TECHNICAL WHITE PAPER

How It Works: Rubrik NAS Protection

Ed Morgan
September 2023
RWP-0504

Table of Contents

INTRODUCTION

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.

AUDIENCE

The intended audience for this document includes sales engineers, field consultants, professional services, partner engineering, and customer architects and engineers who wish to learn more about Rubrik's NAS functionality.

OBJECTIVES

The objective of this guide is to provide the reader with a clear understanding of Rubrik's approach to protecting NAS workloads at scale, and a technical reference to the architecture and design that enable this.

OVERVIEW

Enterprise NAS backup is challenging. Over the past two decades, NAS file systems have grown into petabyte scale in many industry verticals, and file types and use cases have gone beyond general-purpose office files and home directories. Increasingly, NAS stores either very large multimedia files such as 4K video or billions of small files such as bank check images or IoT sensor data, and typically each file type or workload often has a different performance or storage profile.

CHALLENGES

As data continues to grow, managing infrastructure for backup, recovery, and compliance can become incredibly complex. However, enterprises today require IT teams to not only manage all the infrastructure for mission-critical applications, but also to deliver new services that drive competitive advantage. Amidst the complexity, many organizations are looking to leverage public cloud, but face significant challenges in optimizing cost and performance at enterprise scale.

The sheer size of NAS and the variety of content profiles make data protection very difficult for backup administrators. As all enterprise NAS systems have proprietary file systems and are generally appliance-based, there is typically no ability to install a backup agent as would be done on a traditional file server running Windows, Linux, or Unix.

The industry-standard method for protecting NAS, NDMP (Network Data Management Protocol), is not efficient enough to protect large NAS environments. NDMP, as a protocol, was designed and developed 20 years ago[1] when datasets were both considerably smaller than today's workloads, and also when the long-term

1 https://tools.ietf.org/html/draft-skardal-ndmpv4-04

primary destination for backup data was tape. Additionally, NDMP is only a control protocol and does not dictate the format of the data, thus as a result, each data protection vendor sends the backup stream in a proprietary format.

NDMP also requires periodic full backups, otherwise the recovery chain from large sequences of incrementals can drive RTO to an unacceptable level. Also, from a performance perspective NDMP is only designed to support single-stream backups, which can create bottlenecks during the transfer process that a modern, scale-out architecture does not suffer from.

THE RUBRIK APPROACH FOR NAS

Rubrik's approach to building software, from day one, has been focussed on simplicity and elegance. The desire to simplify what has traditionally been extremely complex in data protection software—from consolidating an infrastructure-heavy, siloed architecture to boiling down an unwieldy job-based system defining RPO, replication, and archival into an SLA based policy engine[2].

Rubrik's approach to protecting NAS environments is no different. Key design criteria when building NAS protection into the product included the following:

- Must be vendor agnostic—data can be restored to a dissimilar system from the source
- Data must be stored in the original format—this significantly reduces RTO as we do not need to unpack the data before we can restore it.
- Incremental Forever, reducing backup windows, shortening RPO, and minimizing network consumption
- Full Rubrik functionality is available on the protected dataset
  - SLA based management
  - Instant search
  - Anomaly Detection
  - Sensitive Data Monitoring
  - NAS: Direct Archive
  - Et cetera

ARCHITECTURE AND DESIGN

PHASES OF A NAS BACKUP JOB

Whilst the full lifecycle of a Rubrik NAS backup is more complex than this diagram illustrates, and will be covered further within this document, at a high level it can be broken down into three high-level phases:

- *Scan*
- *Fetch*
- *Copy*

2 https://www.rubrik.com/blog/rubrik-sla-domain-settings-ops/

The Scan Phase

During the scan phase, Rubrik walks the file system of the share to discover files that need to be protected. For a full backup, everything is flagged for ingest, and in the case of incrementals the metadata attributes, including the *mtime* value, of each file is checked to discover if it has been modified or not. It is worth noting that the metadata scan will still capture POSIX permissions changes, even if the file content is not changed and we do not modify the *atime* attribute as part of the scan.

Prior to Rubrik 4.1, traditional POSIX file scans were leveraged to walk the filesystem. With more recent releases of Rubrik, this has changed to use optimized file system libraries to vastly increase performance of the Scan times. During internal testing, we have seen up to 10x increase in performance for NFS backups and up to 3x increase in performance for SMB.

In the example below, we have a 4 deep directory structure which contains *file1*, *file2*, and *file3*.

As the traditional POSIX libraries have to check each level of the directory tree every time a file is checked, scanning these 3 files would require fifteen *stat()* calls:

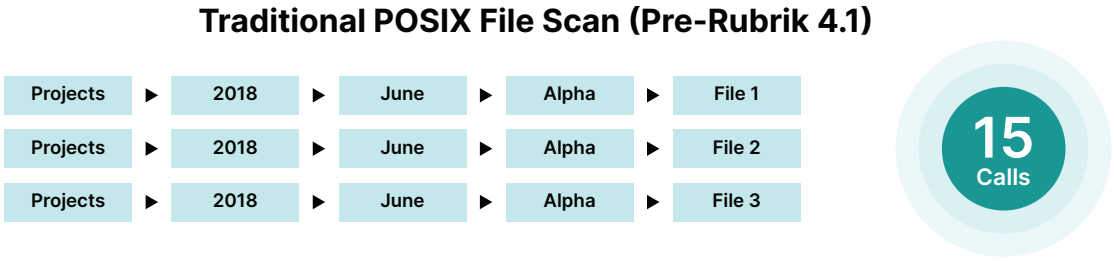

With the introduction of LibNFS in Rubrik 4.1 and LibSMB2 in Rubrik 4.2, each parent directory only needs to be checked once, reducing the previous fifteen *stat()* calls to seven.

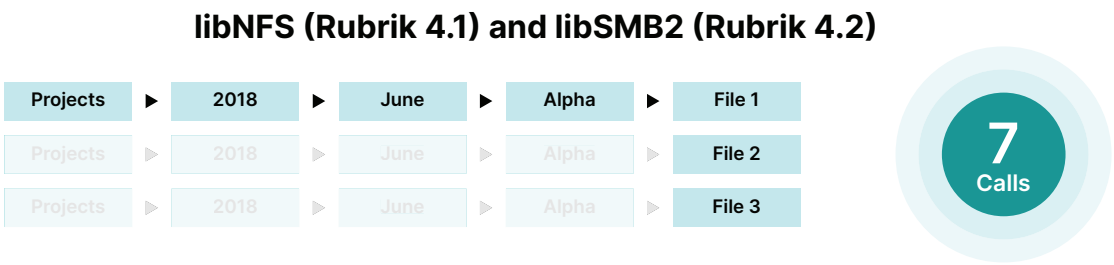

Although the above example uses only a small dataset, as the dataset grows the performance benefits increase with it, and in larger directory structures Rubrik can achieve significantly faster scan rates with the new libraries.

Note: In the case of NetApp FAS and EMC Isilon can leverage the SnapDiff[3] and ChangeList[4] APIs to offload the scan to the array's native file tracking APIs.

3 See the NetApp SnapDiff white paper *RWP-0505* for further information.

4 See the EMC ChangeList white paper *RWP-0521* for further information.

The Fetch Phase

The fetch phase takes a list of files gathered during the scan phase, either via traditional scan or vendor differential APIs, and reads them over NFS or SMB in parallel across multiple nodes within the Rubrik cluster. In the case of scale-out NAS systems, Rubrik also ingests from multiple NAS nodes.

After the initial full backup, Rubrik uses an incremental forever approach and only ingests data that has been modified since the last backup. Data is chunked into 400GB "partitions", which are then distributed across the Rubrik nodes and ingested in parallel into the system.

All files are indexed to allow for granular search and recovery, and then data is written locally on the Rubrik cluster.

The Copy Phase

During the copy phase, data is reduced by only keeping unique blocks, encrypted using AES-256 Asymmetric Encryption, and then is written out to the local Rubrik file system and distributed across the cluster using 4:2 erasure coding.

All data that is written to the Rubrik cluster is, by the nature of the Rubrik *Atlas* file system, stored in an immutable fashion, and is also stored in its source format to allow restores in heterogeneous environments across NAS platforms and vendors.

Depending on the retention settings defined in the SLA, this data is typically tiered off after a short period of time to either NAS, object storage or the public cloud for long term retention.

> **Note:** If the data being protected is leveraging NAS: Direct Archive[5], the data will only be indexed on the local Rubrik cluster before being passed to the Archive repository. No data outside of the metadata index will reside on the local system in this index.

KEY FEATURES

As detailed above, it was imperative when conceptualizing the Rubrik approach to protecting file data that it was treated as a "first class citizen" in the Rubrik ecosystem, and thus that all core functionality available to other datasets is available to NAS workloads being protected. These include those detailed below.

SLA Based Management

Rubrik has, since its inception, focused on simplicity of management. One of the core constructs that enables this is the distillation of traditional, complex jobs and schedulers into the concept of SLA Domains.

Whilst an in-depth look at SLA Domains is outside the scope of this document, at their core the SLA engine allows users to specify their RPO requirements and retention policies, and the system takes care of the rest of the minutiae normally associated with managing data protection.

5 See the NAS: Direct Archive white paper *RWP-0516* for further information.

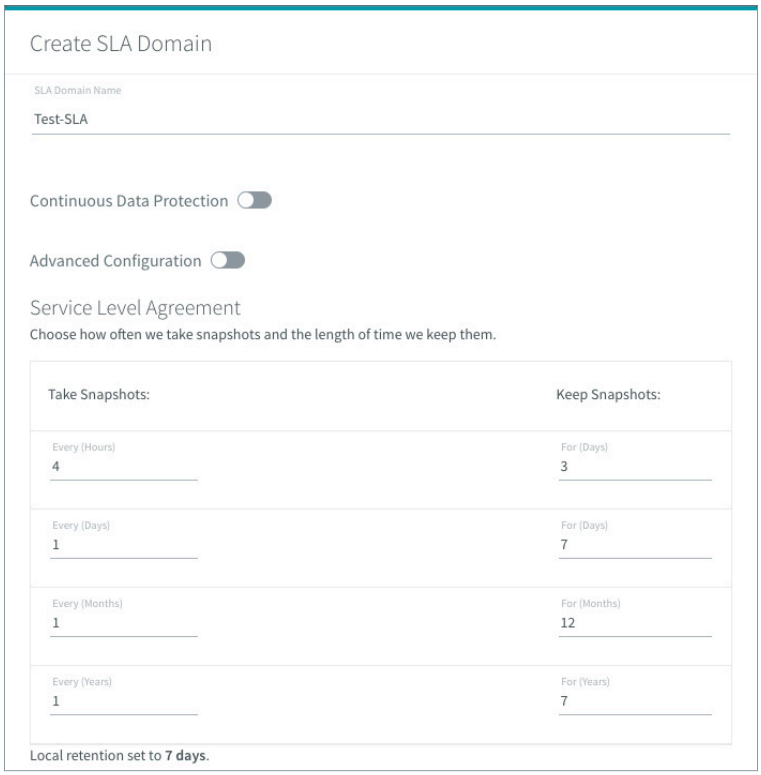

"Google-like" Search

Rubrik indexes all the data managed by the system, allowing for a global search for files across all backups, whether they are stored on-premises or in the cloud, with a predictive search algorithm that delivers suggested results as the user types their query. This allows Rubrik to perform single file restores from archive targets without requiring to pull back full snapshots. The data is dynamically reconstructed prior to retrieval, and only the required dataset is restored, keeping RTO low, and egress charges and network usage minimized.

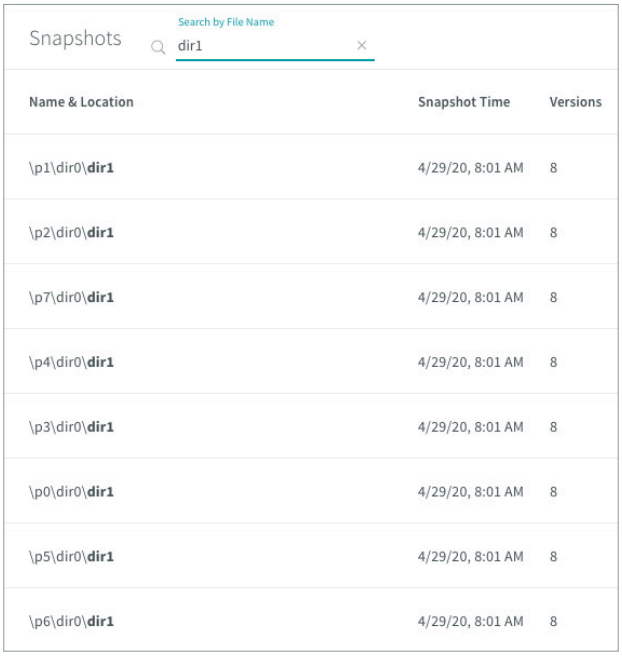

Data Immutability

Rubrik has been designed from day one to store data in an immutable format. An immutable file system, paired with a zero trust cluster design in which operations can only be performed through authenticated APIs, ensures that once data has been committed to the Atlas file system it cannot be deleted, even if exposed via functionality such as Live Mount.

For more information on Rubrik's immutable architecture, please refer to the blog post "The Magic of an Immutable Backup Architecture" on the Rubrik website.

Anomaly Detection

One of the products within Rubrik's SaaS offering is Anomaly Detection. Anomaly Detection applies machine learning against application metadata to establish normal baseline behaviour for the workload, and then proactively monitoring datasets and flagging any activity that varies significantly from the established baseline.

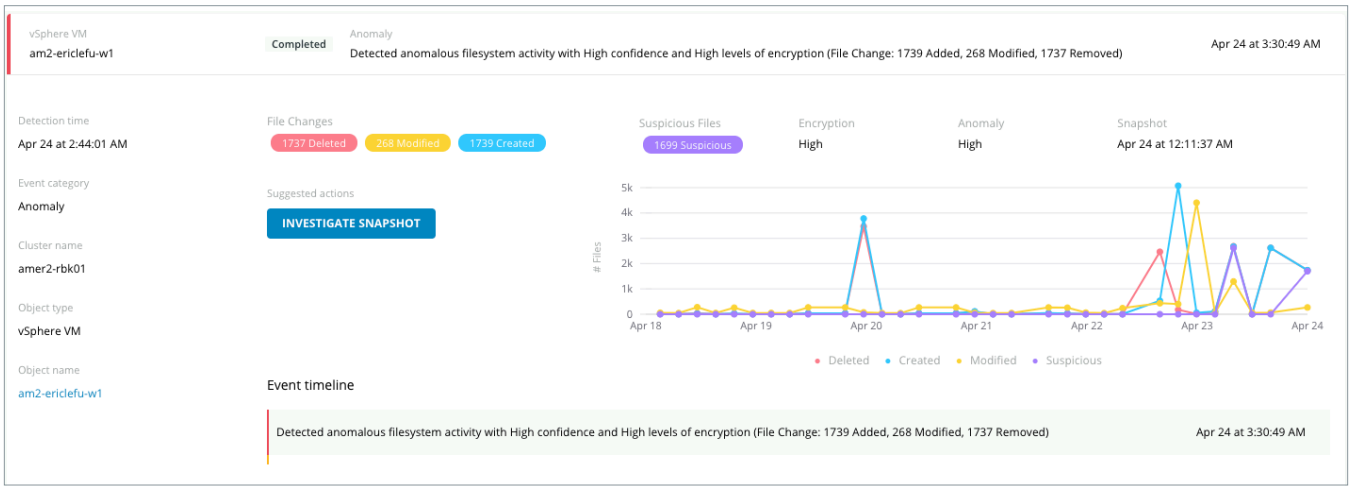

For more information on Rubrik Anomaly Detection, please refer to the Rubrik Anomaly Detection datasheet on the Rubrik website.

Sensitive Data Monitoring

Another of the products within Rubrik is Sensitive Data Monitoring. Sensitive Data Monitoring scans and classifies sensitive data, and leverages either pre-built or custom generated dictionaries and policies to identify common data types from regulations and standards such as GDPR, PCI-DSS, HIPAA, and GBLA. Datasets are then proactively searched and any infractions are alerted to the customer, and can be reported on in the event of a compliance audit.

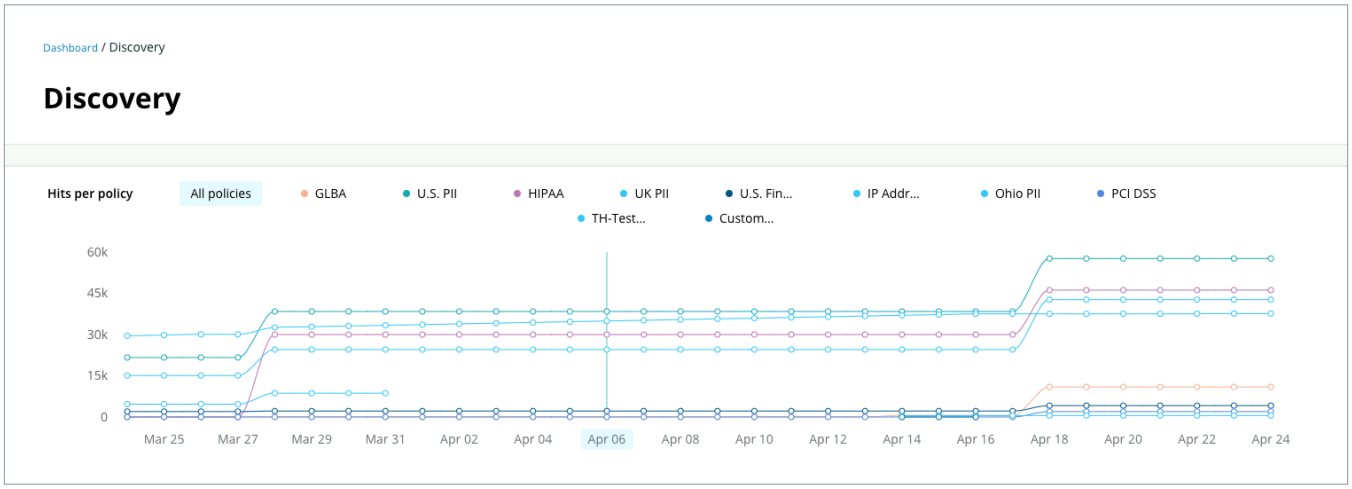

For more information on Rubrik Sensitive Data Monitoring, please refer to the Rubrik Sensitive Data Monitoring datasheet on the Rubrik website.

NAS Direct Archive

NAS Direct Archive was introduced to suit the requirements of customers that have large quantities of unstructured data. Whilst it makes sense to retain local copies of backup data for tier 1 workloads such as VMs and databases, often it is both cost-prohibitive and counterintuitive to retain local copies of large amounts of cold file data on the local Rubrik cluster. NAS Direct Archive allows for unstructured data to be immediately passed through to the storage platform being used for long term retention whilst still being indexed and managed by Rubrik.

Further information and a technical deep dive on NAS Direct Archive can be found in the Rubrik White Paper *"RWP-0516 - How It Works: NAS Direct Archive"* document on the Rubrik website.

HOW IT WORKS

BACKUP WORKFLOW

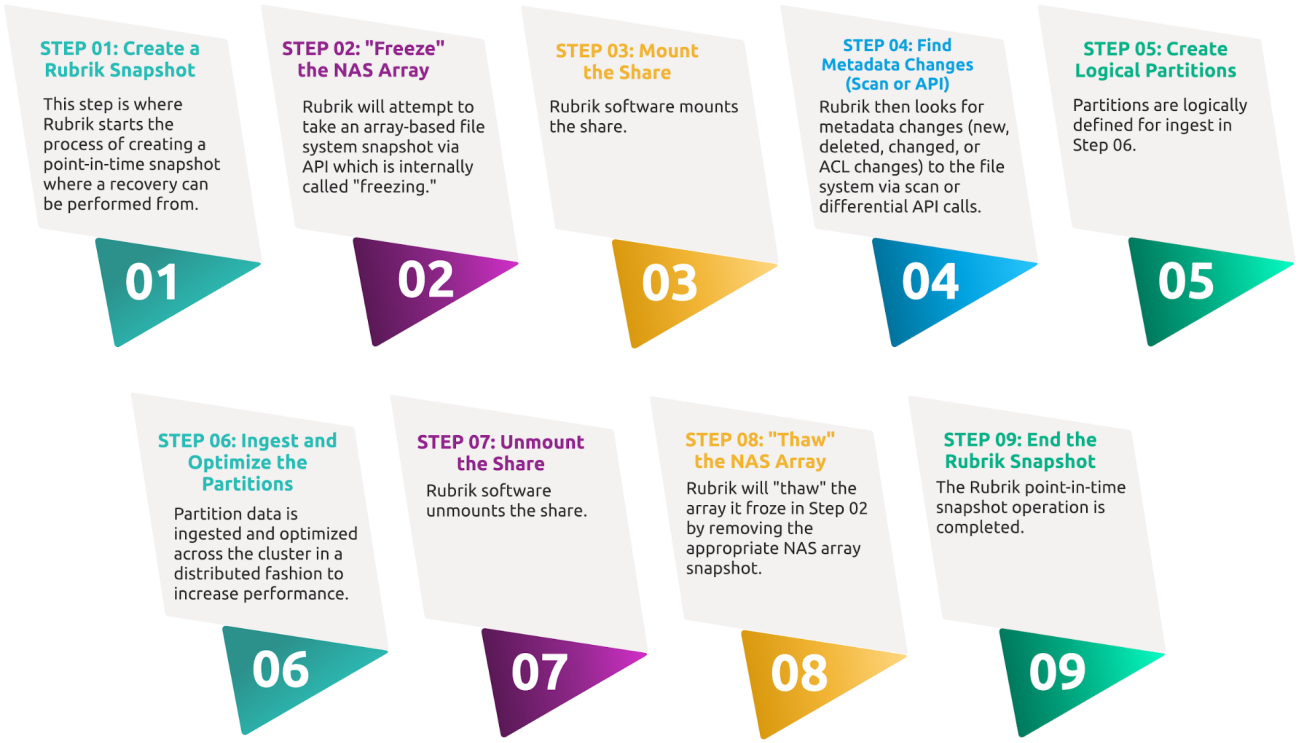

Phase 1: The Scan Phase

The scan phase ascertains what Rubrik needs to ingest, either via the vendor's snapshot differential APIs, or via a traditional scan.

STEP I: CREATE A RUBRIK SNAPSHOT

This starts the process of creating a recovery point.

STEP II: "FREEZE" THE NAS ARRAY (OPTIONAL)

As mentioned above, Rubrik supports certain vendor's snapshot APIs (DellEMC Isilon, Pure Flashblade, and NetApp FAS). If backing up one of these supported arrays, Rubrik will first attempt to create and mount an array based snapshot of the volume being protected to both ensure that opened files can be protected, and to later leverage the vendor's snapshot differential APIs.

If, for whatever reason, this snapshot process fails, Rubrik will instead mount the live volume as opposed to the snapshot created.

STEP III: MOUNT THE SNAPSHOT OR SHARE/VOLUME

If the snapshot was created in Step II, Rubrik nodes will mount these snapshots, or failing that will mount the active filesystem from the NAS. As each Rubrik node in the cluster can participate in this process, the more shares there are, the greater levels of parallelism (and thus performance) can be achieved in the subsequent steps of the backup process.

In a large environment where the organisation has tens of thousands of shares and hundreds of NAS devices, the load to protect them would be distributed across Rubrik nodes without any manual load balancing by customer engineers. Load is evenly distributed across the Rubrik cluster without the requirement for engineers to define a direct relationship between share and Rubrik node.

STEP IV: FIND METADATA CHANGES

A core design principle of Rubrik is that all source objects are protected using an incremental-forever approach. Once a snapshot or share is mounted onto a Rubrik node, the files and directories in that share are scanned. The initial scan will result in a full dataset being backed up, but in subsequent runs either vendor specific differential APIs will be used to compare the array based snapshot created in the previous backup task, or a traditional file scan as detailed in the Architecture and Design section within this document.

Once a list of new or modified files has been established, these will be ingested onto the Rubrik cluster in the next core phase of the protection process, known as the "Fetch" phase.

Phase 2: Fetch

The "Fetch" phase pulls the required data from the source NAS onto the Rubrik cluster.

STEP V: CREATE LOGICAL PARTITIONS

Once the metadata tree is created by the previous task, Rubrik will then split this dataset up into logical boundaries called "partitions". These logical boundaries are defined by size, and allow the Rubrik cluster to fetch multiple partitions concurrently to different nodes within the cluster to increase parallelism and throughput performance.

STEP VI: INGEST THE PARTITIONS

The files identified in the Scan phase are now ingested onto the Rubrik cluster. Maximum performance is achieved through the distribution of ingest onto multiple Rubrik nodes fetching different partitions of data simultaneously.

Phase 3: Copy

The "Copy" phase is the step of the protection process that distributes and commits the source data to the Rubrik file system.

Note: If the data being protected is leveraging NAS: Direct Archive[6], the data will only be indexed on the local Rubrik cluster before being passed to the Archive repository. No data outside of the metadata index will reside on the local system in this index.

STEP VII: UNMOUNT THE SNAPSHOT OR SHARE/VOLUME

Rubrik unmounts the source dataset.

STEP VIII: "THAW" THE NAS ARRAY (IF STEP II WAS PERFORMED)

Rubrik will remove the appropriate array based snapshot from the array.

STEP IX: END THE RUBRIK SNAPSHOT

Data is written into journal files locally to the Rubrik node doing the ingest, which are then verified for integrity and converted to patch files. Patch files are Rubrik's proprietary file format which are then distributed out across the Rubrik cluster using Reed-Solomon 4:2 erasure-coding for resilience purposes, and then committed to Atlas as part of the backup dataset under management by Rubrik. This completes the point-in-time snapshot operation.

6 See the NAS: Direct Archive white paper *RWP-0516* for further information.

RESTORE WORKFLOW

There are four different ways to restore data from Rubrik:

- Download
- Export
- Overwrite Original
- Restore to Separate Folder

Download

When *Download* is selected, Rubrik will reconstruct the data into the required format and then package it up into a zip file that the administrator can download to their workstation.

If this option is selected, it must be ensured that there is enough free space available on the Rubrik cluster to create the zip file. This can be done by viewing the summary dashboard in the Rubrik UI and checking the Capacity Available:

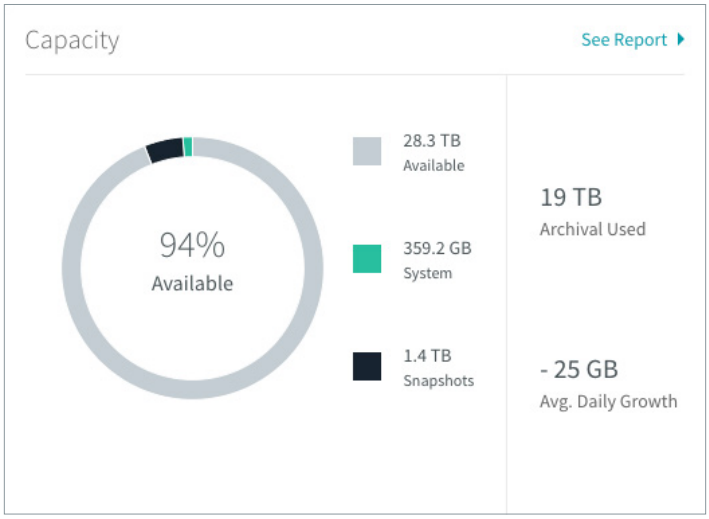

If there is not sufficient space available on the cluster, the Download operation will fail.

Export

The *Export* option allows for the selected files and folders to be restored to a different target to the original source.

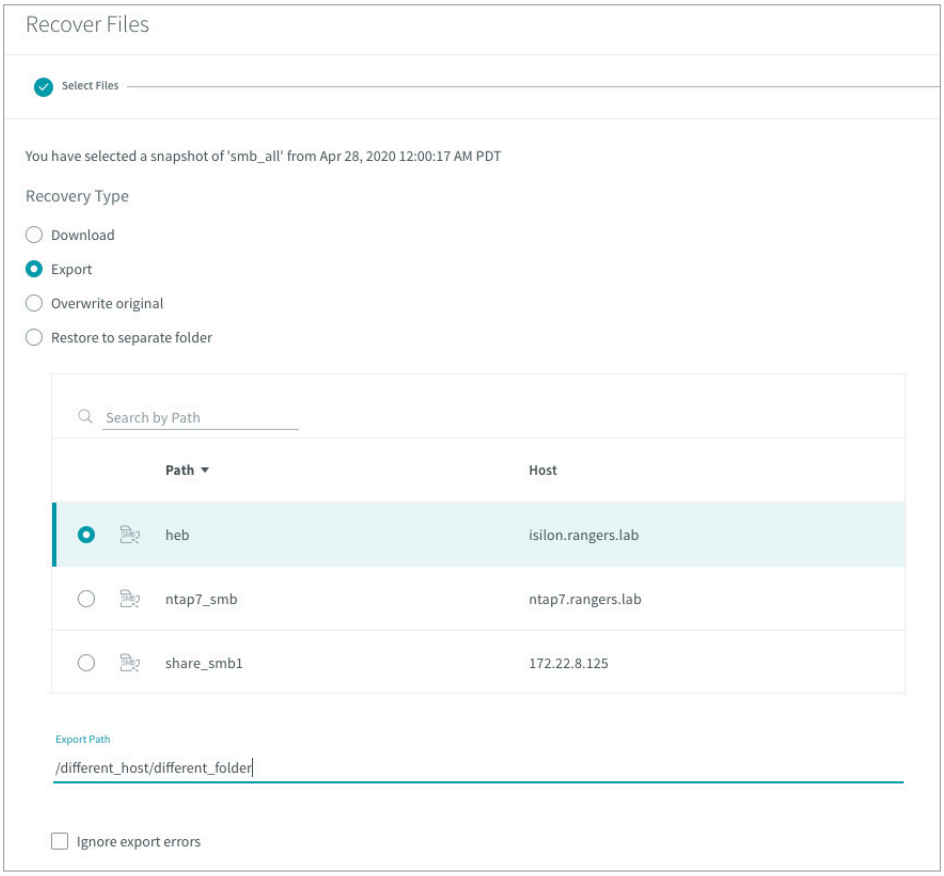

You can restore to any NAS device or file server regardless of vendor. The only requirement is that the data is restored via the same protocol with which it was backed up with.

Overwrite Original

The *Overwrite Original* option will restore the selected dataset back to its original location, overwriting the source files if they still exist.

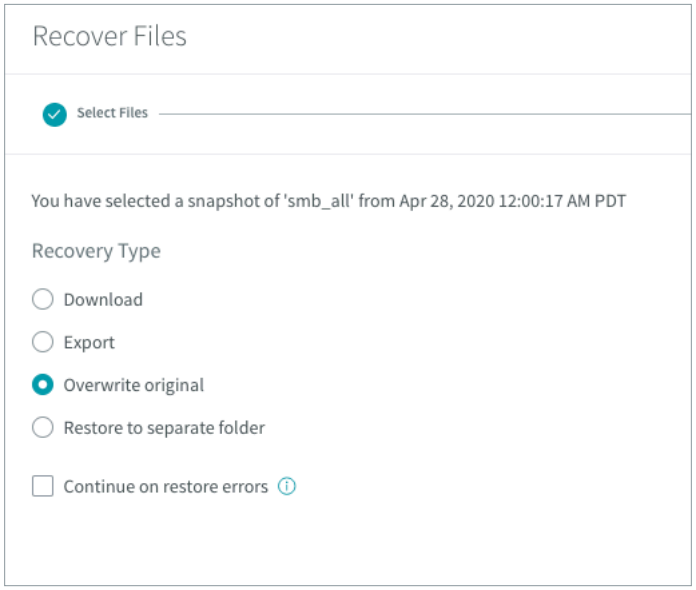

Restore to Separate Folder

Restore to a Separate Folder allows you to restore the dataset back to the original share, but to a new directory.

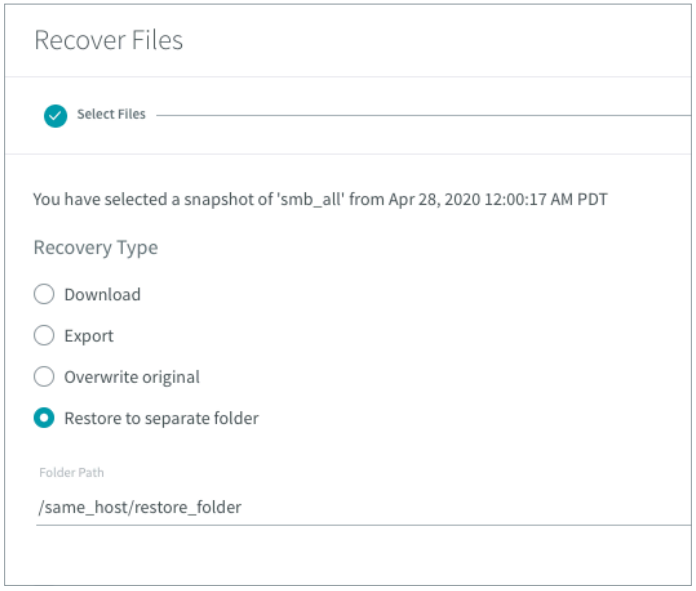

If the specified location does not exist, it will be created as part of the restore operation.

CONCLUSION

This document gives a thorough view into Rubrik's technical and architectural approach to protecting unstructured data in the enterprise, and the benefits that come with a modern, data-centric approach to data protection.

Whilst every effort has been made to have this document be as informative as possible further information on the products detailed can be found via the Rubrik website in the form of additional downloadable Rubrik White Papers, or via the support portal at https://support.rubrik.com.

GLOSSARY

MTIME

The file attribute *mtime* (modification time) indicates the time the contents of the file has been changed.

ATIME

The file attribute *atime* (access time) indicates the time the contents of the file has been accessed.

STAT()

stat() is a Unix system call that returns file attributes about an inode. The semantics of *stat()* vary between operating systems. As an example, the Unix command *ls* uses this system call to retrieve information on files that includes *atime*, *mtime*, and *ctime*.

ATLAS

Atlas is an immutable Filesystem in Userspace (FUSE) that is largely POSIX compliant. Atlas is custom designed to be a distributed and immutable file system for writing and reading data for other Rubrik services. Immutability is provided in multiple ways to ensure data integrity. For more information on Atlas, please refer to "The Magic of an Immutable Backup Architecture" and "Introducing Atlas, Rubrik's Cloud-Scale File System".

ERASURE CODING

In coding theory, an erasure code is a forward error correction (FEC) code under the assumption of bit erasures (rather than bit errors), which transforms a message of k symbols into a longer message (code word) with n symbols such that the original message can be recovered from a subset of the n symbols. The fraction r = k/n is called the code rate.

VERSION HISTORY

Version	Date	Summary of Changes
1.0	May 2020	Initial Release
1.1	January 2022	Update product naming in line with Winter 2021 Release
1.2	September 2023	Product naming and boilerplate updates

rubrik

Global HQ
3495 Deer Creek Road
Palo Alto, CA 94304
United States

1-844-4RUBRIK
inquiries@rubrik.com
www.rubrik.com

rwp-hiw-rubrik-nas-protection / 20230913