COHESITY

Version 1.1

April 2020

Best Practices for NAS Data Protection with Cohesity

Streamline NAS Backups with Cohesity

ABSTRACT

This guide outlines best practices for protecting NAS devices with Cohesity's flexible and scalable NAS backup solution.

Table of Contents

Figures

Tables

Introduction to NAS Data Protection

With network-attached storage (NAS) devices, organizations can store and serve business-critical data. However, NAS devices are vulnerable to virus attacks and hardware crashes. As any threat of data loss affects the business's sustainability and continuity, it becomes imperative to protect this data.

Cohesity provides a platform that eliminates the complexities and operational inefficiencies of traditional NAS protection solutions by unifying your end-to-end data protection and recovery infrastructure — including target storage, backup, replication, disaster recovery, archiving, and cloud tiering. Cohesity eliminates the need for traditional data protection and recovery silos by converging all your backup infrastructure on a single scale-out platform.

Cohesity's web-scale architecture allows your NAS data protection to grow with your NAS infrastructure. Even as it handles billions of files and folders on your NAS devices, Cohesity's indexing engine ensures you will be able to find your 'needle in a haystack' quickly when you need to restore only specific files and folders. When you need to recover from a large-scale disaster or loss of data, or when you need to migrate data stores from one data center to another, you can recover all your NAS data within seconds as a Cohesity View.

Cohesity's Solutions for Protecting NAS

Using Cohesity to protect NAS devices brings with it a host of features that extend well beyond backup. When you use our solution to back up your NAS device's SMB/CIFS shares and NFS mounts, you can combine Cohesity Protection Policies with Protection Groups to replicate your data off-site, archive and tier it to lower-cost storage in the cloud, and send it to tape or the cloud for long-term retention and disaster recovery. What's more, the Cohesity CloudArchive Direct feature allows you to archive your NAS data to cloud storage directly, without having to store it on-premises first.

Cohesity gives you the power to combine your business and legal requirements (Protection Policy) with operational flexibility (Protection Group) to protect your NAS data (Source entities), giving you rich flexibility and granular control over your NAS data protection strategy.

Figure 1: Protect Your NAS Data with Cohesity

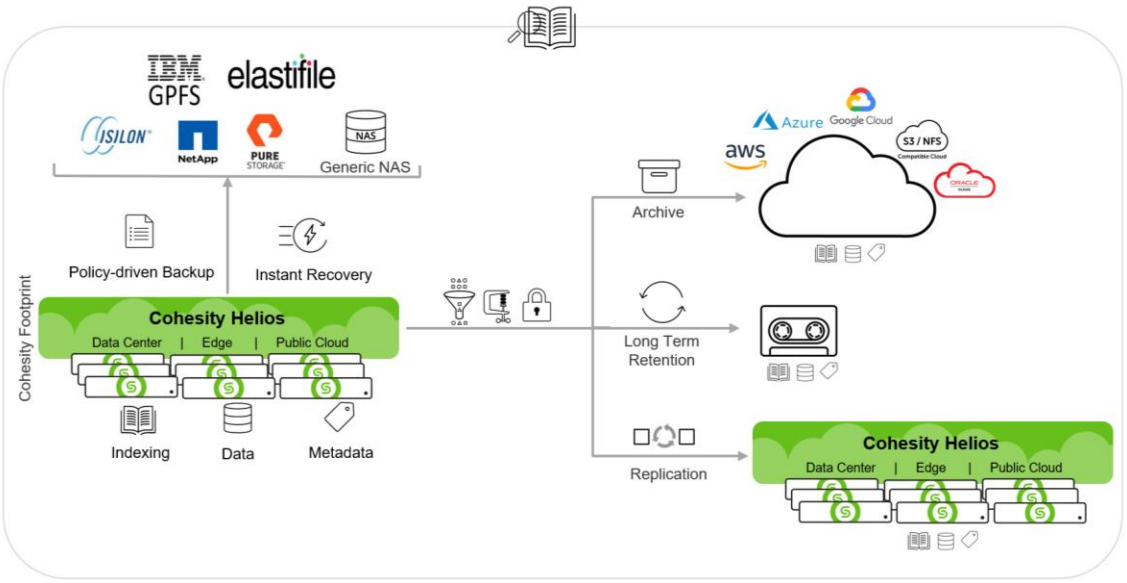

CloudArchive Direct

To protect NAS data at scale without the need to store the actual data on-premises, Cohesity has built CloudArchive Direct for NAS, a cost-efficient solution that processes and sends the data directly to lower-cost storage on External Targets using object store in the public/private cloud or over NFS. By eliminating the need to store a copy locally before archiving, the footprint/capacity requirements of Cohesity cluster are dramatically reduced. Only the metadata and indexes, which enable quick search and recovery, are stored on the local Cohesity cluster. In this solution, the entire NAS dataset (the data along with metadata and indexes) is stored only on the External Target.

CloudArchive Direct is a policy-driven feature with seamless integration with all major cloud vendors like AWS, Azure, GCP, Oracle, or any S3-compatible object store. It can also be configured with on-premises compression and encryption to achieve maximum storage efficiency and security.

NOTE: On-premises encryption and compression is supported in Cohesity version 6.5 and higher.

Figure 2: Making NAS Data Archival Cost-effective with Cohesity's CloudArchive Direct

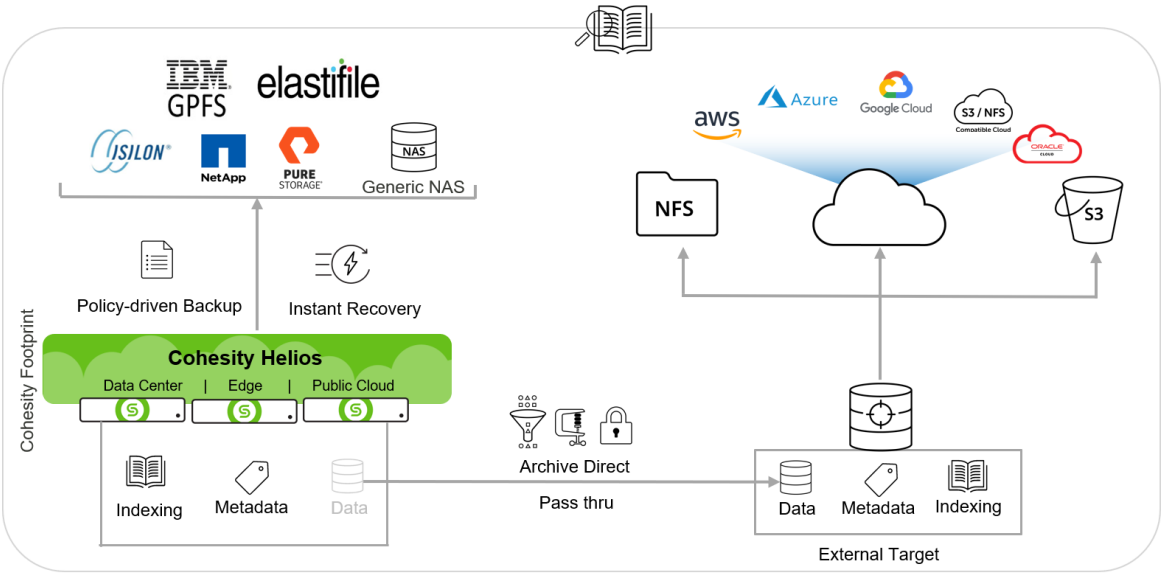

Protect NAS with Cohesity

Protecting NAS devices with Cohesity is a straightforward process that involves a few tasks:

1. Create or configure a user on your NAS device with the minimum backup and restore permissions.
2. Use that user to register the NAS device as a source on Cohesity Platform.
3. Choose (or create) a Cohesity Protection Policy.
4. Create a Cohesity Protection Group with the NAS data objects you need to protect.
5. Search and restore your protected NAS data.

Figure 3: Set Up NAS Data Protection with Cohesity

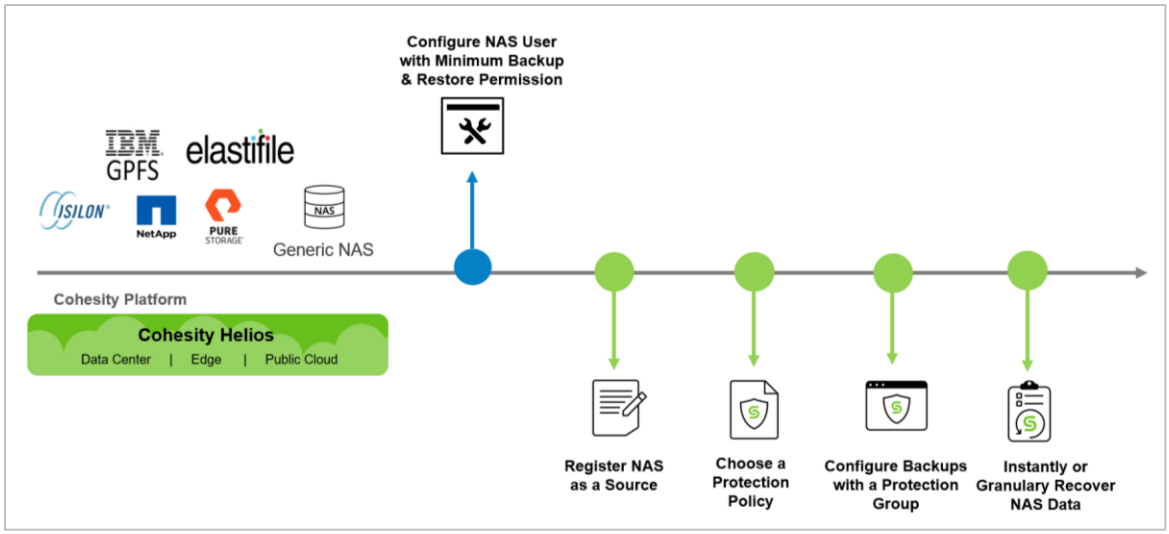

Register Your NAS with Cohesity

Cohesity offers two approaches to protect your NAS data:

- API integration (snapshot-based backup)
- Without API integration for all NAS (mount-point-based backup)

For snapshot-based backup, Cohesity provides native integration using each major NAS vendor's native integrated APIs to protect the NAS data using snapshots. The vendor-specific Cohesity NAS adapters for this integration enable Point-in-Time (PIT) snapshots and change file tracking. Depending on the level of integration in place with different NAS vendors, Cohesity uses Native Change File Tracking or the File Runner for incremental backups. For instructions, find your vendor in Register or Edit NAS in the online Help.

For all other environments, you can use an SMB/CIFS or NFS mount path to register your NAS device — the vendor-agnostic, or 'Generic NAS,' approach. For more, see **Mount Point** in Register or Edit NAS in the online Help.

There are tradeoffs between the two approaches, as illustrated in Figure 4 below.

Figure 4: Cohesity NAS Data Protection Approaches

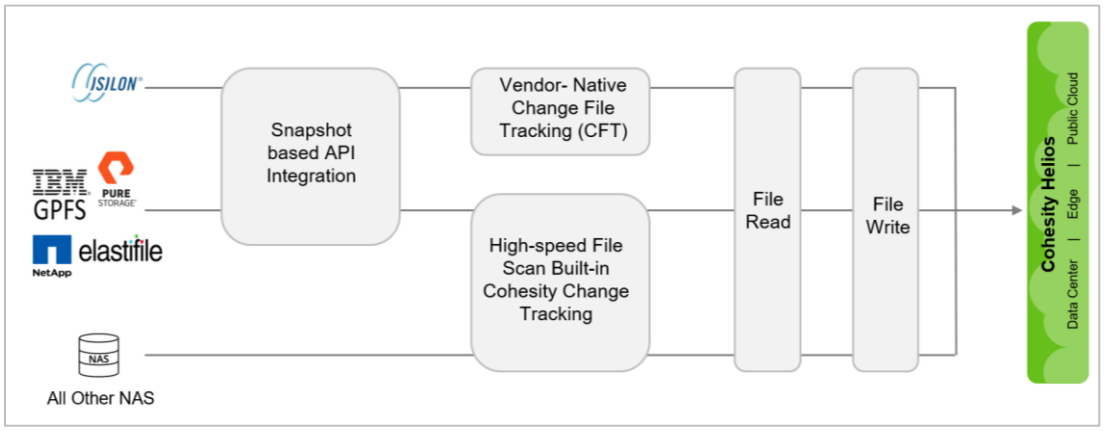

NAS Source Considerations for Cohesity

As you set up to protect your NAS devices with Cohesity, consider these configuration and best practice recommendations for best results:

- If you use the generic, vendor-agnostic integration, all vendors and versions of NAS are supported as long as they support NFSv3 or SMB1,2.x. If you are using a vendor-specific adapter, refer to Table 1 to ensure you have the minimum required version.

Table 1: Supported NAS and Cohesity Versions and Adapters

COHESITY NAS ADAPTER TYPES	VENDOR VERSION	COHESITY VERSION
Isilon OneFS	8.0.x, 8.2.x, 9.x*	6.3.1b+
NetApp ONTAP	8.2+, 9.1, 9.2, 9.3, 9.5, 9.6, 9.7*, 9.8*, 9.9.1*	4.1+
IBM GPFS	5.0.2.0+	6.4.1+
Pure FlashBlade	2.2.4+	6.1.1+
Elastifile EFS	3.1.0.31+	6.3.1b+
Generic NAS	All vendors supporting NFS (v3) and SMB (v1+)	Any

NOTE:

- *Isilon OneFS 9.x is supported on Cohesity version 6.5.1f onwards.
- NetApp 9.7 is supported on Cohesity version 6.5.1b onwards.
- NetApp 9.8, 9.9.1 is supported on Cohesity version 6.5.1d onwards.

Please consult Cohesity Support to verify proper permissions are set.

- After ensuring that the source and the target are on supported software versions, open the recommended firewall ports between your NAS device and Cohesity cluster to enable network communication between source and target.

  NOTE: For more on the specific firewall ports used by a Cohesity cluster, see Open Firewall Ports.

- Ensure the NAS volumes to be backed up are supported.

Table 2: Supported NAS Volumes

NAS STORAGE	VOLUME/DIRECTORY TYPE	VOLUME SUB-TYPE	SUPPORTED
NetApp	Flex Volume	Normal Flex Volume	Yes
		SnapLock Enterprise Volume	Yes
		SnapLock Compliance Volume	Yes
		Encrypted Volume Storage	Yes
		Encrypted Volume (Using KMS)	No
	FlexGroup Volume	N/A	No
	Data Protection Volume	SnapMirror Destination Volume **Supported with Cohesity version 6.4.1c and 6.5.0a+.*	Yes
		SnapVault Destination Volume **Supported with Cohesity version 6.4.1c and 6.5.0a+.*	Yes
Isilon	Non WORM	N/A	Yes
	Data Protection Directory	SyncIQ Destination Directory	Yes
	WORM	SmartLock Enterprise Directory	No
		SmartLock Compliance Directory	No
Pure FlashBlade	Filesystem	SMB-enabled Filesystem	Yes
		NFS-enabled Filesystem	Yes
		HTTP-enabled Filesystem	No
Elastifile	Instance	NFS Instance	Yes
IBM GPFS	Fileset	Independent Fileset ** Only NFS based fileset is supported.*	Yes
		Dependent Fileset	No

- Set up the user account with backup and restore privileges that will be used for registering the NAS storage device with Cohesity. Cohesity recommends you create a custom role with all the requisite privileges and then assign the created role to the user. Cohesity supports both local and domain users.

  NOTE: For information on backup and restore privileges, see Ensure Adequate Privileges for Cohesity on the Source in the online Help.

- Create SMB/CIFS shares or NFS exports on the NAS volumes that will be backed up.
- Cohesity DataProtect uses the native source array snapshot capabilities in NAS to snapshot the volumes. Point-in-Time snapshots are taken and mounted locally for faster backups. To ensure that snapshots are captured successfully, make sure that:

  a) Your snapshot license is enabled on the NAS device.

  b) The snapshot reserve is set to a non-zero value.

  c) The `.snapshot` directory is visible.

  d) Enough space is available on your NAS device for snapshot operations.

NAS Data Backup Considerations

To ensure a smooth setup, follow the guidance and best practices below for configuring Cohesity for NAS data protection.

- While registering NAS with Cohesity:
  a) Ensure the Host IP or Name allows data access via NAS protocols (SMB, NFS).
  b) Ensure the user account has the required privileges. If you plan to protect NAS SMB shares, use the local or Active Directory user credentials that allow at least read access to the SMB share and full access to recover.
  c) If you use a domain username, use the ***<domainname>******<username>*** format.
  d) For multi-tenancy, register your NetApp Vservers instead of a NetApp cluster.
- Cohesity includes three standard Protection Policies: **Gold**, **Silver**, and **Bronze**. For their default settings, see Manage Policies in the online Help. If the default settings of the standard Protection Policies do not meet your needs, you can create a customized Protection Policy.
- While creating a Protection Group:
  a) You can select the default SMB share or NFS export for the NAS device (for example, **`/ifs`** for Isilon) as an object in your Protection Group, but Cohesity recommends that you select a specific share/export, to have better control over access and share settings. This will also help NAS data backup operations to use the same protocol as the protocol that is used to access the data, as multiprotocol NAS volumes/directories can be backed up using only one protocol, SMB or NFS. In addition, this approach helps keep your backup environment manageable.
  b) Avoid adding the same volumes in multiple Protection Groups.
  c) Cohesity recommends to turn on **AutoProtect** to protect new NetApp volumes that are added to a selected parent Vserver or cluster automatically.
  d) Isilon Changelist enables OneFS to compare the snapshot from a prior successful backup to the snapshot taken for the current backup to quickly find changes without performing a lengthy file scan of the targeted directory structure. This can significantly improve backup times when the data change rate is low. To leverage the Isilon Changelist API for faster incremental backup, enable **Use Isilon Change List** in your Protection Group. However, if this option is enabled, NAS backups might take longer where the data change rate is high.
  e) Select the appropriate **QoS Policy** for your operational requirements. There are two default QoS Policies: HDD or SSD.
    i. Backup HDD (default): The Cohesity cluster writes the data directly to HDD for this Protection Group. Cohesity recommends HDD.
    ii. Backup SSD: The Cohesity cluster writes the data directly to SSD for this Protection Group. Only choose this QoS policy if you need fast ingest for a small number of Protection Groups.

f) Use Pre & Post Scripts to perform customer-driven processing tasks. **Pre script** is executed before a Protection Group runs to perform specific *before-backup* operations. **Post script** is executed after a Protection Group run, to perform specific *after-backup* operations. If configured, scripts are run every time an object is backed up by a Protection Run.

   i. If the business requirement is to abort the backup on the first occurrence of any error during a Protection Run, disable the **Skip File on Errors** option.
   ii. If possible, do not include or exclude files from backup. Including or Excluding files can affect backup performance, due to the filtering overhead during file scans.
   iii. To be able to recover at granular levels later, enable **Indexing** while creating a Protection Group.
   iv. Update the default **SLA** values according to your business requirements.
   v. For multi-tenancy setup, create different Protection Groups for different Vservers.

- If you have Komprise-enabled NetApp Volumes for stub files (offline files), then starting with Cohesity version 6.4.1, you can enable the internal configuration on Cohesity Platform to back up stubs created with Komprise on a NetApp volume without rehydrating the data. Any other stub files, created by other means, are backed up by rehydrating the data. Contact Support to configure custom settings.

NAS Data Recovery Considerations

To ensure a smooth NAS data restoration operation with Cohesity DataProtect, follow these best practices and guidance:

1. Ensure the user account that is configured in Cohesity Platform during registration has full control on the target volume where SMB restore is being performed.
2. In cases of file- and folder-level recovery:
   a) If the target is case-insensitive, file and directory names with case differences will be treated like files and directories of the same name. This can cause one file or directory to overwrite another during the restore.
   b) Ensure there is no directory resting on the target of the same name as the file to be recovered. Otherwise, the target location directory will be replaced with the file being recovered.
   c) If **Overwrite Existing File/Folder** is enabled, directories with the same name are merged and files with the same name are overwritten.
3. Before you can recover storage volumes, ensure that snapshots of those volumes exist on the Cohesity cluster.
4. If you want to repurpose the data without time delays (when creating a dev/test environment, etc.), use the recovery option **Restore as Cohesity View**. When recovering to a Cohesity View, Cohesity DataProtect clones the selected backup to a new View within the Cohesity cluster and provides you instant access to it.
5. In cases of data unavailability or loss at the source, set the target to **Original NAS Location** to recover the data back into the source from which it was backed up.
6. Recover to an alternate NAS location under circumstances such as source unavailability or source data migration. In such a case, set the target to **Alternate NAS Location**.
7. If you want to recover at the file or directory level, use the **Recover > File and Folder** option.

Performance Recommendations

To achieve the best possible performance, follow these best practices:

1. Consider the impact of the NAS source directory structure on backup performance:
    a) If you are backing up multiple directories that contain small files, consider setting up a separate schedule for each directory, as too many read/write operations can create a bottleneck for backup performance.
    b) Limit the depth of nested subdirectories in your file system.
    c) Limit the number of files in a directory. Distribute files across multiple directories instead of including a large number of files in a single directory.
2. Configure multiple Protection Groups when configuring NAS backups, with each Protection Group capturing a portion of the volumes and/or directories. Cohesity doesn't recommended to attempt to back up the entire NAS storage device through a single Protection Group, as it can affect backup performance.
3. Use different Protection Policies with Protection Groups to run the backup tasks at different times, so that they do not overlap. Running Protection Runs simultaneously can affect the performance of either or both.
4. Bandwidth is an important component in overall backup and restore. Cohesity recommends that you use a minimum of a pair of 10 GbE connections between the NAS device and Cohesity Platform, which provides good bandwidth and resilience.
5. Tune NAS backup performance by enabling the internal configuration as described in the NAS performance tuning Support article.

Resources

Find more information about Cohesity's solutions for NAS in these reports and solution briefs:

- Modernizing NAS Backup and Recovery (third-party report)
- Enterprise Data Protection for Pure Storage FlashBlade with Cohesity NAS Protect (webinar)
- Cohesity NAS Protect for Pure FlashBlade (solution brief)
- Protecting NetApp or any NAS with Cohesity DataProtect (video)
- NAS Data Protection (solution brief)

Your Feedback

Was this document helpful? Send us your feedback!

About the Authors

Ruby Garg is a Technical Marketing Engineer at Cohesity. In her role, she focuses on NAS and VMware Data Protection.

Document Version History

VERSION	DATE	DOCUMENT HISTORY
1.0	Jan 2020	Original document — internal release
1.1	Apr 2020	First release

ABOUT COHESITY

Cohesity radically simplifies data management. We make it easy to protect, manage, and derive value from data -- across the data center, edge, and cloud. We offer a full suite of services consolidated on one multicloud data platform: backup and recovery, disaster recovery, file and object services, dev/test, and data compliance, security, and analytics -- reducing complexity and eliminating mass data fragmentation. Cohesity can be delivered as a service, self-managed, or provided by a Cohesity-powered partner.

Visit our website and blog, follow us on Twitter and LinkedIn and like us on Facebook.