druva

How to solve for **NAS backup challenges**

Minimize complexity and costs with reliable cloud data protection

Introduction

Network-attached storage (NAS) systems are increasingly popular due to their simplicity of operation, easy backup and recovery features, and quick ability to scale simply by adding disks. A good NAS system can provide most all the benefits of an on-premises private cloud.

Of course, the NAS system itself still needs data protection, which can be costly and complex. Ask admins at enterprises with substantial NAS deployments and they'll highlight these **frustrating obstacles:**

- **Slow performance** — Slow backup and difficult restores cause missed RPOs and RTOs.
- **Security** — Ransomware is a very real threat and is actively targeting backup infrastructure. The security team worries about detection. But the backup team needs to make sure the backup data is isolated and adequately protected.
- **High cost** — Storage and third-party infrastructure costs are unpredictable and rising.
- **Difficult to manage** — Administration and infrastructure management consumes too much time.
- **Backup storage analytics** — Large amounts of NAS data are often old, inactive, or not required for backup. Unfortunately, backup and storage teams don't have the time or the right tools to identify this waste.

The amount, costs, and complexity of data that people bring into their NAS systems makes data protection and management a unique challenge. Fortunately, there is a way to effectively protect a NAS system and ensure you get the greatest value from its data.

This eBook explains why NAS data is important and what the requirements are for NAS data protection. You'll also get recommendations for a simple, reliable, and cost-efficient solution. Once implemented, your organization can increase the predictability and transparency of costs, improve business resiliency, lower TCO, and enhance financial flexibility.

Why NAS data is so valuable

All of your enterprise data is valuable — but it's vital you understand how your NAS data is different, and why it's some of the most important data in your environment.

NAS servers hold huge amounts of data, most of that data changes infrequently, and many users and applications store different types of data. As a result, it's very difficult to understand what type of data you have on your NAS system. **That means:**

- **Cyberthreats like ransomware are particularly insidious** because they seek to exploit devices on your corporate network. If someone's uploaded trouble in an unused file they want to archive, it could be days, weeks, or even months before the problem surfaces.
- **With more-and-more privacy and data management regulations, any data mismanagement in NAS backups can become a costly liability.** If you lose any kind of sensitive or personal data, you may miss the loss — but a regulator won't, and the fines can be stiff.
- **That same regulator will also ask if you're complying with data sovereignty laws —** for application and backup data.
- **New workloads are loading up your NAS systems** with data from custom and one-off applications specific to your vertical.

There's always more data for your NAS!

What makes NAS data protection different?

It's the files. If you have hundreds or thousands of VMs, dozens or hundreds of databases, thousands of snapshots, thousands of users, hundreds of applications, and custom applications that rely on NAS, you have billions of files. That means petabytes of data, a considerable haystack. Protecting data for basic recovery, disaster recovery (DR), and compliance purposes on a daily, weekly, or monthly basis takes on a new meaning at this scale. The traditional approaches to NAS data protection and long-term retention (LTR) no longer apply.

The haystacks of data in your NAS systems are bound to include data that includes rich intellectual property and that is subject to multiple regulations, from PII and health information to financials and legal-hold data. To maintain regulatory compliance and the data's intellectual property, the data needs tight management, and with billions of files, that means managing metadata.

You have more metadata in your NAS systems, more information about your information, than anywhere else in your enterprise. Indeed, the underlying culprit to slow NAS data backup is metadata. How do you quickly find new or changed data in a file system with 100s of millions of files? While native NAS storage snapshot technology can identify changed blocks, backing up NAS to a different storage medium for off-site protection, storage optimization, and lower-cost, long-term compliance is challenging. It's managing this metadata that makes NAS data protection fundamentally different from traditional backups and disaster recovery.

Given the volume of unstructured data stored on NAS, organizations need three key items for effective data protection:

- **Rapid, reliable data protection**
- **Long-term retention**
- **Visibility into the NAS environment**

What it takes to protect NAS data

Rapid, reliable daily protection is the first requirement for NAS systems. If you can't protect data quickly, the rest doesn't matter. Meeting this daily requirement, especially at scale, requires a fast backup technology that leverages optimized, incremental-forever backups, and flexible recovery options — simple file recovery and rapid full recovery. However, organizations also need to move a copy of this data offsite for both DR and ransomware protection.

Organizations typically accomplish this using snapshots, snapshots and replication, or snapshots and some form of disk backup. Storage snapshots are built into the NAS system, and they are fast, simple, and reliable. However, over time, you can end up managing thousands of snapshots. Getting these daily backups into secure, offsite storage and off the primary storage system is done with disk-to-disk backup solutions, or by sending backup data to the cloud. Whatever solution you're looking at, it has to embrace the metadata and the scale of the system.

Rapid, reliable daily protection

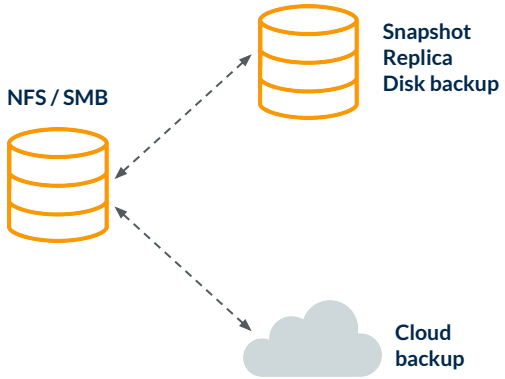

What it takes to protect NAS data

Long-term retention of NAS data has its own set of requirements. It must be low cost, offsite (for ransomware and DR protection), durable for many years (5-100 years), and searchable. Because the data is diverse and unstructured, many organizations keep all of it because they aren't sure what to exclude and the costs of deleting data too early can be expensive.

But no one wants to hold years of snapshots on a NAS system or a backup appliance. That leaves tape or the cloud as options for long-term retention. Tape is still alive for providing low-cost, off-site cold storage. However, cloud solutions such as AWS Glacier deep archiving typically outperform both tape and backup appliances.

Long-term retention

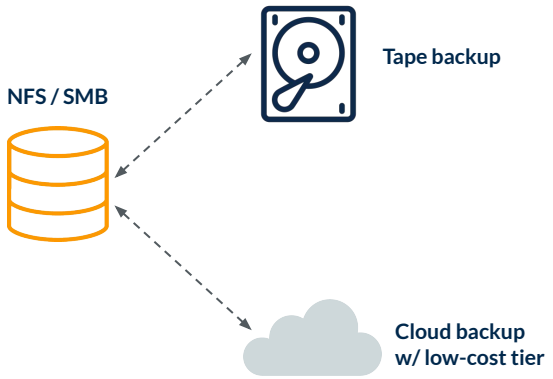

What it takes to protect NAS data

The third requirement for modern NAS data protection is visibility. You can't just put data away somewhere and never think about it again. Visibility into your NAS data means understanding growth trends, monitoring for cyberthreats, and staying compliant with a myriad of regulations.

By enabling analytics, your enterprise can actually take advantage of its haystacks of data. Analytics can detect unusual data patterns and alert you to ransomware. It can manage regulated content. Analytics provides trends and other analyses for smarter business decisions.

The two approaches for analytics are one, using a standalone tool which scans primary storage or secondary storage systems and two, leveraging cloud backup with built-in metadata analytics. If you're putting NAS data in the cloud today for offsite or long-term retention, you should investigate how you can apply analytics to that pool of information to gain greater visibility. Performing analytics in the cloud is not a bad idea considering some of the most powerful machine-learning and AI high-performance computing (HPC) analytics infrastructures are cloud-based today.

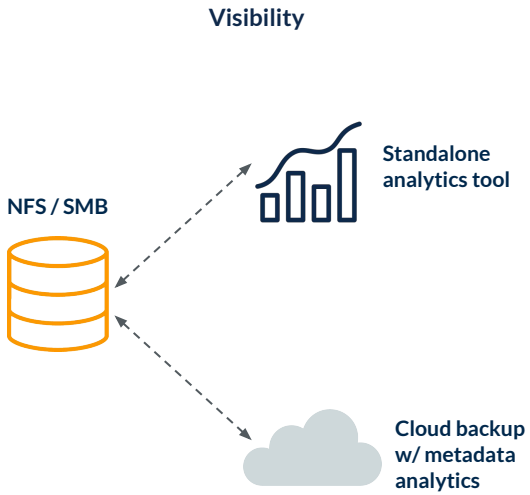

NAS data protection requires a SaaS approach

Instead of thinking, how do I solve for high-performance NAS backups in order to get data offsite? How do I solve for long-term retention? How do I solve for metadata analytics and visibility? Instead of approaching the NAS data protection challenge as three independent issues, approach it holistically. Ask how can I achieve these goals with less complexity, infrastructure, and cost. The answer is in the cloud. Cloud data protection offers the flexibility to augment or replace your current NAS data protection approach without the same overhead or investment required with traditional approaches.

- **Your existing NAS system offers local snapshots and replicas that are fast and easy for short-term recoveries.**
- **For short-term NAS backups and long-term retention, a SaaS solution can automatically tier backups across storage types, cost-effectively keeping some data warm and other data cold.**

With a cloud data protection solution, analytics can extract the greatest value from all your data and metadata, using either your own resources or platforms such as Amazon Athena, Kinesis, or Redshift. Beyond the NAS system, there's no hardware or network management on your end. When you have to scale up, you don't have to buy and install more disks for your NAS. You can automatically move storage between warm and less-expensive cold tiers. And, high-performance computing resources for your analytics are virtually unlimited.

Although there are challenges to overcome, there is one effective solution you can turn to — Druva.

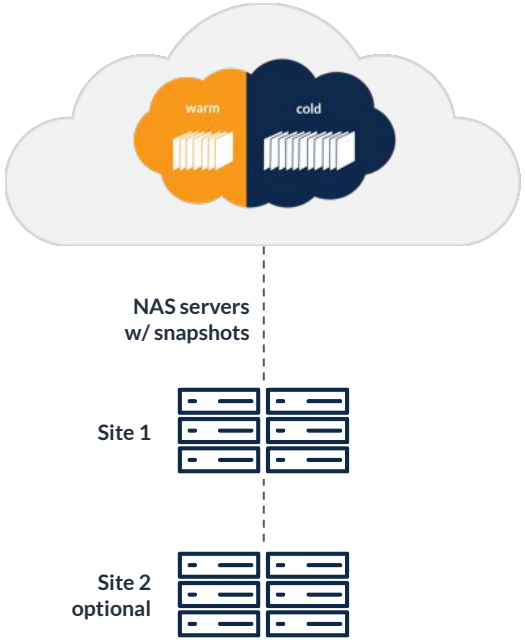

How Druva NAS data protection works

For short-term backups, NAS systems are fine. It's long-term retention, cost-effectively juggling cloud resources between warm and storage with fast recovery, where Druva NAS data protection is unmatched. Some data can be tiered and go to warm storage to retain fast, granular file recovery. It can then move to cold storage later after a specific period of time or a period of no access. Other data that is infrequently accessed can go directly to cold storage. Druva manages it all — you don't have to set anything up, and you just pay for the storage you actually use.

The key to these efficiencies is the Druva cloud storage system which leverages a metadata-centric architecture to enable deduplication, security, and analytic capabilities. It effectively manages billions of files independently, the petabytes of data from thousands of users and hundreds of applications that you need to store — the reason why NAS data protection is different from protecting VMs or databases.

A cloud backup platform with a metadata-centric architecture also enables precious insight about what's happening in your NAS backup environment. For example, analyzing metadata can show trends such as how different data types are growing and where data growth is occurring without the need for additional software, hardware, or work from your team.

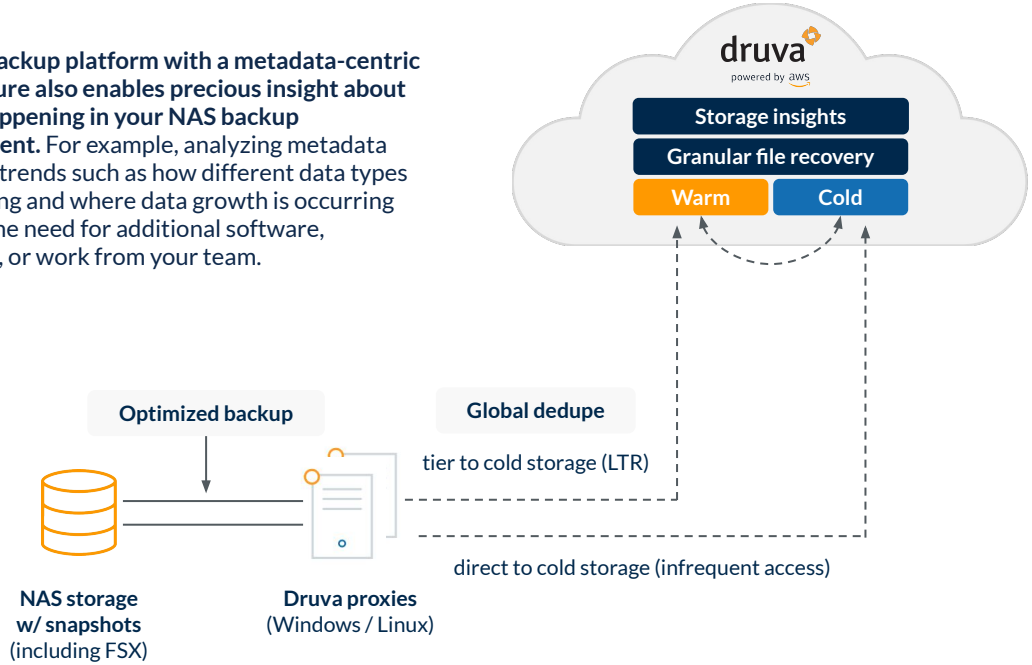

Takeaways

Your organization will undoubtedly face cyberthreats, regulations, and omnipresent data growth. Along with these obstacles, you'll encounter unique NAS challenges with billions of files and diverse, unstructured data types generating an enormous quantity of metadata. You'll need short- and long-term protection as well as clear visibility into your environment.

If you approach the challenge of protecting NAS systems as individual problems, you may miss seeing how they all are, in fact, tied together. It's important to seek out a third-party data protection solution like Druva, that can help overcome your NAS data protection and retention challenges. Combining your NAS system snapshots and replication with Druva cloud backup will streamline your environment, lower your TCO, and simplify long-term retention.

Druva provides a cloud-based data protection service, ranging from backup and recovery to providing cyber resilience. You'll receive all-inclusive services with no need to manage hardware, software, or the associated cost and complexity — enabling you to improve business resiliency, optimize financial flexibility, and increase the predictability and transparency of costs.

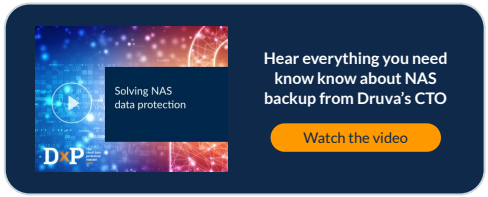

aws marketplace

Find Druva in AWS Marketplace

Get started

druva

Sales: +1 888-248-4976 | sales@druva.com

Americas: +1 888-248-4976
Europe: +44 (0) 20-3750-9440
India: +91 (0) 20 6726-3300
Japan: +81-3-6890-8667
Singapore: +65 3158-4985
Australia: +61 1300-312-729

Druva™ delivers data protection and management for the cloud era. Druva Cloud Platform is built on AWS and offered as-a-Service; customers drive down costs by up to 50 percent by freeing themselves from the burden of unnecessary hardware, capacity planning, and software management. Druva is trusted worldwide by over 4,000 companies at the forefront of embracing cloud. Druva is a privately held company headquartered in Sunnyvale, California and is funded by Sequoia Capital, Tenaya Capital, Riverwood Capital, Viking Global Investors, and Nexus Partners. Visit Druva and follow us @druvainc.