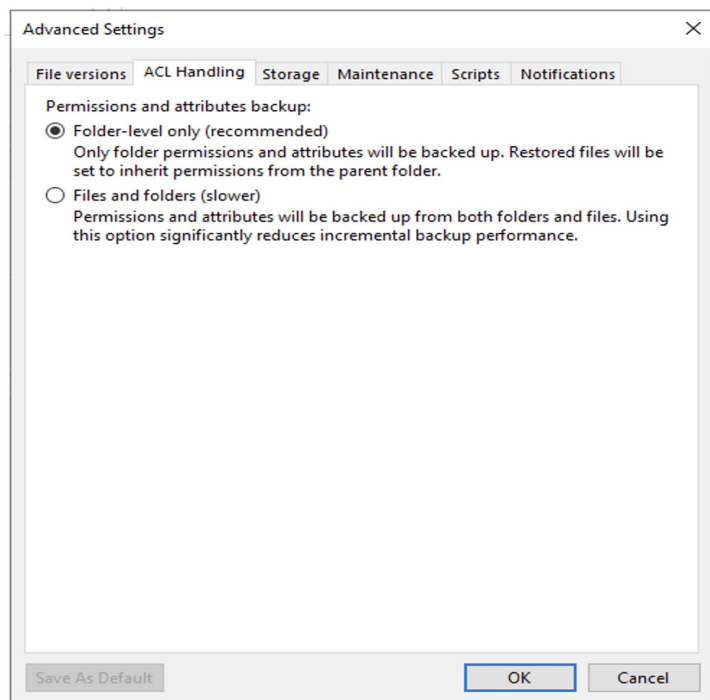


FIGURE 31 ACL handling option



## General Snapshot Best Practices

### Create Storage Snapshots on File Share Managed Directories

You can create storage 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 Protection Policy for Storage Snapshots

Protection policies define how often to create snapshots, how long to keep them, and how to name them. For each managed directory on the FlashArray that you want to use snapshots to protect, you must add the directory to the policy as a member. Policies can have multiple member directories; all will follow the same schedule, retention, and snapshot naming. The key point to note is that the retention can't be longer than the snapshot frequency.

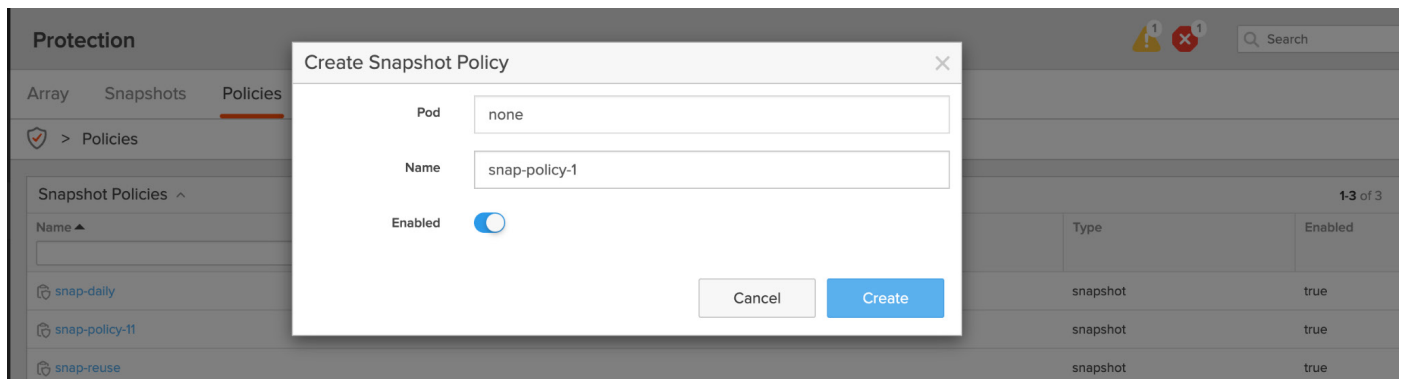


FIGURE 32 Snapshot policy creation

Under member selection, click the more options button, then click **Add Member**.

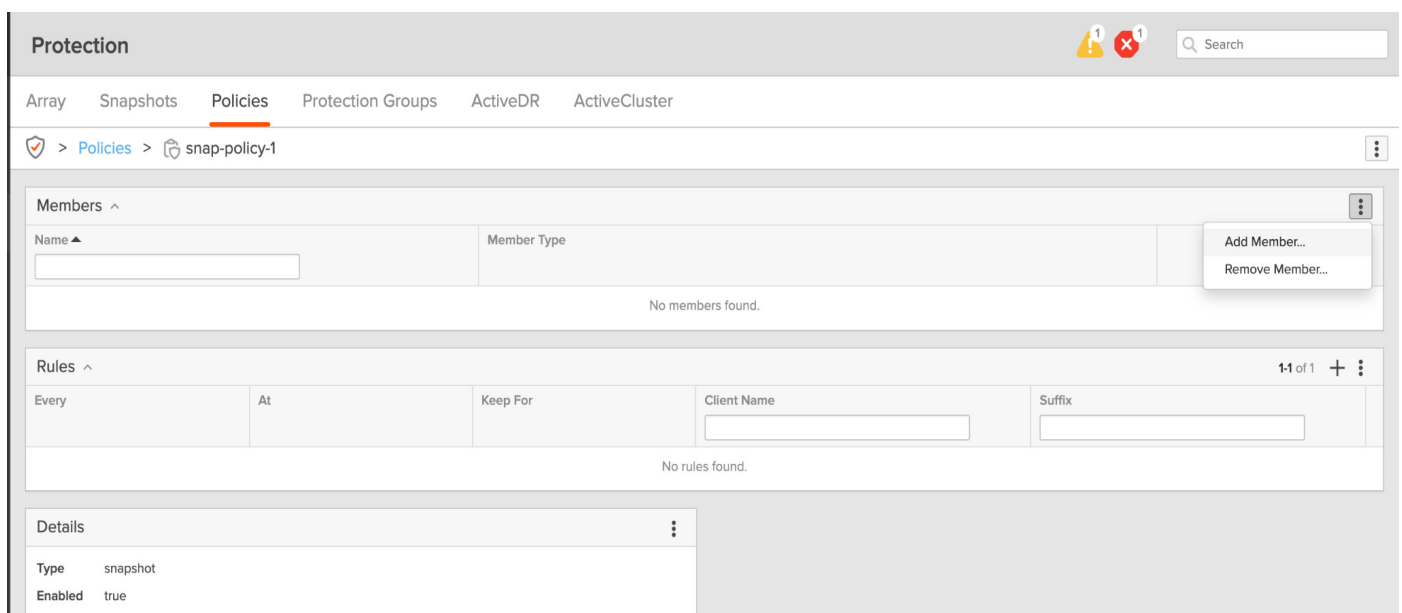


FIGURE 33 Protection policy view

Add the directory to the snapshot policy. All member directories share the same schedule, retention, and snapshot naming.

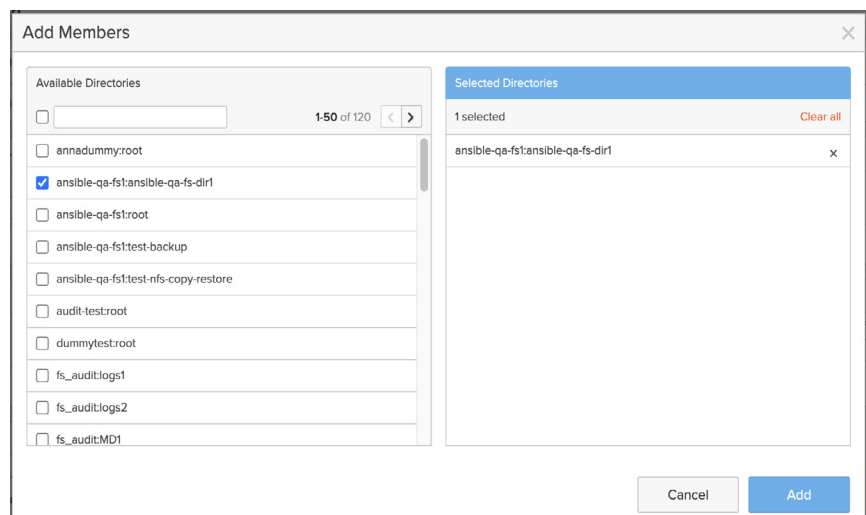


FIGURE 34 Member addition to protection policy

The snapshot rule automates your data protection by defining the schedule, retention, and naming for snapshots within a protection group. For your specific needs, the rule would be configured as follows:

- **Schedule:** Take a snapshot daily at 6:00 PM.
- **Retention:** Keep each snapshot for one day before it is automatically eradicated.
- **Naming convention:** Snapshots will be named with the prefix vbr and the suffix bkp.

This rule ensures a consistent, automated backup for your volumes and file systems, creating a new recovery point every day at 6:00 PM and holding it for 24 hours.

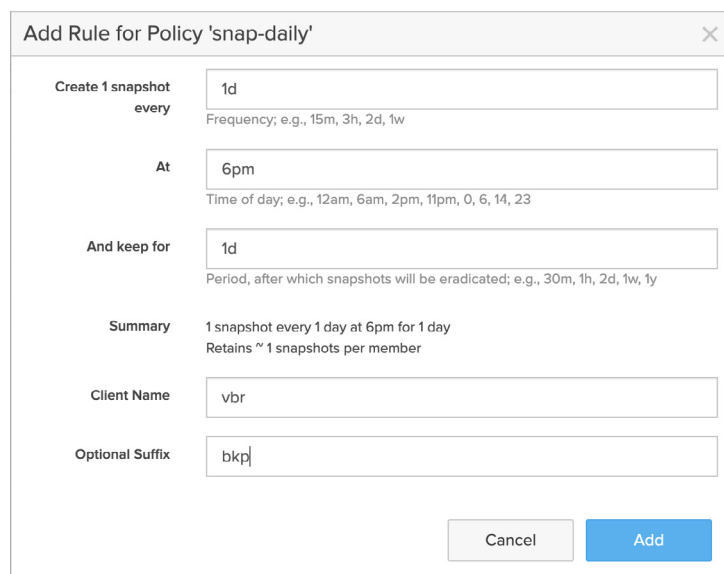


FIGURE 35 Snapshot policy rule definition

## 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.

## Conclusion

Real-time Enterprise File on the Pure Storage platform can meet demanding unstructured data needs with FlashArray File Services. FlashArray File Services arranges the data in file systems and managed directories, which have policies applied to them.

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 Storage 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 Storage [unified block and file services](#).
- Find out how to deploy Pure Storage as a Veeam repository in the [Veeam and Pure Storage Security Blueprint](#) reference architecture.