rubrik

PURE FLASHBLADE DATA PROTECTION WITH RUBRIK

EXECUTIVE SUMMARY

As data continues to grow, supporting billions of files or more creates an increasing strain on backup and recovery environments. Traditional methods of backing up large amounts of file data can exacerbate server CPU and memory usage to the point of degraded application performance.

Rubrik's native NAS protection and universal fileset technology provide a simple and effective method to protect and recover NAS file data regardless of customer environment size. NAS protection leverages a common fileset concept as Linux, AIX, and Windows hosts and intentionally does not rely on Network Data Management Protocol (NDMP). Not using NDMP is a conscious design choice which enables the Rubrik approach to be vendor agnostic, provide true incremental forever even for NAS, and store data in a "native" format.

FlashBlade™ from Pure Storage is an extremely fast and dense, unstructured data platform that provides rapid backup and recovery of archival data, along with other primary unstructured data use cases. FlashBlade is tuned to handle small, metadata-heavy to large files, random or sequential access patterns, up to 10s of thousands of clients all requiring real-time response, without the need to constantly retune the storage platform.

From the ground-up, hardware to software, it is built with a massive parallel architecture to deliver maximum performance for any modern workload today and in the future. FlashBlade delivers superior performance in a small form factor. It is tuned to deliver multi-dimensional performance for any data size, structure, or access, delivering greater savings in power, space, and cooling costs.

The goal of this white paper is to characterize the performance of protecting primary file data residing on Pure FlashBlade and to help customers understand the recovery point and recovery time objectives (RPO, RTO) associated with a combined Pure and Rubrik solution.

COMMON CHALLENGES OF PROTECTING NAS DATA

Customers typically face two common challenges when protecting NAS data:

- "Unsatisfactory backup performance or elongated recovery times are challenges that backup administrators typically encounter when protecting file data."
- "Large file servers with billions of files are often particularly hard to handle as they may take hours to days for a full backup."

HOW RUBRIK SOLVES THESE CHALLENGES

Rubrik takes a modern incremental-forever approach obviating the need for periodic full backups. Performing incremental-forever backups reduces nightly backup windows especially when there is a need to protect millions or billions of files. From a restore perspective, each backup looks like a full backup whether restoring a single file, a folder, or a full volume. This approach works not only for small enterprise NAS systems or file servers but for large NAS platforms as well. Operational management is minimized, thus providing valuable production time for administrators to focus on more closely aligned business objectives.

THE FIVE PHASES OF RUBRIK'S NAS DATA PROTECTION

The first concept to remember is that Rubrik's architecture is distributed in a peer-to-peer cluster design. Nodes can participate in any role, including backup, recovery, managing software, metadata operations, and providing APIs.

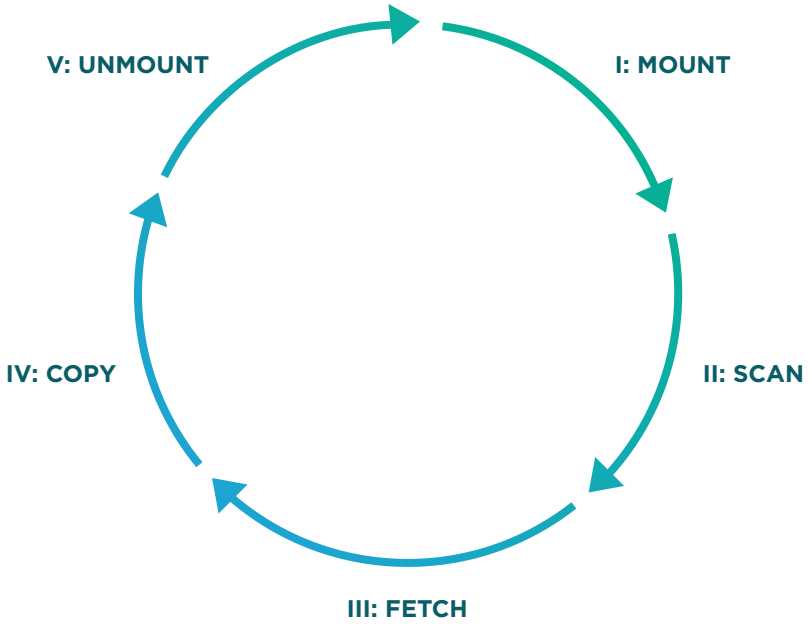

Phase I: Mount

Rubrik nodes mount NFS or SMB shares directly from the NAS devices or file shares. If an organization has tens of thousands of shares and hundreds of NAS devices, the load to protect them would be distributed across Rubrik nodes without customers having to define a direct relationship between share and Rubrik node - Rubrik knows of the share and an available node in the cluster will mount the share at the appropriate time.

Phase II: Scan

One of Rubrik's design principles is that all source objects are protected using an incremental-forever approach. How is that achieved with NAS share data given that NDMP is not leveraged using Rubrik?

Once a share is mounted onto a Rubrik node, the files and directories in that share are scanned. The initial scan will result in the full data set being backed up, but subsequent scans will identify only the incremental files and directories that need to be protected and they will be pulled across onto the Rubrik cluster in a phase that is known as the "Fetch" phase.

Phase III: Fetch

The identified files and directories in the "Scan" phase are then pulled onto the Rubrik cluster. To reiterate, only incremental files and directories are fetched. The goal of this phase is to pull data onto the cluster as quickly as possible from the source NAS device. Maximum ingest performance is achieved through the distributed Rubrik nodes fetching different chunks of data simultaneously. While each single stream fetching is quite impressive, it is the combined ingest throughput of all of the participating Rubrik nodes that creates the overall throughput performance.

Phase IV: Copy

The goal of the "Copy" phase is to erasure-code and spread data across the cluster in a distributed fashion to maintain data integrity and resiliency while minimizing data capacity.

Phase V: Unmount

Once the data is protected across the Rubrik cluster, the share is unmounted.

WHY TEST?

It is always important to help partners and customers understand what kind of performance they can expect from a combined solution between any two or more vendors. Pure and Rubrik agreed to do some joint lab testing and from the onset of the project felt it would be important to publish the results so everyone can understand them.

One important note: Performance is an area that continuously improves due to new innovation. The numbers and respective code bases shown below (and what Pure and Rubrik are measured against) are as of July 2018.

WHAT SHOULD BE TESTED?

The basis for this testing was centered on "real-world" data sets. All test data sets were designed around environments we see from our customer interactions.

The first step was to create a simple file generator tool with parameters regarding:

- Compressible versus non-compressible data
- Directory-depth settings: How many subdirectories should a structure have?
- Average file size: The tool should be able to create both small and large files

Test #1: *How quickly can Rubrik scan small files on a filesystem that has a complex directory structure?*

Customers often times have subdirectory structures of depth five or more with hundreds of millions to billions of files. Traditional systems can be very slow in these environments. Software development and AI/ML use cases have these types of directory structures and file scan times are the main performance-limiting step in protecting these environments.

Test #2: *How quickly can Rubrik ingest large files?*

One of Pure FlashBlade's use cases is "Rapid Restore" for large data sets like database dumps. After data is written to Pure FlashBlade, it can be managed via Rubrik's archive, search, and index capabilities. The question is how fast can this data be ingested into Rubrik.

Test #3: *How many files can Rubrik manage in a single fileset?*

This is more of a theoretical test. The likelihood of a single fileset being used to manage hundreds of millions of files is unlikely, but customers want to know what the highest level limit is.

TEST ENVIRONMENT

Network Details	Pure FlashBlade Details	Rubrik Cloud Data Management Cluster Details
Two Arista 7050 10GBps switches in MLAG configuration	7x8 Configuration (7 blades, 8 TB each in capacity)	8-node cluster
———	Purity v2.1.8	Running CDM v4.1.0
———	Connected to two QSFP+ 40GBps ports	All Rubrik nodes connected via 10GBps Twinax cables

TESTING RESULTS

Test #1: File Scan Performance

In this test, the goal was to maximize the scan rates in increasing levels of complex directory structures. The number of subdirectories defined in these sets of testing is the number followed by the letter "D". For example if a directory structure has (6) levels of subdirectories, that is defined by in this document as "6D".

Here are the test parameters:

Directory depth configuration:	Fileset configuration:
• 1D, 2D, 4D, and 8D	• Each "fileset" was created with files averaging 10KB in size • 3.125M files per fileset • Tested from 1–64 filesets concurrently • Ramp from 3.125M–200M files

Test #1 Results

The expectation was that increasing the number of directory depths would lead to lower maximum scan rates. For example, a directory complexity of 1D would be easier to scan files than at 8D. This proved to be true. The Rubrik 8-node cluster was over 80,000 files/sec for 1D, and ~55,000 files/sec for 8D and values in between for 2D and 4D.

It was not clear what the maximum file scans per second would be given that "traditional" NAS systems can be somewhat limited in metadata performance. The 80,000+ files/sec per 2-briks far exceed what was originally anticipated.

As of July 2018:

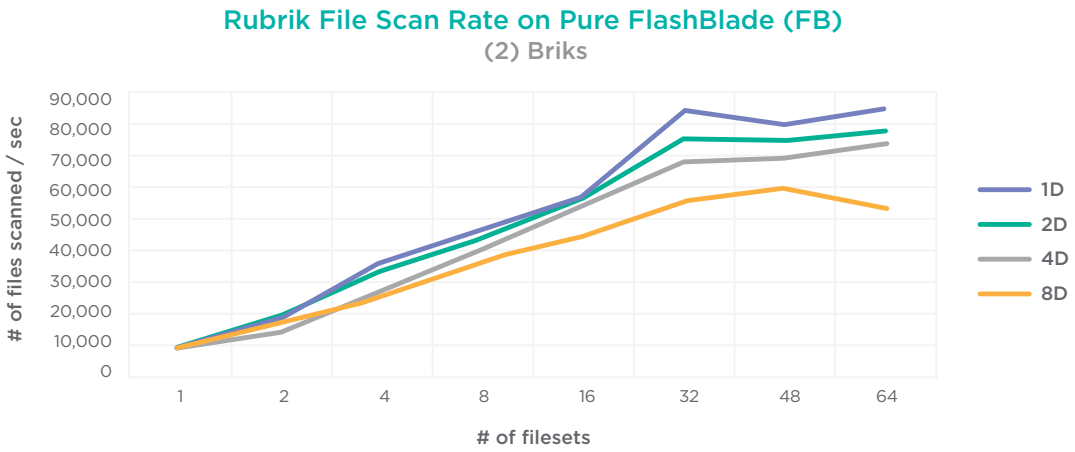

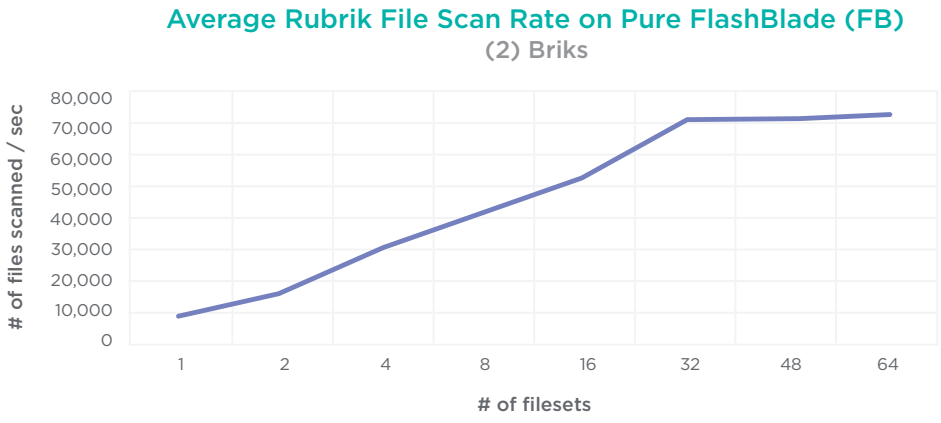

Test #1 Update

The maximum file scan rate in the initial testing (~10,000 files per second per Rubrik node) was impressive for both Rubrik and Pure engineering teams. Rubrik engineering was not satisfied with this number and sought to find more performance enhancements to further maximize this. The next Rubrik release (v4.2) has enhancements that have improved scan rates to 40,000 files per second.

Here is a table that shows how many files Rubrik can scan in 12 and 24-hours per day per Brik.

(File scans) per second per node	# Files scanned (12-hour/day) per Brik	# Files scanned (24-hr/day) per Brik
10,000 (old)	1.728 billion	3.456 billion
40,000 (new)	6.912 billion	13.824 billion

Both Pure FlashBlade and Rubrik scale linearly. With 10 Briks, users can scan 69,000,000,000 to 138,000,000,000 files per day.

Test #2: Data Ingest Performance

In this test, the goal was to see how much data could be ingested from Pure FlashBlade to Rubrik. The goal was to configure both systems naturally. We used default settings for both Pure FlashBlade and Rubrik, and the network ports were configured without any added help.

The test parameters are as follows:

- Directory depth configuration: 1D
- Eight directories created
- Each directory with 300 files and 675 GB capacity
- 5.4 TB for total data set

Test #2 Results

The ingest testing results with default settings were above expectations. Without tuning, the ingest performance per Brik was 1.67 GBps, with peaks around 2 GBps per Brik. For VMware, Rubrik states ingest rates of 1.2 GBps but have tested up to 2 GBps with tuning and setting values for non-default values. Pure FlashBlade ingest rates were 40% higher than for VMware out-of-the-box.

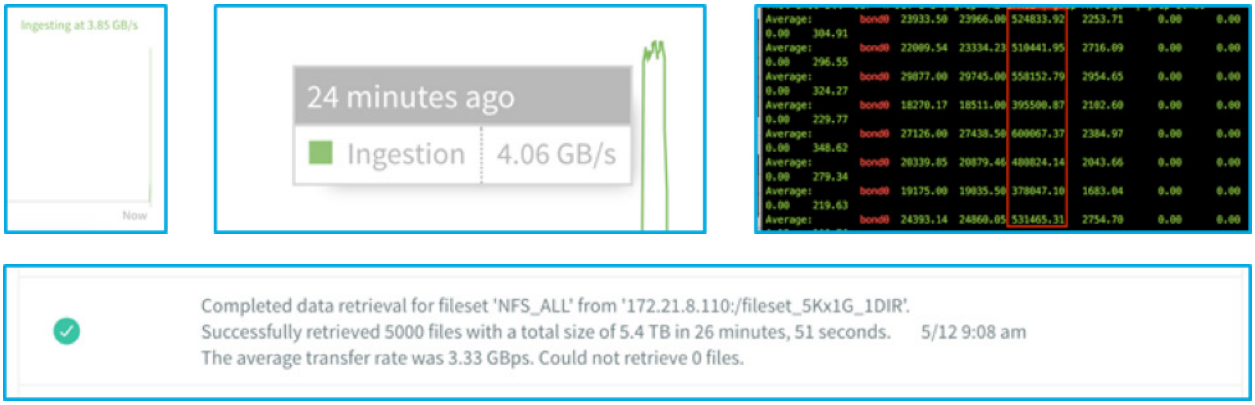

Test #2 Update

No further tests were performed.

Test #3: File Scan Performance

This test was mainly designed to see how many files Rubrik could scan in a single filesystem running on Pure Flashblade until performance was compromised. Of course, this number is just a point-in-time maximum as it has been steadily increasing since Rubrik's NAS support in v3 was released.

The test parameters are as follows:

- Directory depth configuration: 10D
- Average file size: 100KB
- Number of files (Ramp): 5M, 10M, 25M, 50M, 100M, 150M, 200M, 300M, 400M, 500M, 750M, and 1B
- 5.4 TB for total data set

Test #3 Results

In the early days of Rubrik's NAS support, there were best practice limits to the number of files in a single fileset. The recommendation was 10M files in v3 and 20-30M files in v4. Given the improvements to both the scan and fetch phases in v4.1, the expectation of the number of files for this round of testing was that maybe we could do 100M files in a single fileset.

We slowly ramped the number of files in the single fileset configuration from 5M to 10M to 25M to 50M to 100M. All jobs finished without issues. We then accelerated the fileset growth to 150M, 200M, 300M, 400M, 500M, 750M, and finally 1B files.

All jobs up to 750M files finished without issue, but there were slight performance issues at 1B.

Test #3 Update

Scanning 1B files in a single fileset on Pure FlashBlade should not be an issue. A single node can handle scanning 1B files in approximately seven hours, assuming 40,000 files per second. However, it remains to be seen if a customer would actually want to perform such an operation. Rubrik engineering is looking to improve the number of files that a single fileset can support.

Performance Update as of October 16, 2018

Testing on the latest Rubrik code release (4.2-GA) we can clearly demonstrate that a single fileset under Rubrik protection can scan files in at over 40,000 files per second. The following fileset with just under 4,000,000 files was scanned in 82.30 seconds!

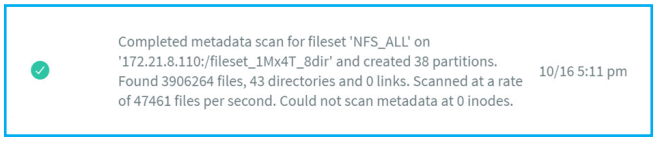

CONCLUSION

In just a few short pages, we have clearly demonstrated that the paradigm of backup and recovery has dramatically shifted from what has traditionally been a pain point in both large and small enterprises, to an easy to setup and manage, holistic solution. FlashBlade provides the density, scalability, and unparalleled performance to handle your most demanding unstructured workloads, while Rubrik allows for exceedingly simple and smooth recovery of entire directories, all the way down to individual files. Rubrik and Pure Storage have partnered to provide an integrated, performant, and simple-to-manage solution that will enable IT professionals to seamlessly protect and restore content for the Pure Storage FlashBlade™ array.

These complimentary products will enable customers to truly have the best of both worlds: FlashBlade for performant, dense, scale-out storage for unstructured workloads combined with Rubrik for robust data protection and granular point-in-time recoveries of individual files and entire directories. The combination of Rubrik and Flashblade delivers the entire suite of data performance, protection, and agility which end users have been craving for.

rubrik

Global HQ
1001 Page Mill Rd., Building 2
Palo Alto, CA 94304
United States

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

Rubrik delivers a single platform to manage and protect data in the cloud, at the edge, and on-premises. Enterprises choose Rubrik's Cloud Data Management software to simplify backup and recovery, accelerate cloud adoption, and enable automation at scale. As organizations of all sizes adopt cloud-first policies, they rely on Rubrik's Polaris SaaS platform to unify data for security, governance, and compliance. For more information, visit www.rubrik.com and follow @rubrikInc on Twitter.

20181121_v1