



TCO REPORT

NAS File Tiering

Economic advantages of
enterprise file management

Executive Summary

Every organization is under pressure to meet the exponential growth in demand for file storage capacity. Surveys show, however, that 60% or more of data on Tier 1 storage is either dormant or rarely used. Organizations can now achieve significant savings by moving that inactive content to a secondary storage tier.

While the concept of storage tiering is well known, it has not been widely adopted in the past due to various limitations. New storage technologies now overcome those limitations, making tiering an attractive option to reclaim capacity on Tier 1 storage systems and reduce backup costs and time requirements — often resulting in overall file storage cost savings of 50%.

Of particular note, new solutions now enable these savings with zero impact on user data access.

In this paper, we compare the total cost of ownership (TCO) of traditional NAS to the TCO of traditional NAS augmented by file tiering with Cloudian object storage.

KEY TAKEAWAYS

- 60% of data on Tier 1 NAS is typically dormant or very infrequently used
- Object storage provides a durable, cost-effective platform for long-term storage
- New tools provide for transparent migration of inactive information
- Management features ensure zero disruption to user data access
- Overall cost savings exceed 50% of storage and backup costs

Storage Growth Isn't Slowing Down

Aside from the massive explosion in data — from 4.4 zettabytes (ZB) in 2013 to a projected 44 ZB in 2020, and 180 ZB by 2025 according to IDC— the profile of the data has changed. New types of content have emerged. It's often unstructured data, captured at the edge or based on activity that touches multiple systems. Other examples include:

Media Archives

1080p and 4k resolution videos, 360 degree video, and 60 fps content all require significantly more storage than previous generations. A single hour of 4k video requires over 300GB of storage.

Bioinformatics

Genomics and similar scientific research now generates extremely large raw data files (sometimes multiple PB) that then need to be analyzed. The original data sets usually need to be maintained.

Financial and Consumer Data

Fraud detection data streams, high volume transactions, and customer data streams all need to be captured and stored.

Big Data

Analytics outputs, log, machine, and sensor data are just some examples of the massive data outputs coming online.

Engineering Files

Build files, build artifacts, and multiple versions of software for every imaginable product.

Web Content Storage

General web content, including media, need to be stored for access. Older data also needs to be tiered.

User-generated Data

Consumer and business users create their own data — documents, photos, videos, social posts, activity data from wearables and more.

At an estimated \$2 per GB for 3 years of primary storage, this can add up. But the cost of the storage itself is only the beginning.



Why Backup and Replication with NAS is so Expensive

- An initial file or data is created on primary storage.
- Almost as soon as it's created, data is added to the snapshot schedule.
- Snapshots have specific interval retention policies (e.g. hourly snapshots are kept for a week).
- Every change is retained in a snapshot.
- Active files have a daily change rate of about 20 percent, so a 1MB file becomes 2MB of snapshots within a week.
- Backup retention — which is separate from snapshots — adds to that. Daily and weekly backups are retained for months, while monthly backups are retained for years.
- Multiple copies of inactive files live on — for a long time.
- The end result: Over the life of a file, enterprises use 10 to 20 times more storage than the original file size!

The Hidden Overhead: Backups and Snapshots

Backup copies and snapshots can quickly double or triple capacity requirements, leading secondary storage costs to often exceed primary storage costs. Data retention policies — a necessary part of doing business — also mean companies keep a lot of data on expensive primary and secondary storage tiers even when it's inactive. Sixty to 70 percent of data is inactive after 6 months, but copies of files are kept on the primary storage.

Backup retention also has a big impact — monthly backups are usually retained for years, often on expensive backup tiers (for example, Data Domain). Storage costs for backup and snapshots are an ongoing issue. While primary storage can cost \$2/GB over 3 years, it's indispensable for nearline, fast access

Object Storage for Cost-effective, Highly Durable Data Tiering

Two factors now conspire to make data tiering a highly viable solution for NAS-offload. The first is the emergence of highly cost-effective Enterprise Object Storage systems.

Object storage has long been a key part of cloud platforms like Amazon S3, Google Cloud Platform, Facebook, and Netflix. It is now becoming a part of the enterprise data center landscape as well, with multiple vendors offering on-prem object storage systems, or “private cloud” storage.

For enterprises, object storage offers a low cost entry point that can start small and then scale in much smaller increments than traditional storage. Enterprise object storage options range from \$0.005 to about \$0.01 per GB, per month, or about one-third the traditional SAN or NAS systems.

Object storage is also flexible in that it can be configured for the level of data durability, availability, performance, and accessibility to meet an organization's needs. Data durability of nine-9s are achievable, with built in redundancy and cross region replication.

Transparent Access for Migrated Data

The second factor is the emergence of file management software that automates the tiering process to object storage. These solutions, such as Komprise data management software, identify Tier 1 data based on user-defined attributes such as file age, frequency of access, owner, and file type.

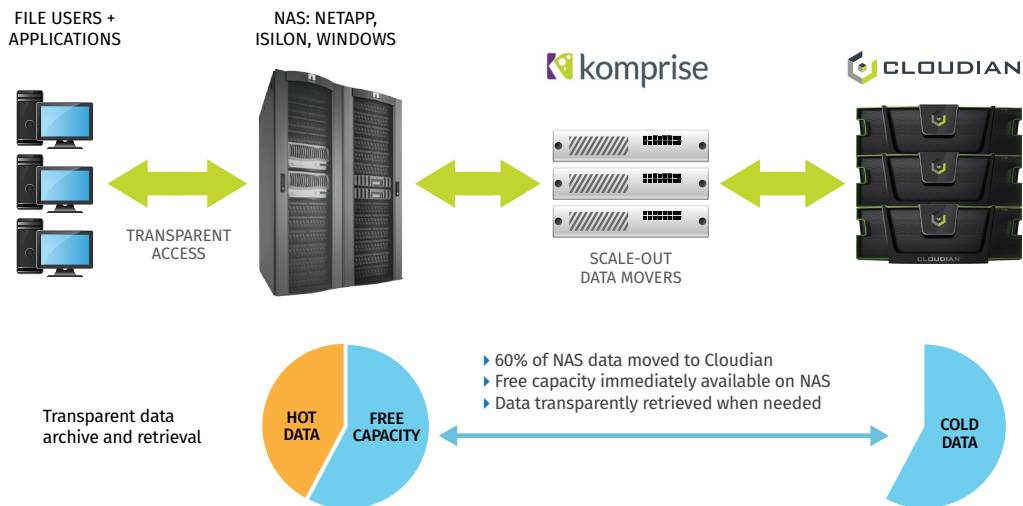
The selected data is then migrated to object storage. If the user requests the data, it is transparently retrieved from object storage. If the user changes the file, it is re-hydrated to the original filer. The entire process is invisible to the user. There is no change in data access.

Object storage has long been a key part of cloud platforms like Amazon S3, Google Cloud Platform, Facebook, and Netflix. It is now becoming a part of the enterprise data center landscape as well.

to active data. But secondary storage can add another \$2 per effective GB — largely because multiple copies of data are created in backups and snapshots. Software, power, cooling, administrative overhead, and maintenance add another 30 to 50 percent to the cost. Even with strong deduplication and compression technology, it's expensive to maintain the data.

Backup software also drives significant cost. It is typically licensed on the capacity being protected, with costs per GB that can approach that of the NAS system itself. Moving that capacity to a secondary tier will have a proportionate savings in the backup software license requirements.

One final issue that every storage manager is familiar with — traditional NAS and backup systems have to be provisioned well ahead of demand, so storage often sits idle for 2 or 3 years.



Clouidian provides cost-effective storage for dormant or inactive files. Komprise data management software identifies files on Tier 1 NAS based on user-defined attributes. Users can transparently retrieve files when needed.

TCO Example: Traditional NAS vs. Clouidian — 54 Percent Savings

To understand the savings associated with object storage, we can create a hypothetical but typical example of a customer environment and evaluate costs over 3 years. The environment has the following characteristics:

- 200 TB of primary storage
- 60 percent of that data is inactive
- 400 TB of secondary storage (1 copy of primary data for replication; 1 or more copies for snapshots and backups)
- Backup and replication software licensed for 200 TB

In the above example, even though 60 percent of the data is inactive, the costs of storing and backing up the data are significant. That doesn't even account for potential data growth, and with total data storage requirements growing at about 30 percent a year, enterprises will need to have access to significantly more storage in the future.

By comparison, Clouidian costs between \$0.18 to \$0.36 per GB over 3 years. Since Clouidian is on-premises storage, there are also no data transfer charges when data needs to be retrieved. In the

3 Year Cost: Traditional NAS & Backup

Primary NAS Storage @ \$1.90 / GB x 200 GB	\$380,000
Secondary NAS Storage @ \$1.90 / GB x 400 GB	\$760,000
Backup Software @ \$2.00 / GB x 200 GB	\$400,000
Total 3 Year Cost	\$1,540,000

3 Year Cost: Move Inactive Data to Clouidian

Primary NAS Storage @ \$1.90 / GB x 80 GB	\$152,000
Secondary NAS Storage @ \$1.90 / GB x 160 GB	\$304,000
Backup Software @ \$2.00 / GB x 80 GB	\$160,000
Clouidian Object Storage @ \$0.36 / GB * 120 GB	\$43,200
File Management Software	\$50,000
Total 3 Year Cost with Clouidian	\$709,200

The savings of \$830,800 translates to a 54 percent lower TCO.

example below, 60 percent of the data has been moved off traditional NAS systems to Cloudian. We assumed \$0.01 per GB per month for Cloudian, though costs can be lower for some implementations. We also included costs for file management software such as Komprise that can automate data tiering to object storage.

The savings of \$830,800 translates to a 54 percent lower TCO.

Benefits of Object Storage for Tiering

Cloudian offers enterprise object storage that seamlessly scales from TBs to PBs in a single storage environment.

Highly Available

Data redundancy is built-in. Every Cloudian cluster is HA with no single point of failure

Efficient and Secure

Compression and encryption features included.

No Performance Bottlenecks

Throughput grows as your cluster expands

Available as Software or as Pre-configured Appliances

Deploy on industry-standard standards, or select fully-supported, pre-configured appliances.

With Cloudian, enterprises can achieve over 50% cost savings compared to traditional NAS systems. Cloudian saves costs in three significant ways:

Reduced Cost Storage

Cloudian provides high-availability, enterprise class storage at 70% less cost than traditional NAS.

60% Lower Backup Cost

Data migrated to Cloudian no longer consumes backup license costs. That data is protected via Cloudian's built in redundancy.

Simple Offsite Storage

If offsite archival storage is required, Cloudian's built-in cloud integration makes it easy to replicate data either to remote Cloudian devices or to public cloud services such as Amazon or Google.

See Your Savings Before You Start

Get a free storage assessment from Cloudian. We will analyze your file storage and provide you with a written estimate of what you will save.

Call Cloudian today to get started, or visit www.cloudian.com/freestorage.



Cloudian, Inc.
177 Bovet Road, Suite 450
San Mateo, CA 94402
Tel: 1.650.227.2380
Email: info@cloudian.com
www.cloudian.com