



TECHNICAL WHITEPAPER

Protecting NAS at Scale with Rubrik

TABLE OF CONTENTS

AUDIENCE	3
EXECUTIVE SUMMARY	3
NDMP IS NOT THE ANSWER	3
A Short History of NDMP	3
NDMP Limitations.....	4
PROTECTING NAS WITH RUBRIK.....	5
Three-Phase Approach	5
Scan	5
Scan with No API Integration.....	6
Scan with Snapshot API Integration	8
Scan with Snapshot and File Change APIs Integration	9
Fetch	10
Copy	10
Rubrik Direct Archive.....	11
Global Search and Rapid Recovery	12
BENEFITS OF USING RUBRIK.....	13
CONCLUSION	13
ABOUT THE AUTHORS	13

AUDIENCE

This white paper is intended for Data Center Architects, Systems Administrators, Storage Administrators, and Backup Administrators who have responsibility for protecting and recovering Network Attached Storage (NAS) data. To meet the demands of today's enterprise data growth, Rubrik provides an alternative approach to legacy NAS protection solutions. This paper will help users make an informed decision about what approach to take in managing their NAS environments by walking through the innovations that Rubrik has invested in our data management solution.

EXECUTIVE SUMMARY

Data has been growing exponentially and will continue to do so for the foreseeable future. Today, much of that data lives in enterprise NAS environments that are growing at a rapid rate and protecting this data requires a next-generation data management solution that is designed and built to protect terabytes to petabytes of unstructured data.

Rubrik offers a next-generation solution that meets the requirements of today's growing enterprise data protection needs. Customers using Rubrik can realize the benefits of reliable backups, rapid recovery of data, and the flexibility of a heterogeneous NAS data management solution.

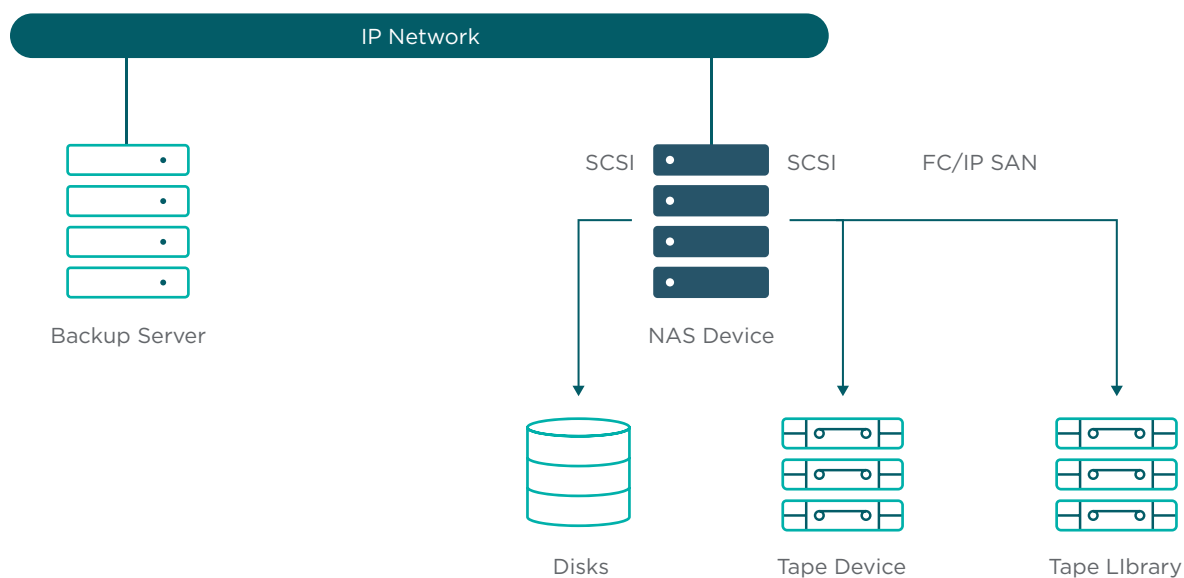
NDMP IS NOT THE ANSWER

The Network Data Management Protocol (NDMP) is currently the de facto approach for protecting NAS data. However, an increasing number of enterprises are exploring new approaches to protecting their data as their environments grow and they are experiencing the challenges and limitations associated with NDMP.

A SHORT HISTORY OF NDMP

More than two decades ago, NAS pioneer NetApp and backup vendor Intelliguard collaborated together to try and solve an issue that was becoming increasingly vexing for NAS users - the inability to reliably protect their data. Up to that point, NAS platforms such as NetApp Filers were being backed up by having backup servers mount NAS shares and then moving the backup data to locally attached tape devices or to a networked tape library. This solution was fraught with problems ranging from low performance due to the need to read every file over a POSIX interface, performance bottlenecks created by having to send data over a single mount point, and the complexity of having to manage multiple devices.

NDMP was first proposed by NetApp and Intelliguard in 1995 to address these challenges in protecting NAS platforms. NDMP is officially defined as an "open standard protocol for network-based backup for network-attached storage." The protocol specifies a separation of the control path from the data path. Control traffic passes from the backup application to the NAS platform over an IP network while data traffic flows from the NAS platform to a storage medium over SCSI or over a Storage Area Network (SAN). NDMP also defines a mechanism for allowing a backup application to initiate and manage backup jobs running across multiple NAS devices. Each NAS device would then be responsible for preparing its files to be protected and for copying files directly to a locally attached or network attached storage device.



The advantages of NDMP included the following:

- Backup traffic could be offloaded to locally attached or fibre channel attached devices, avoiding the bottlenecks created when trying to stream data across what was then typically 100 mb networks.
- Eliminated the need for data management vendors to write special agents or unique device drivers to be installed for each NAS vendor's solution. Each NAS vendor would build interfaces that adhered to the NDMP standard.
- Centralized data management so that a single instance of a backup application could initiate and manage data protection for multiple NAS devices.

NDMP LIMITATIONS

Two decades later, NDMP is still the primary solution offered by most data management vendors. While it has been the standard approach for NAS protection for over two decades, NDMP has limitations that has only become more pronounced as the amount of data to be protected has increased over time.

- NDMP is designed to support single stream backups for each NAS device which has created performance bottlenecks during the data transfer process, particularly as file systems has grown larger along with the number of unstructured files stored.
- To mitigate against performance bottleneck issue, resulting from single stream backups and metadata scanning of large file systems, many backup vendors using NDMP offer image-level backups. However, recovery of the entire file system is then required even when only a single file needs to be recovered.
- Designed primarily to use tape as the backup target, backups using NDMP typically require periodic full backups to limit issues with slow restore times which can occur with long chains of incremental backups.
- NDMP does not specify a data format for NAS backups, leaving that to the discretion of individual NAS vendors. This has created platform incompatibilities that prevent data protected from one NAS platform to be recovered to a different NAS platform. Customers are effectively locked in to a specific NAS platform and even if they choose to move to another platform, customers will likely need to maintain

some instances of the previous platform in case older files have to be recovered. Even if a customer commits to a single vendor, any changes in the data format from one version to another may necessitate maintaining multiple versions of the same vendor's solution.

PROTECTING NAS WITH RUBRIK


To ensure that customers are able to reliably protect their data and recover quickly from data corruption and data loss, Rubrik takes an approach that obviates the need for complex NDMP implementations.

THREE-PHASE APPROACH

NAS protection with Rubrik utilizes a three-phase approach. Each phase is optimized using modern techniques such as snapshot API integration, data partitioning, and parallel file streaming.

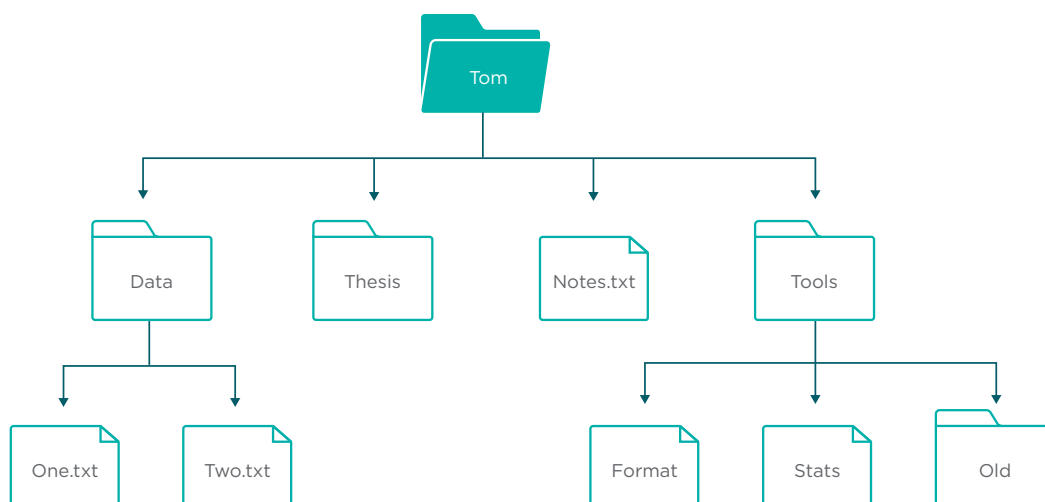
The Rubrik three-phase approach for NAS protection includes:

1. **Scan** - Rubrik identifies which files need to be protected for full or incremental backups.
2. **Fetch** - Rubrik takes the list of files from the Scan phase and reads them over the NAS protocol.
3. **Copy** - Rubrik compresses, encrypts, and writes the data to the Rubrik cluster or to an archive location on-premises or in the cloud.




Scan

At a high-level, all NAS backups begin with a scan of one or more Rubrik filesets which enables Rubrik to create a list of files that it should protect during a given backup run. A fileset is a user-specified grouping of files and can comprise of an entire NAS share or a subset of a share. Customers can use simple “include/exclude” expressions to create filesets that include only files that meet a specific parameter.



In the above example, a user can create a fileset for the entire “Tom” NAS share, a fileset for only the “Tools” folder, or a fileset that includes or excludes files with a “txt” extension.

Rubrik offers multiple options for scanning a NAS fileset. Each option offers different feature sets which are made available to data management vendors, like Rubrik, via API integration. Which option is used by Rubrik with a particular NAS vendor depends on the integration that is in place. Vendors, such as NetApp and Isilon, have made their snapshot (S) API and file change (FC) API available, enabling faster and more efficient file scanning. Other vendors, such as Pure Storage, have made their snapshot (S) API available to enable application consistent snapshots. For vendors where these APIs are not available, Rubrik has implemented an optimized approach to file scanning without leveraging API integration.

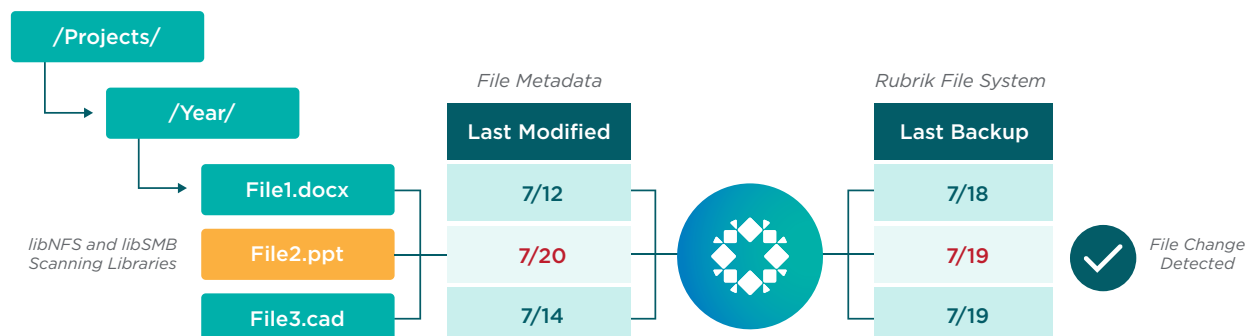


Scan with No API Integration

File scanning has historically been the biggest bottleneck to NAS backup performance. This challenge has only grown as the amount of data has grown. Previous approaches, such as image-level backups, has increased backup performance by avoiding file scanning altogether. But that approach has made file-level recovery complex and slow while often locking customers into proprietary backup formats. Modern approaches such as snapshot and file change API integration has provided much needed improvement for those NAS platforms with API integration in place with data management platforms.

When snapshot API integration is not available, Rubrik will mount the NAS share to be protected vis NFS or SMB and read the metadata of each file enumerated in a fileset. The first backup of a file system will always be a full backup which means Rubrik will scan and index every file in the fileset. For an incremental backup, Rubrik perform the following operations:


1. Scan the metadata of each file to determine its mtime (the last modified time) as well as other relevant file attributes.
2. Compare the mtime of a file against the last backup time.
3. Mark a file to be protected if the file has been modified since the last backup.



The operation Rubrik uses to check the mtime of a file is the stat() system call, which is used in Unix to determine information about a file. In a Unix operating system, stat() is implemented via the “ls” command and returns file information such as last modified time and last accessed time. Traditionally, stat() has been implemented in NAS using POSIX libraries. This has contributed to previous performance bottlenecks when protecting a NAS file system, due to inefficiencies in how file scanning is performed.


Using a POSIX library, scanning a NAS share traditionally meant that a stat() call has to be made at each level of the NAS directory tree down to the endpoint of each file. In the example below, scanning the file system requires 12 stat() calls, 4 per file.

NAS Tree



To gain greater efficiencies in the scanning operation, Rubrik implemented stat() using the LIBNFS and LIBSMB libraries. Both are part of the LIBNFS open source project and offers libraries that provides POSIX-like functions, such as stat(), that are more efficient than POSIX. Instead of stat() calls at every level of the directory tree, Rubrik only needs to do that once while scanning all the files in that same tree. In the example below, scanning the file system requires only 6 stat() calls to scan all 3 files.

NAS Tree

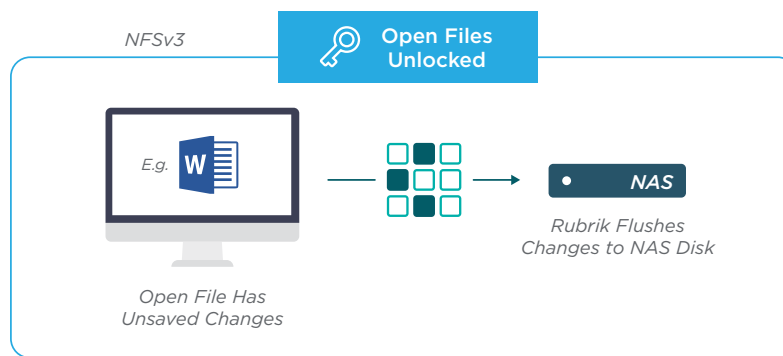


For very large NAS file shares, the performance gains can be significant given the lowered scan time and reduced overhead of the scanning operation.

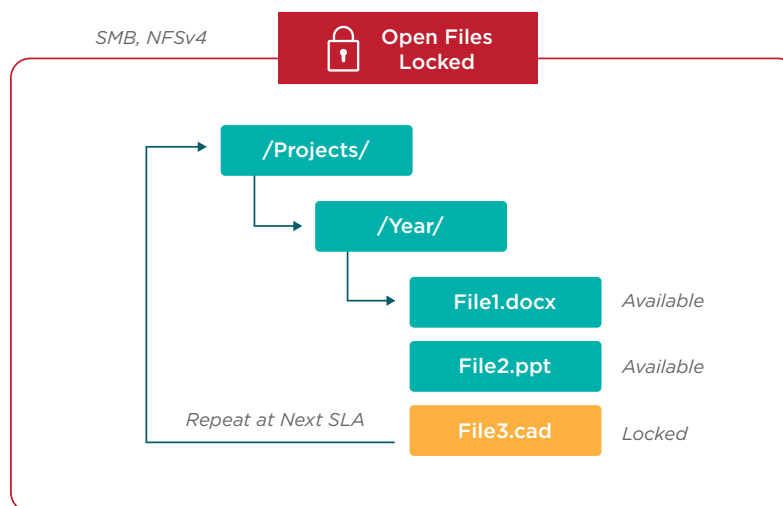
Since Rubrik may have to scan an active file system when NAS snapshot API integration is not available, there is a strong possibility that open files may be encountered. Rubrik addresses protecting open files differently depending on the file protocol used to access the share. Currently NFSv3, NFSv4, and SMB are supported.

For filesystem running NFSv3, open files are unlocked and Rubrik will protect the file in one of two ways:

1. If the application support the operation, Rubrik request any file changes stored in memory be flushed to disk before backup.
2. If the application does not support request for changes to be flushed to disk, Rubrik will ignore any file changes stored in memory and backup the version currently saved to disk.



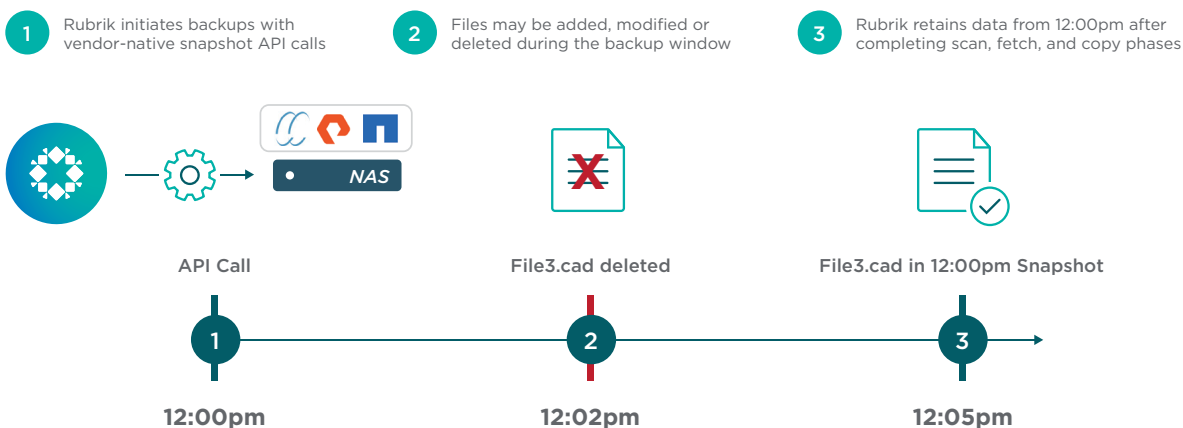
If supported by the user application, a NAS file system running NFSv4 or SMB can lock open files. When Rubrik encounters a locked file, it will bypass the file and attempt to protect it on its next backup run. If not supported by the application, an open file will remain unlocked and Rubrik will protect it in the same way as file stored using NFSv3.



Scan with Snapshot API Integration

Some NAS vendors provide a snapshot API that Rubrik has integrated with to provide additional benefits during the scan process. Prior to Rubrik performing a scan, an API call is made to the NAS platform to take a snapshot of the file system to be protected. Rubrik then mounts the snapshot, instead of the active file system, for file scanning.

By leveraging a snapshot, Rubrik is able to protect a consistent point-in-time copy of a file system without the need to use file locking. Changes stored in memory will be flushed to disk prior to the snapshot being taken. If files are deleted from or modified in the file system during a backup operation, they will be retained in the snapshot and protected by Rubrik.

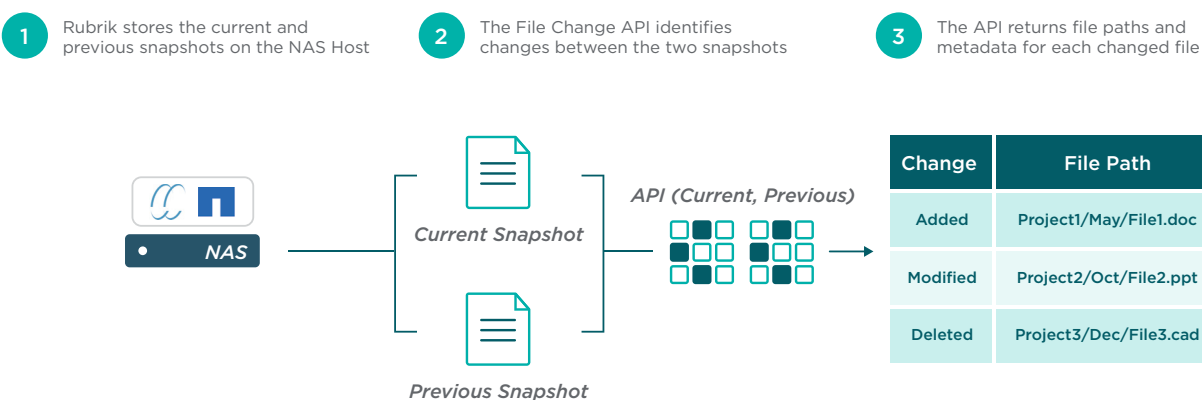


As is the case when there is no API integration with the NAS platform, Rubrik will scan a snapshot using optimized stat() calls.

Scan with Snapshot and File Change APIs Integration

In addition to snapshot API integration, some NAS vendors also provide file change API integrations. Using a file change API eliminates the need to do a traditional scan for changed files. Instead, the file change API will return to Rubrik a list of files that have been added, modified, or deleted since the last backup on their respective NAS platforms.

Rubrik has integrated with Isilon's ChangeList API and NetApp's SnapDiff API to provide faster and more efficient scanning with additional vendor API integration planned. Note that snapshot API integration is required with file change API integration.



The backup workflow with the snapshot API and the file change API is as follows:

1. Rubrik invokes the NAS vendor's snapshot API to create the first backup. The snapshot will be a full copy of the active file system.
2. Rubrik mounts and protects the full snapshot using our optimized scan method
3. On the next run, Rubrik invokes the snapshot API again to create a new snapshot.

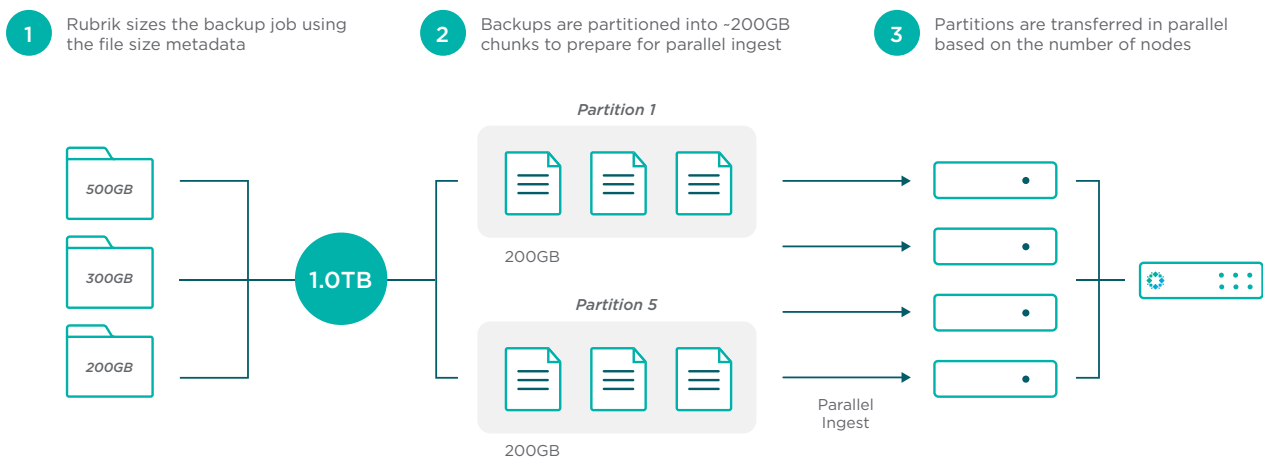
4. Rubrik invokes the NAS vendor's file change API (ChangeList for Isilon or SnapDiff for NetApp) to compare the current snapshot against the previous snapshot.
5. Rubrik receives, via the file change API, a list of changes since the previous snapshot was taken. This step obviates the need for Rubrik to conduct a scan of the file system.
6. Rubrik performs an incremental backup that only include the changed files.
7. Rubrik invokes the snapshot API to delete the older of the 2 remaining snapshots.

Fetch

Once the final list of files to be protected has been determined during the scan phase, Rubrik reads the files over either the NFS or SMB protocol. Fetching of files is done using a number of optimization techniques.

After the initial full backup, Rubrik takes an incremental forever approach to protecting NAS file systems. Only files that have not been modified or added since the previous backup will be fetched, significantly reducing the required backup window.

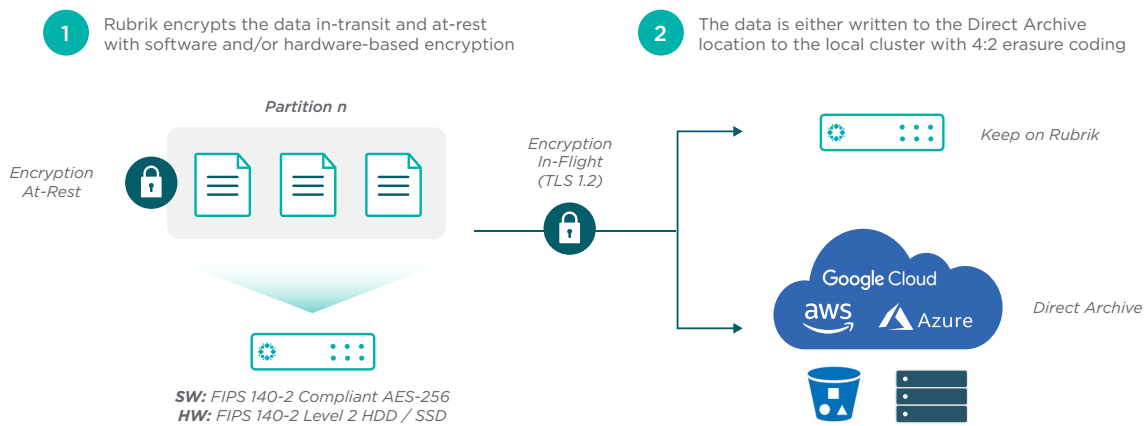
To take advantage of Rubrik's parallel architecture, backups are divided into ~100 to ~200 GB partitions which then can be ingested over parallel streams to different Rubrik cluster nodes. In the example below, three filesets with varying sizes are fetched by Rubrik and divided into five ~200 GB partitions. Each partition is then streamed in parallel to one or more Rubrik nodes.



During fetching, all files are indexed by Rubrik to enable global search and rapid recovery. Files are also broken into blocks and fingerprinted before being copied to Rubrik to enable more efficient storage of data.

Copy

The last phase in the backup process is the copy phase where files from the NAS platform are streamed to the Rubrik cluster or directly to an archive location.



During the Copy phase, all data are encrypted using AES-256 Asymmetric Encryption and streamed in parallel to the Rubrik cluster. All data streams are encrypted, providing encryption in transit as well, and all backups are written to disk on the Rubrik nodes in their encrypted-at-rest state.

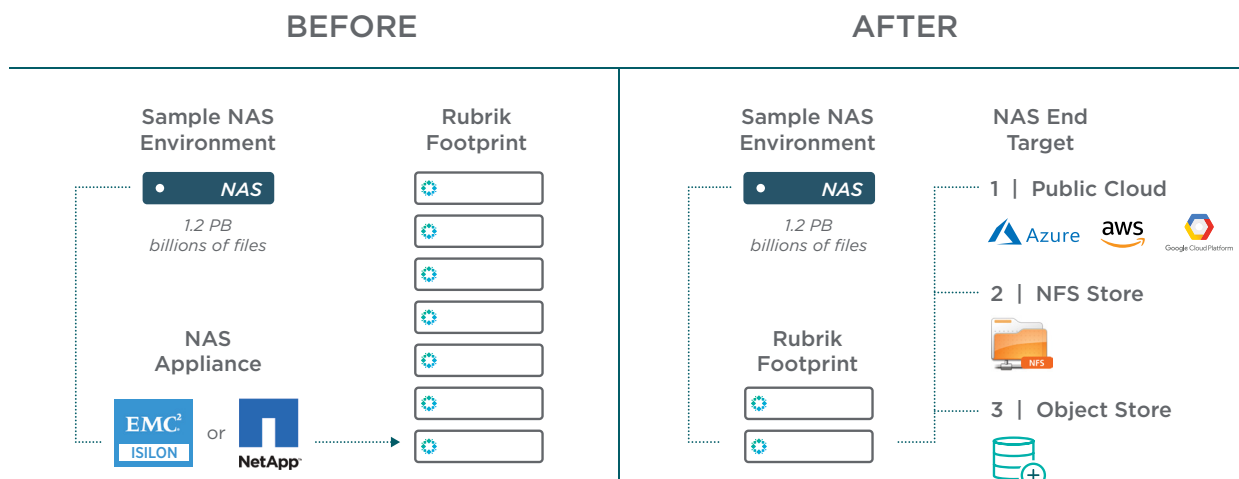
Unlike NDMP, files are not saved in a proprietary format but retains its source formatting and can be easily restored to any vendor's NAS platform.

RUBRIK DIRECT ARCHIVE

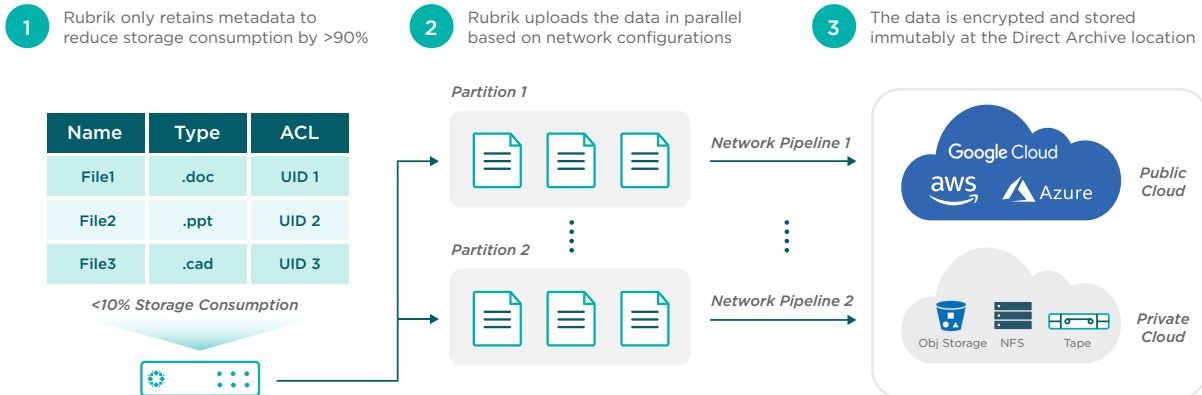
NAS backups can be stored locally on the Rubrik cluster or sent directly to an archive location during the copy phase. Many enterprise customers with large scale NAS environments prefer the latter approach to save on capital expenditure.

Rubrik Direct Archive provides the option for customers to save NAS backup data directly to an archive location without having to first store it in Rubrik but still retain the benefits of global search and rapid recovery. The archive location can be any of the following:

- A public cloud object storage service such as Amazon Simple Storage Service (S3), Microsoft Azure Blob Storage, or Google Cloud Storage
- A private cloud object storage solution such as NetApp StorageGRID
- An on-premises NFS store

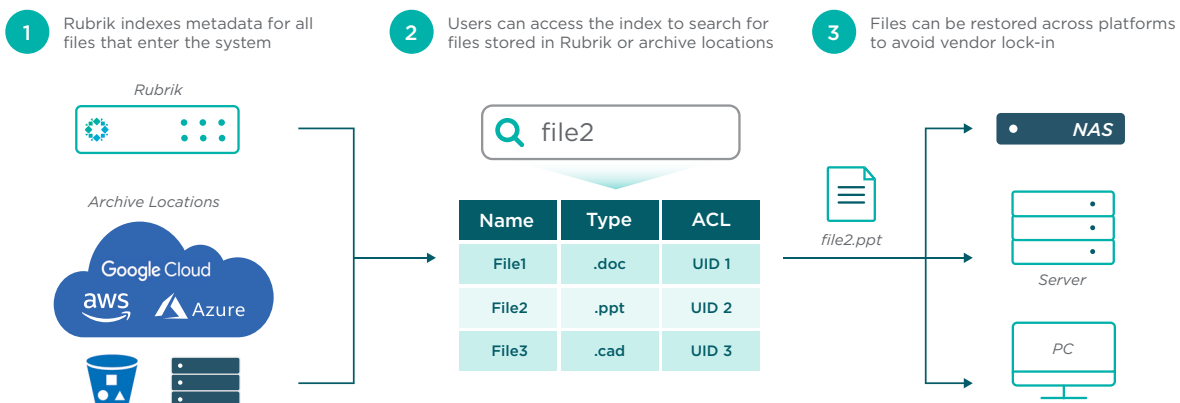


During fetching, Rubrik will index and store metadata about the files being protected as it would with backups that are being copied and stored locally on Rubrik clusters. The metadata is stored locally on Rubrik and is typically <10% of the capacity of the entire fileset. The data itself is encrypted and streamed to a designated archive location. Since the metadata is retained on the Rubrik cluster, customers can leverage global search and rapid recovery wherever the backup data resides.



GLOBAL SEARCH AND RAPID RECOVERY

Ultimately, a NAS protection solution is only as good as the ability it provides user for finding and recovering their data. By optimizing the use of file metadata, Rubrik simplifies and accelerates file recovery by NAS.



All files being protected are indexed during fetching and the metadata is stored locally on the Rubrik cluster, regardless of where the data itself is stored. Users are able to use the Rubrik console or CLI to search the index for any files using the filename, type, and other attributes. There is no need to rescan the entire file system to locate the files to be recovered. The location of a file can be on the local Rubrik cluster or in an archive location.

Once the files to be recovered has been located, they can be restored to the original NAS platform, to another NAS platform, or a server using the same file protocol as the original NAS platform. The file location is abstracted from the user and the file itself can be located on any node in a Customer's Rubrik clusters on in any archive location on or off premises.

BENEFITS OF USING RUBRIK

The Rubrik approach to NAS protection benefits users in a number of areas:

- **Reliability** - Rubrik ensures that customer NAS data is consistently protected by optimizing the data protection process and leveraging Rubrik's scale-out architecture for parallel streaming of backup data. At a time when the risk to data is higher than ever, due to exponential data growth and new forms of malware, Rubrik customers know they can rely on their backups to recover from incidents and disasters.
- **Rapid Recovery** - "Time is Money" is a truism that is as valid today as it was decades ago. Every moment that a business is down or unable to access critical data equates to lost revenue and lost opportunities. The issue is particularly acute in a large NAS environment with millions of files where all files may have to be recovered or a single file has to be located and recovered. Rubrik customers can take advantage of Rubrik's scale-out architecture for rapid recovery of data when recovering an entire file system. Using Rubrik's global search capability, customers can recover a subset of files rapidly, without having to manually search through a file system.
- **Flexibility** - By foregoing NDMP and storing NAS backup data in a standard format, Rubrik customers have the flexibility of being to protect any NAS platform and restoring the data to any other NAS platform. Customers are not locked in to a specific NAS vendor and their implementation of NDMP. This opens up the possibility of migrating between NAS platforms or leveraging a lower-cost NAS platform for secondary NAS storage.
- **Built for Growth and Scale** - Rubrik's scale-out architecture enables customers to protect their data in times of rapid growth. Every node added to a Rubrik cluster provides not only additional capacity, but additional performance. As new nodes are introduced, more data streams become available that can be used to partition NAS backup data for parallel ingestion of data. Customers can also leverage the unlimited capacity of the public cloud by using NAS Direct Archive to send data directly to public cloud storage.

CONCLUSION

Rubrik is the leading next-generation data management platform and is providing an increasing number of enterprises with solutions for protecting data in their data centers and in the public cloud. Customers can leverage the ongoing innovations of the Rubrik platform to ensure their data is protected as they enter new markets and expand their businesses.

ABOUT THE AUTHOR

Kenneth Hui is a Senior Solutions Architect in Technical Marketing at Rubrik. He has 20+ years of experience in IT, designing and administering technical solutions for commercial and enterprise customers. Ken has experience working in data centers and with cloud providers. His role at Rubrik is to create content to educate fellow technologists on these technologies and how to leverage them to solve customer challenges.