

Guidebook


# **The 9 Challenges of NAS Backup and how to overcome them**




# Executive Summary

For ONTAP users, backing up network-attached storage (NAS) data at scale is a challenge, but that doesn't mean it should take days or weeks. In this guidebook we take a closer look at NAS backup technologies, their limitations, and the reasons **why they cannot always meet an organization's backup needs**.


We discuss:



The nine challenges of  
NAS backup



Why organizations  
fail to meet their  
backup objectives



How to backup  
your PB-scale data  
100X times faster

We will also introduce a great alternative, [NetApp Cloud Backup](#): the ONTAP native solution to backing up NAS data at enterprise-grade scale.



# Table of Contents

Executive Summary	2
Storage Admins Are Getting Overburdened with Backup	4
The 9 Challenges of NAS Backup	6
How to Meet Your Backup Needs 100X Faster: Solving NDMP-Based Backup Challenges	12
Cloud Backup as a Solution for NAS Challenges	21
Case Study: Multinational Insurance Company Backs Up Petabytes of Data to Azure	23
Don't Let NDMP Backups Slow You Down. Try Cloud Backup	25



# Storage Admins Are Getting Overburdened with Backup

Consider what the average storage admin's backup workload looks like. Let's introduce you to Anne. Anne is a storage admin at a large American financial services company, where she is responsible for the backup (and recovery) of **10 PBs of file shares of business-critical data**. The organization has been using network-attached storage (NAS) as its storage architecture since NAS was **first introduced almost 25 years ago**. The NAS backups are streamed to tape media over the company's IP network using the Network Data Management Protocol (NDMP). In order

to meet the organization's RPO requirements Anne has to carry out one full backup per week. Incremental backups (only files that have changed) are performed twice daily. **Anne spends almost the entirety of her workday overseeing these backup operations**, testing the backups, and ensuring that recovery procedures meet the company's RTO requirements. Because so much of her time is invested in these backups, other storage-related tasks that need her attention might wind up being neglected.

*The organization has been using network-attached storage (NAS) as its storage architecture since NAS was first introduced almost 25 years ago.*




# Storage Admins Are Getting Overburdened with Backup



Here are some of the questions that concern Anne on a daily basis:

- Did the last backup succeed? How can I know that the backup is viable for recovery?
- Will the back up take more time than planned, due to network congestion or growing volumes of data?
- Did the backup finish before the next backup window? If not, then the organization's business-critical file share is not properly protected.
- When is my next maintenance window and how can I leverage it to close backup and recovery gaps?
- Why did the backup from two weeks ago affect the production environment?
- How do I integrate a new workload?
- How long will it take me to restore a single file?



Anne  
Storage Admin

So now that we've seen some of the day-to-day issues that Anne has to deal with, let's figure out what's at the root of the problem.

**What are the backup challenges that are making Anne's work so difficult?**

# The 9 Challenges of NAS Backup



# The 9 Challenges of NAS Backup

Like most storage admins, **Anne has been backing up NAS data for a while.** There are a number of different ways she can do this, from using NDMP to leveraging other storage systems for backup purposes. But things are getting more complicated for Anne since her company, like many other organizations, now has a greater amount of data on hand to back up. Her NAS data volumes are now easily in **petabyte-scale dimensions and the previous approach is no longer effective**, whether it comes to costs, time, or reliability.

There are nine major challenges to backing up NAS, especially when dealing with large volumes of data:

## \$ High overhead costs

The NDMP-based NAS backup architecture requires at least one dedicated management server on which the backup software is installed. Installation of the backup software, training on how the NDMP-based software works, and maintenance of this backup server all **contribute to additional admin work.** To use secondary systems for backup takes considerable CAPEX spending, which is unrealistic considering the scale.

**To find more about the hidden costs of NDMP-based NAS backup, [check out how the costs add up using our TCO calculator here.](#)**

## ⏮ Slow performance

NDMP was designed for single stream backups per NAS node. This can slow down performance if a backup application does not have a multiplexing solution or if dealing with an NDMP version that doesn't support multiplexing. Without performance optimization, **backing up 10 TB of NAS data using NDMP can take as long as two weeks.**

*Anne's NAS data volumes are now easily in petabyte-scale dimensions and the previous approach is no longer effective, whether it comes to costs, time, or reliability.*



# The 9 Challenges of NAS Backup

## ♥ Reliability

Due to the long time it takes to back up large amounts of NAS data, backup windows will be missed. This means that the backups that are retained are likely to have significant amounts of data missing. This can make **NDMP-based NAS backups an unreliable way to replace the primary data** set if something goes wrong.

## 🕒 Inefficiency

NDMP-based NAS backup requires recreating the full, baseline copy of the dataset on a fairly frequent basis. A full backup takes a lot of time and can encroach on the next backup window. ONTAP users who rely on NDMP-based NAS backups also face the loss of ONTAP's cost- and space-reducing storage efficiencies, such as data deduplication and compression. **Without these efficiencies in place, backups take longer to create, longer to transfer, and consume more storage space**—all of which add up to increased costs.

## ⚡ System stress

**NDMP-based NAS backups can put a lot of stress on the production environment** due to the amount of processing required to create backups. The increased computational capacity that is needed to prepare backups can be especially challenging when it's done within mission-critical production systems, since it negatively affects the overall system performance. Using a secondary system dedicated for backup to avoid this problem can become very costly.

***“backups take longer to create, longer to transfer, and consume more storage space—all of which add up to increased costs.”***





# The 9 Challenges of NAS Backup

## Restore issues

**NDMP-based NAS backups have poor restore performance** when restoring from long chains of incremental backups. That's why it is recommended to perform one full backup for every nine incremental backups. This best practice puts a strain on network resources and is very inefficient for large volumes of data.

## Infrastructure constraints

**Infrastructure constraints can limit the effectiveness of backup solutions.** In many cases, this is simply a problem of scale: using NDMP or secondary storage systems on-prem will simply **not give you the scale-out capabilities of the cloud.** In other cases, the topology itself is preventing it.

Topologies supported by NDMP-based NAS backup vendors require backup media to be directly connected to the NAS server, to the NDMP MediaAgent server, or to a remote third server connected to

the main NAS filer. Indexing for NDMP solutions also requires a robust infrastructure, including the storage required in the local server, which will also consume network bandwidth. Managing all of this infrastructure is another considerable challenge for Anne.

## Security

NDMP doesn't encrypt backup data while it's in transit, only when it's at rest. Without end-to-end encryption, data is exposed while it's being transferred via the network, making organizations vulnerable to data theft, viruses, ransomware attacks, and other cyber threats.

# The 9 Challenges of NAS Backup

## ↻ Additional NAS Backup Challenges


Besides NDMP-based NAS backup challenges, using secondary ONTAP systems for backup come with some specific hurdles:

### **The difficulty in meeting a 3-2-1 backup strategy**

This goes for any other strategy that involves data mirroring, multiple backup media, and offsite storage. The 3-2-1 strategy is to keep three copies of the data: two in different formats, and one copy offsite. Backing up to another ONTAP system doesn't meet the requirement to use a different format. (Note: NDMP-based solutions also fail here since it requires a VERY complex operation that includes backing up to tapes and disks, then physically sending them to other locations)

### **Maintaining the storage resources of the backup destination**

Consider all the upkeep and costs for a primary ONTAP. All of the initial spend, power, floorspace, and other routine maintenance costs of an ONTAP system are duplicated for a redundant system for backup copies. Plus, the primary and secondary systems will have to be the same type of systems: that makes upgrading complicated.



# The 9 Challenges of NAS Backup

**In summary, Anne and her fellow backup admins can't rest easy when relying on NDMP-based NAS backup for large volumes of data, she also can't count on leveraging a secondary storage system that isn't optimized for backup.**

Even if she follows all the backup best practices and procedures, if a failure takes place, Anne might not be able to recover data fully or in a timely manner. In a time when ransomware and data privacy have become major issues that's a considerable risk to take.






# How to Meet Your Backup Needs 100X Faster

We've seen how NDMP-based backup solutions can present some serious challenges when it comes to creating quick, dependable backups. But there are solutions to these shortcomings that have been developed.

# How to Meet Your Backup Needs 100X Faster

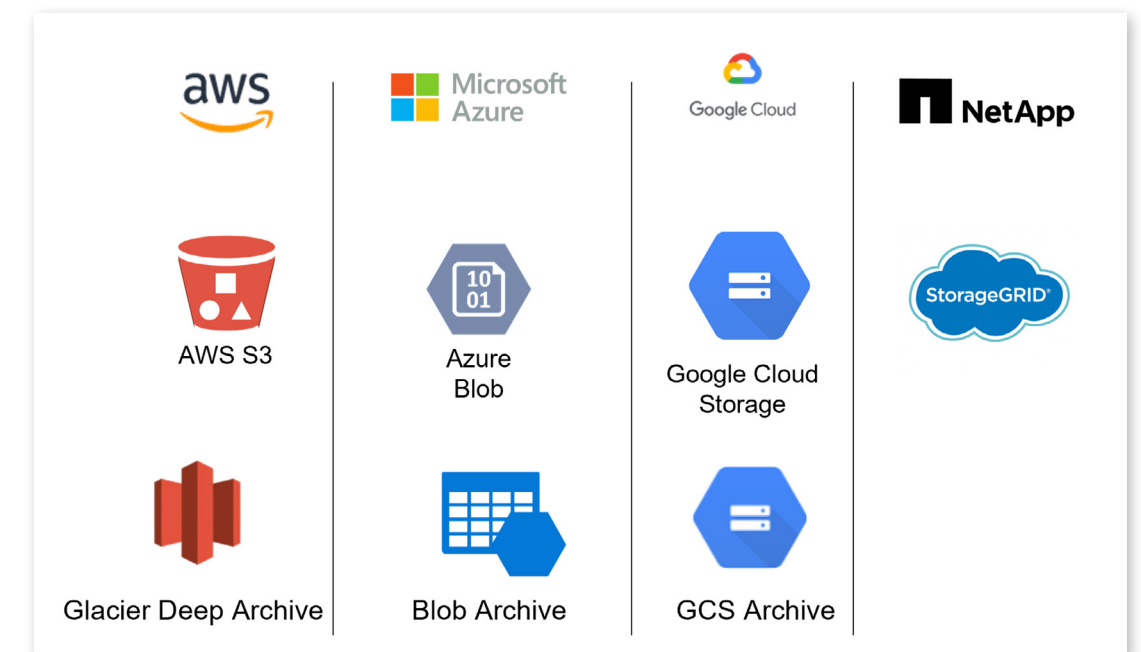


## Leveraging object storage

Object storage is highly valued for its ability to inexpensively store any amount of data, with differing access tiers. The value to an infinitely-scalable storage service can clearly be seen given that object storage was the first product that Amazon made available when it began to offer AWS, in the form of Amazon S3. It's possible that the entire cloud revolution has been thanks in part to object storage.


For backup, **object storage is particularly useful in that there are multiple ways to use it.** Active data can remain stored inexpensively, colder data can be tiered to archive storage, and it can be stored redundantly in more than one region by default, making it reliable (the average SLA for object storage is 99.999999999%). Object storage can also be leveraged inside data centers without the need for the cloud at all. This is useful for creating private clouds, which is relevant when there is the strictest need for privacy and security.

If your primary storage is ONTAP-based, **leveraging object storage for backup also helps meet the 3-2-1** backup criteria of having a copy in a different format. If it's object storage in the cloud, it meets the criteria for keeping an offsite copy.



*There is a wide variety of object storage options.*

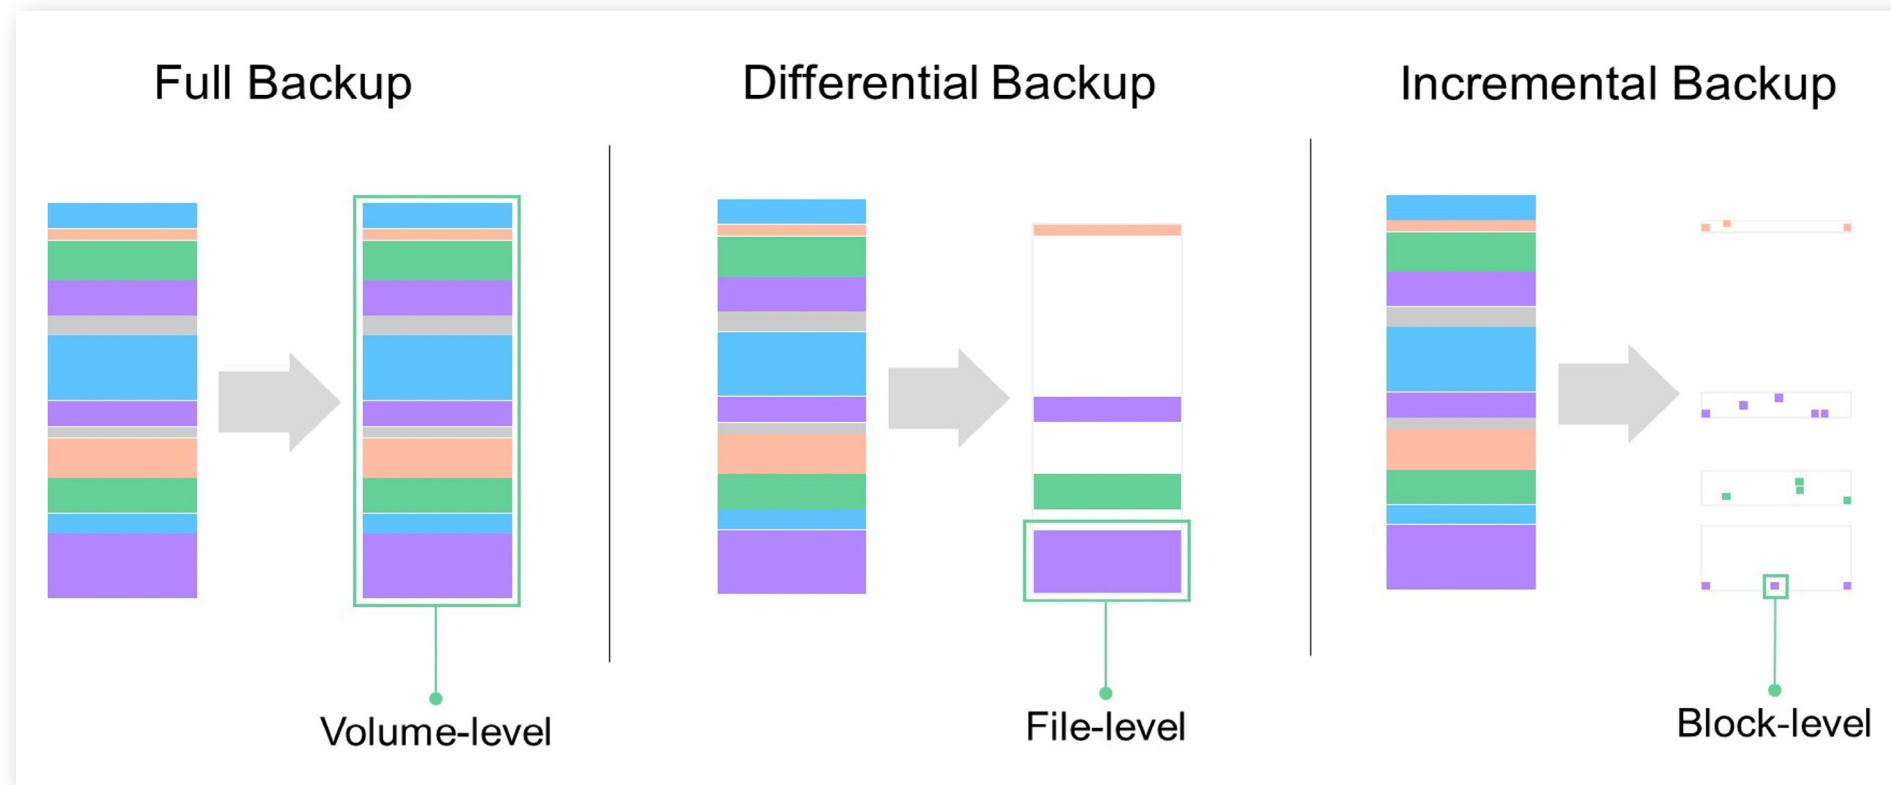
# How to Meet Your Backup Needs 100X Faster



## Incremental forever technology

This powerful technology tackles one of NDMP's biggest limitations which is the required periodic full backups. **Incremental forever backups mean less data traffic on each backup, not having to schedule time-consuming periodic full backups subject to errors, meaning there's a much lower chance of missing a backup window.**


With incremental forever, there's no need to worry about taking a full backup for every other nine incremental updates like with NDMP-based solutions. Incremental-forever chains allow users to skip long and error-prone full backups. This makes life easier for data protection teams and storage admins, especially when dealing with large datasets.



*Creating full backups and file-level backups takes longer than incremental, block-level back up.*



# How to Meet Your Backup Needs 100X Faster




## Block-level backups

**Block-level backups only update data blocks that have been added or changed since the last backup. This makes average scanning times and backup times much shorter than with NDMP.**


This advantage becomes even more valuable when the number of files in the file system is high or when there is a lot of unstructured data, such as media files or all types of office documents. In an era of scale-out NAS this is definitely beneficial. These faster backup windows apply both for full backups and incremental ones.

Faster backup times allow for more frequent backup cycles. This means that **Recovery Point Objective can be shortened and restores can be even more granular.** Combined with incremental forever technology, block-level backups can be created in nearly no time at all, virtually eliminating the backup window. Files are instantly updated on demand. **This is a huge advantage over NDMP backups.**



*Full backups vs. incremental backups*

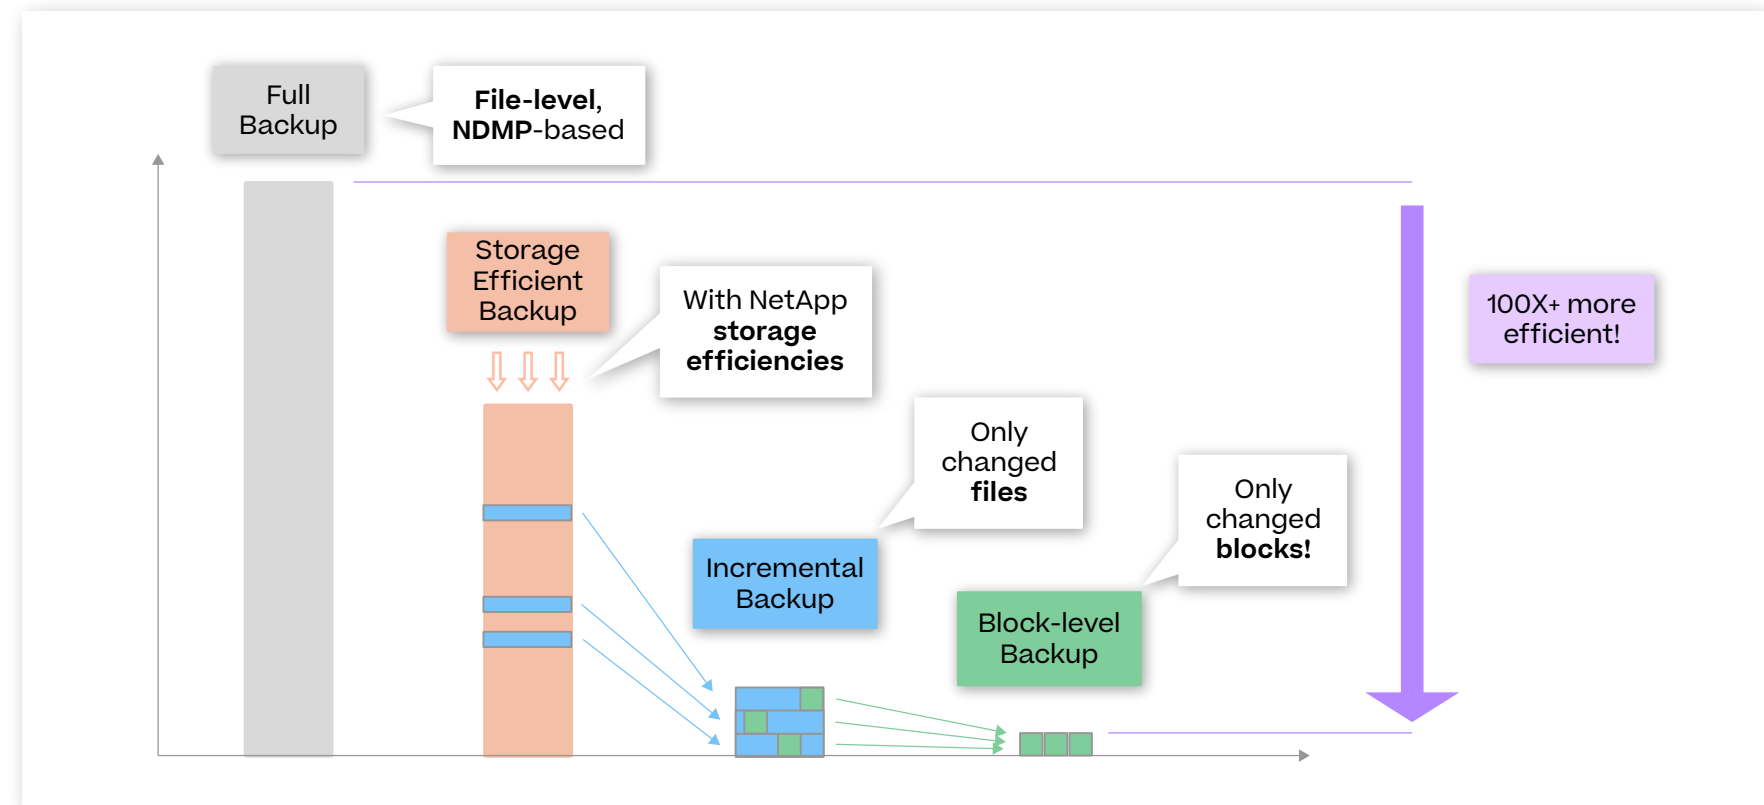
# How to Meet Your Backup Needs 100X Faster



## Retaining storage efficiencies

NetApp users rely on the storage efficiencies ONTAP provides that keep storage footprint and costs low for their primary data sets, and the same is true of their backup data. NDMP isn't able to retain these efficiencies, but there are technologies that **can make sure backup copies of the data are as efficiently stored as the source data. This makes restores faster than with file-level backups.**

Retaining these storage efficiencies not **only reduces data storage footprints at the destination but also reduces the network bandwidth** when transferring the data to the cloud. Compare this to an NDMP-based scenario where a separate index server has to be kept in order to perform deduplication on the data before transmitting it to the backup media. That would represent an additional layer of management even if it is a product aligned with ONTAP's deduplication and compression.



How storage efficiency affects backup file size.



# How to Meet Your Backup Needs 100X Faster




## Cloud and managed service (SaaS) advantages

Storage admins may balk at the idea of relying on the cloud and cloud-based managed services, but this is probably the worst approach to take when it comes to ensuring data is backed up properly. Using the cloud, someone like Anne is able to do everything she wants to do with NDMP-based solutions without the difficulty and without the missed backup windows.

From a cost perspective, it means trading CAPEX for OPEX, so backup can be separated from big budget purchases. But the availability of managed services is where storage admins can get the most benefits, as these will remove many of the repetitive and time-consuming tasks that made it harder to do their jobs. **This includes automatic updates, scheduling capabilities, automatic alerts, and much more.**

**The cloud as such offers a new tool in the storage admin's kit, one that is powerful, cost-effective, and innovative.**



## Easy restores

One of the drawbacks of NDMP is the inefficiency of its restore processes and inability to restore properly after many incremental backups take place. This defeats the entire purpose of incremental backup, since the only way to work around it is to create new, full backups on a frequent basis.

A better solution is a technology that **can restore much more effectively, with incremental backups posing no issue.**

**Restores can be done to a different ONTAP system—both in the cloud and on-prem—and can be done at a volume level or at a file level through a browsable catalog.** These restores preserve the original access permissions on the files. There's no need for a client, extra permissions, or a dedicated service account.



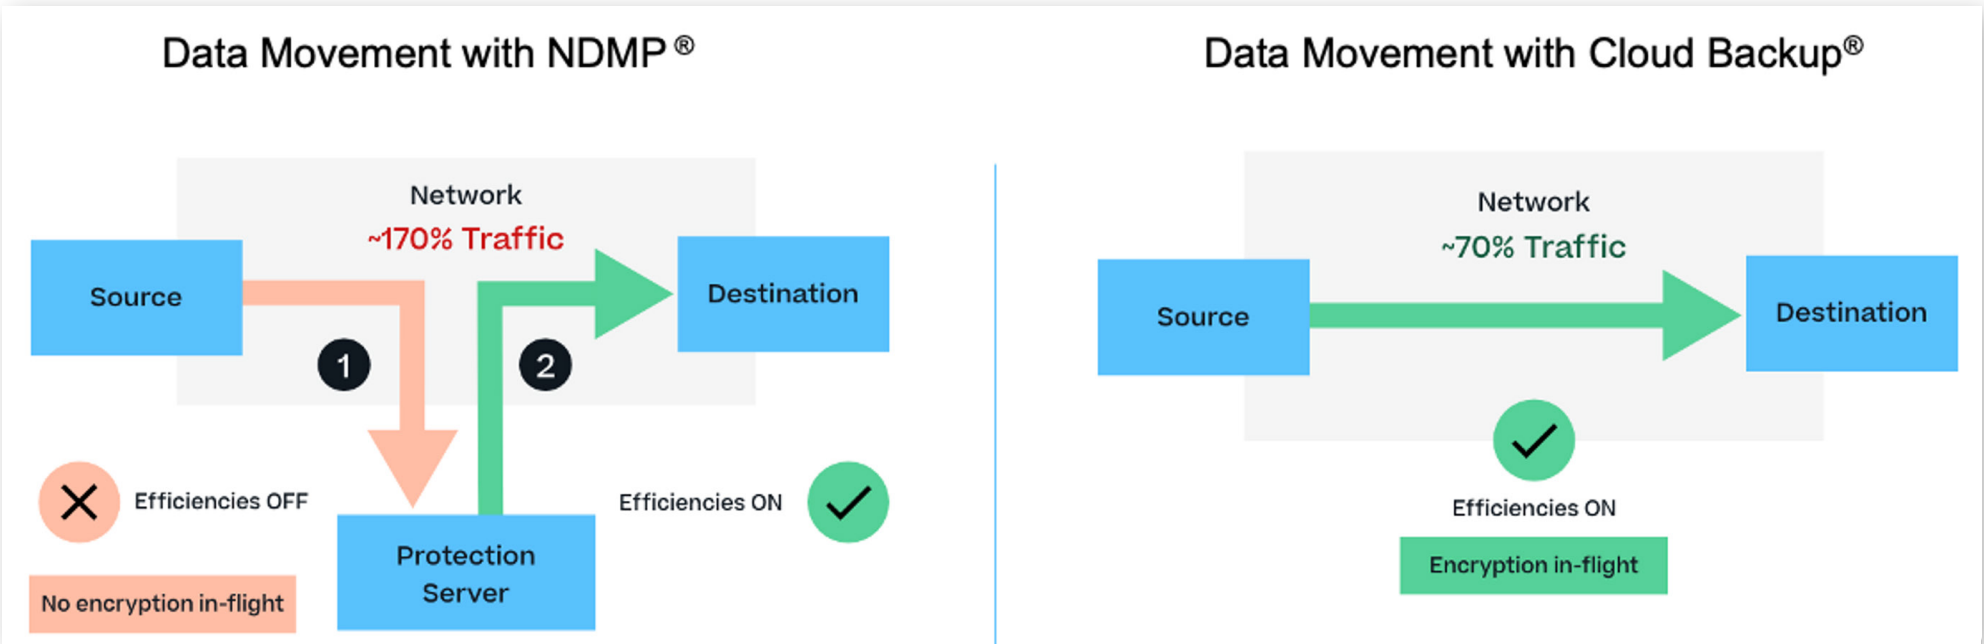
# How to Meet Your Backup Needs 100X Faster



## Direct backups


It's important to be able to directly create backups from the original data source to the backup target on object storage, **with no intermediary**. This allows backup apps to work with the major public cloud providers, offering the option to move the data to cloud object storage and lower-cost archival storage tiers.

Working with this more streamlined architecture eliminates the stress that NDMP-based backup solutions put on applications. Since backups are easily automated, this also can reduce the complexity in existing backup creation processes.



*No middleman: Back up directly from ONTAP to the destination*

# How to Meet Your Backup Needs 100X Faster

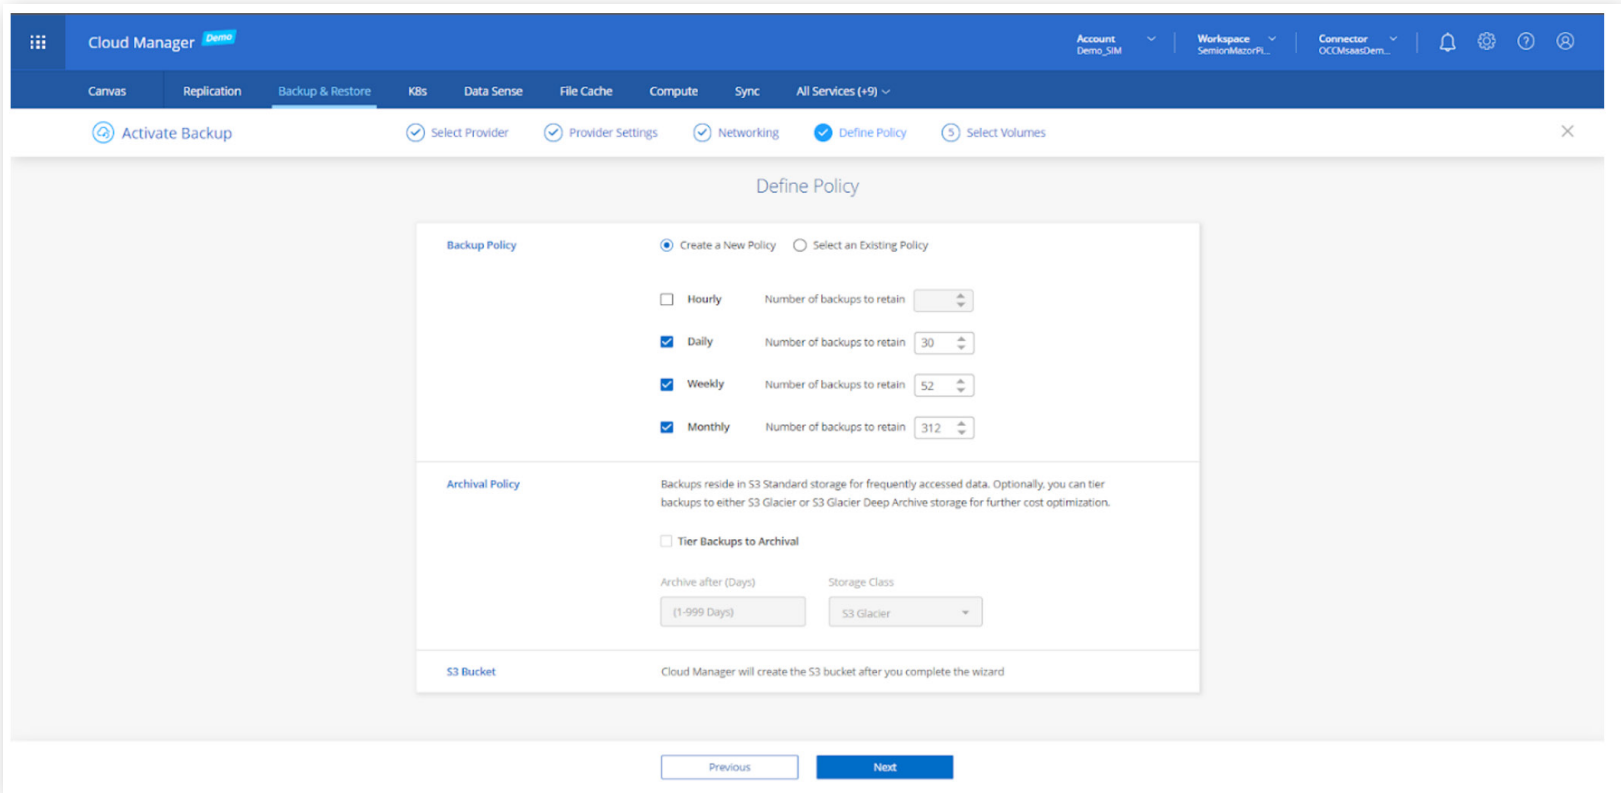



## Automation and Ease of use

Backup apps now have **easy-to-use and unified web-based UIs** which allow users to manage data protection schemes from a single control pane. Automatic backup schedules and policies mean it's possible to meet the most demanding recovery objectives without additional overhead. **APIs can also make it possible to carry out these tasks programmatically**, allowing an infrastructure-as-code (IaC) approach to backup.

**Ease of use is the key here.** Whether relying on code or a GUI for backup, ONTAP users already have set interfaces that they use to manage their clusters. In large organizations, there may even be additional user interfaces to juggle. The last thing that a backup solution should require is the addition of one more window to keep open at the workstation.

Again, this isn't about making Anne feel like there's less that she has to do, but about giving her more capabilities to get her job done than she had before.





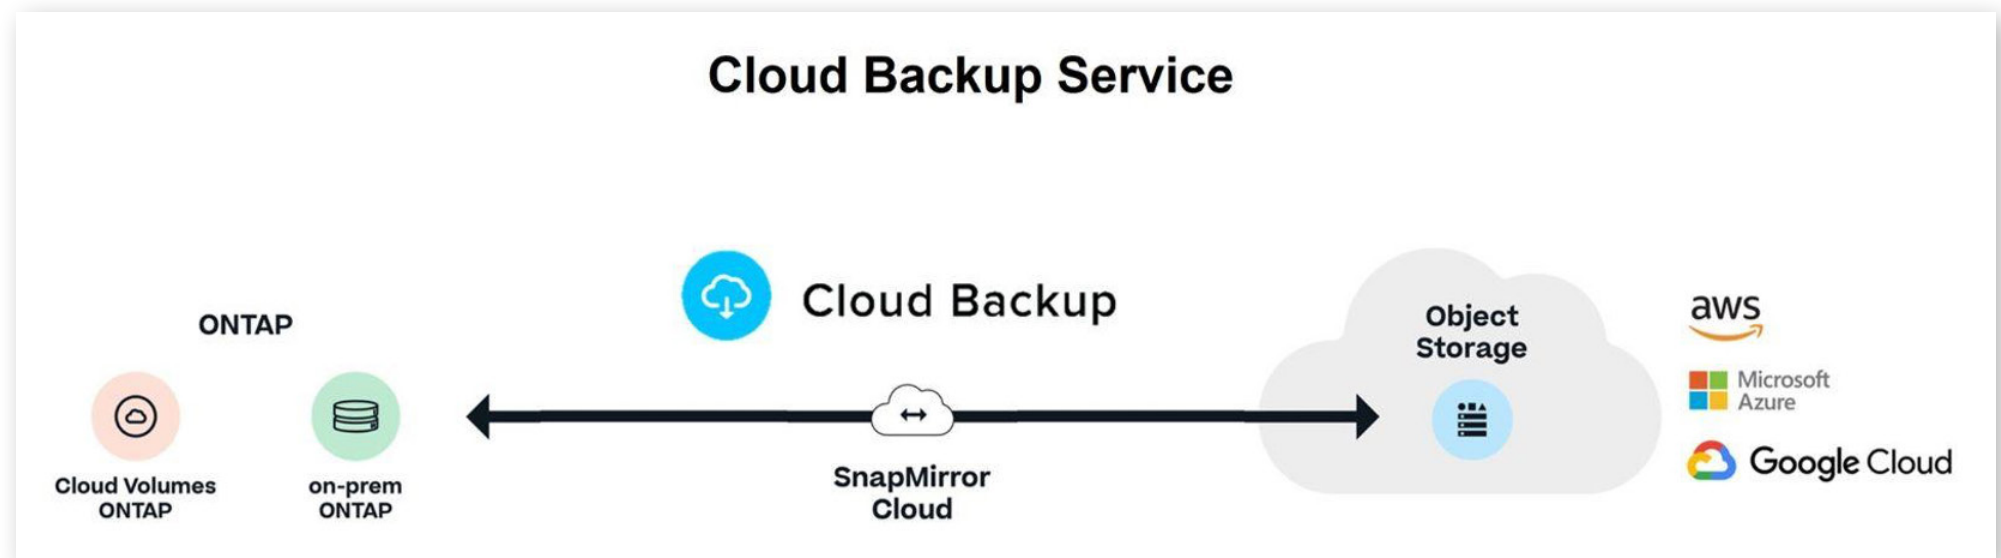
As you can see, NAS backup options have evolved and overcome NDMP limitations. Some of those offerings have even gone post-NDMP to satisfy NAS backup requirements, resembling data management solutions more than just NDMP backup applications.

**It is in this context that we want to introduce NetApp Cloud Backup, which is the best option, even in this competitive context, for backing up ONTAP-based NAS.**



# Cloud Backup as a Solution for NAS Challenges

NetApp Cloud Backup creates block-level, incremental forever backups that are stored in object format in the cloud on AWS, Azure, or GCP or on-prem in NetApp StorageGrid appliances. The Cloud Backup service offers all of the advantages of a modern NAS backup solution fully incorporated into the Cloud Manager ecosystem and the NetApp ONTAP environment.



*The Cloud Backup architecture*

# Cloud Backup Benefits

## **Incremental forever, block-level backups.**

With incremental forever backups, Cloud Backup cuts down operational complexity, shortens the backup window, and reduces costs.

## **Retaining storage efficiencies.**

ONTAP users can be assured that the same efficiencies that keep footprint and costs low for their data are retained in the backup copies of the data. This makes restores faster than with file-level backups.

## **Full ONTAP integration.**

Instead of juggling between multiple vendors, Cloud Backup is an integral part of the NetApp ecosystem, and works seamlessly with ONTAP.

## **Direct backups.**

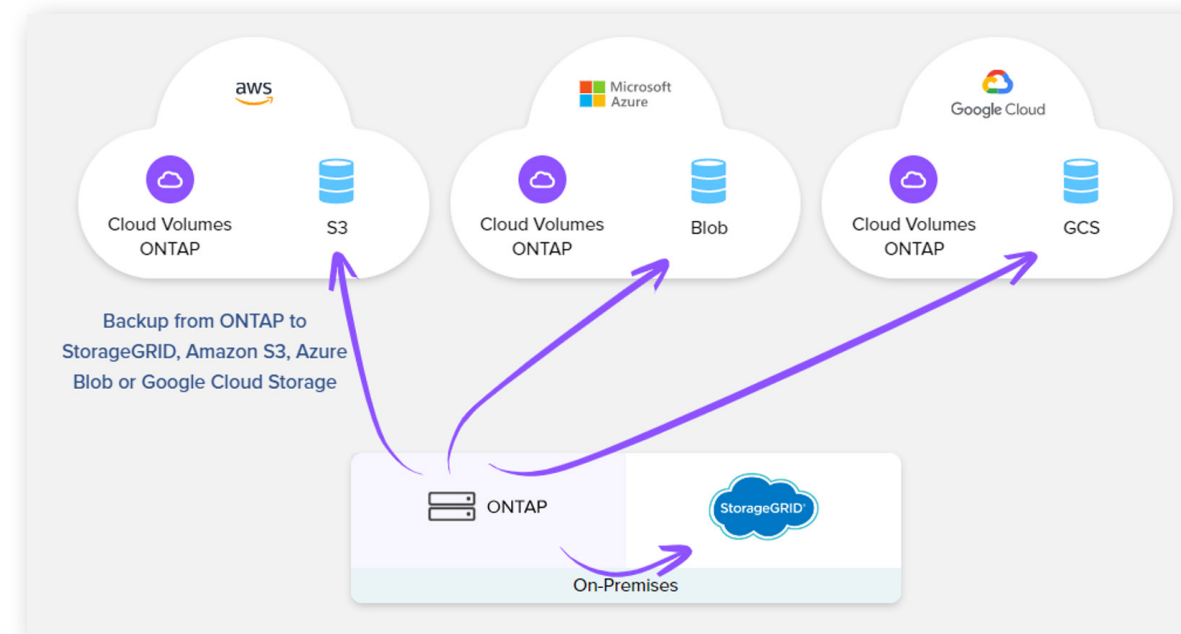
Cloud Backup copies are directly created from the original data source to the backup target on object storage, with no intermediary.

## **Set and forget.**

Automatic backup schedules and policies mean that Cloud Backup helps meet the most demanding recovery objectives without additional overhead.

## **Fully Managed or Software-only.**

Cloud Backup is also available through a software-only option, where no inbound internet connectivity is needed. All updates to the software are done manually, and users are able to leverage backup for the mostly restrictive privacy concerns, such as in private cloud dark sites.



*Cloud Backup options*



# A Cloud Backup Case Study

## Multinational Insurance Company Backs Up Petabytes of Data to Azure

This US-based multinational insurance company offers plans for general liability, workers' compensation, business owners, cyber liability, and more. This company is leveraging Cloud Backup Service to back up critical on-premises data to Azure.

The company relies on its on-premises ONTAP storage systems to manage their considerable volume of application data and legal files. This data needs to be backed up—as an insurance company, backups of data need to be retained for 7 years or more. To meet this requirement, the company had previously been relying on an NDMP-using vendor, which would send the backups to Azure.

This solution was not efficient for the company. Due to the large amount of data—some 1.3 PBs of data—backups were slow, and in some cases backup windows were missed because of restarts and backup operational failures. The impact of these issues was that they were not able to meet their RTO and RPO service levels. They needed a better solution, and they found it in NetApp Cloud Backup.





# A Cloud Backup Case Study

Cloud Backup offered a better solution for this company because it was able to provide them with a number of clear benefits:

## **Familiar technology**

The company's IT architecture already relies on NetApp SnapMirror® and SnapVault® technology, two technologies that they could extend to the cloud with NetApp Cloud Backup.

## **Lightweight solution**

Traditional NDMP backup solutions require multiple appliances that need to be sized in advance and managed. Cloud Backup is activated with one click via Cloud Manager, and it is a SaaS offering with no sizing or ongoing management overhead.

## **Faster Backups**

The previous NDMP solution could only back up entire files and required a two-step process, sending backup files to a mediator before they could reach the cloud. Cloud Backup Service in comparison only backs up changed data blocks after the initial backup is created, and sends data directly to the cloud in a single step.

## **More cost effective**

Other traditional backup solutions are required to perform full backups of the entire volume periodically in order to support long-term retention policies. Cloud Backup Service is 'incremental forever', which means that no interim full backups are required. This translates to a lower backup storage footprint and reduced network requirements.





# Don't Let NDMP Backups Slow You Down. Try Cloud Backup.

**NetApp understands ONTAP better than anyone else, which is why the best backup solution for ONTAP systems is NetApp Cloud Backup.** Designed by NetApp specifically for ONTAP, Cloud Backup automatically creates block-level incremental forever backups. These copies are stored in object format and preserve all ONTAP's storage efficiencies. Your backups are 100X faster to create, easy to restore, and much more reliable than with any other solution. Storage admins like Anne finally have a better way to back up.

Cloud Backup simplifies the entire back up process. It's intuitive, quick to deploy, and managed from the same console as the rest of the NetApp cloud ecosystem. **Whether you're looking for a less expensive way to store your backups, a faster, more capable technology than NDMP, or an easy way to enable a 3-2-1 strategy, Cloud Backup offers the best backup solution for ONTAP.**

## SEE YOUR SAVINGS

Try the Cloud Backup  
TCO calculator here →

## WATCH A DEMO

Check out Cloud Backup  
in action →

## START A FREE TRIAL

Sign up here to get started  
with Cloud Backup today →

Refer to the Interoperability Matrix Tool (IMT) on the NetApp Support site to validate that the exact product and feature versions described in this document are supported for your specific environment. The NetApp IMT defines the product components and versions that can be used to construct configurations that are supported by NetApp. Specific results depend on each customer's installation in accordance with published specifications.

**Copyright Information**

Copyright © 1994–2023 NetApp, Inc. All rights reserved. Printed in the U.S. No part of this document covered by copyright may be reproduced in any form or by any means—graphic, electronic, or mechanical, including photocopying, recording, taping, or storage in an electronic retrieval system—without prior written permission of the copyright owner.

**Software derived from copyrighted NetApp material is subject to the following license and disclaimer:**

THIS SOFTWARE IS PROVIDED BY NETAPP “AS IS” AND WITHOUT ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE, WHICH ARE HEREBY DISCLAIMED. IN NO EVENT SHALL NETAPP BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

NetApp reserves the right to change any products described herein at any time, and without notice. NetApp assumes no responsibility or liability arising from the use of products described herein, except as expressly agreed to in writing by NetApp. The use or purchase of this product does not convey a license under any patent rights, trademark rights, or any other intellectual property rights of NetApp.

The product described in this manual may be protected by one or more U.S. patents, foreign patents, or pending applications.

RESTRICTED RIGHTS LEGEND: Use, duplication, or disclosure by the government is subject to restrictions as set forth in subparagraph (c)(1)(ii) of the Rights in Technical Data and Computer Software clause at DFARS 252.277-7103 (October 1988) and FAR 52-227-19 (June 1987).

**Trademark Information**

NETAPP, the NETAPP logo, and the marks listed at <http://www.netapp.com/TM> are trademarks of NetApp, Inc. Other company and product names may be trademarks of their respective owners.

NA-000-0423

