

Where's Your Third Copy?

Protecting Unstructured Data

Introduction

As more organizations put their critical data on NAS systems, a complete, well thought out data protection strategy is critical. Data is constantly at risk from corruption as well as accidental, erroneous, or malicious updates and deletions.

In a recent survey, 65% of companies reported lost productivity due to data loss or outages – suggesting that current protection efforts are inadequate.



Traditional data protection strategies, though essential, are insufficient in completely protecting an organization's NAS data.

Companies that are looking for more than 'hope-as-a-strategy' are implementing third (or even fourth) data copy strategies to give them the best chance of recovering from data loss and corruption.

Your NAS Protection Strategy

You can think of NAS data protection strategies as layers of protection. Each layer provides a copy of data which can be used to restore availability should the worst happen. And each layer also provides for different Recovery Point Objectives and Recovery Time Objectives as well as different recovery scopes ranging from individual bits of data, or files, to complete data sets.

These data protection strategies typically consist of local copies (RAID, Snapshots, Local Backups) and off-site copies (replicated copies and synchronized copies).

Your NAS Protection Strategy

Beyond the production copy, a typical strategy is to regularly copy data to a secondary system in an offsite facility. From this system, data can be restored or the users and applications can be redirected to that secondary system if the production system becomes unavailable or its data unusable.

Data outage costs run at over \$300K/hr

But for most organizations, loss of data from their primary and Disaster Recovery (DR) site means a complete loss. Even if they can restore from an off-site tape, it could be days or weeks later, creating a substantial business impact given that average data outage costs run at over \$300K/hr.

What's Special About NAS?

While replicating block data is relatively straightforward – simply copy raw blocks from one system to another with no need for an understanding of the actual data or how it is used, structured or controlled – the replication of NAS data is far more complicated.



With NAS storage, beyond just having a copy of the actual data, the remote site must also contain a regular mirror of the file system structure, the shares and exports required for end user and application access, and all permissions (NFS and/or SMB) that control file access.

Without all this information applied, it would be nearly impossible to have a secondary or tertiary NAS system take over the responsibility of serving data to end users and applications when the source data becomes unusable.

The Limits of Backup and In-Family Replication



Although many organizations protect their NAS data with local copies and implement in-family off-site replication, data still remains at risk due to the inherent limitations of data protection strategies.

The Limits of Backup and In-Family Replication

Local Copies

Typically local copies of data used in a data protection scheme fall into two categories:

Snapshots offer near instant recovery but are very expensive as they consume local disk storage space for each copy created, so relatively few copies are kept meaning limited options in terms of historical copies.

Backups take far longer to recover and present an extended outage event as data must be restored to the NAS platform from a different storage media such as another disk system or tape.

Both of these solutions create point-in-time copies which is ideal, but recovery depends on a valid copy being available which may not be the case as both snapshot and backup copies are expensive to maintain and will be aged out and discarded over time. In addition, both solutions are vulnerable to a single site loss or outage.

The Limits of Backup and In-Family Replication

Off-Site Copies

In-family replication solutions provide the ability to instantly failover to a replica in case the production data system is unavailable or its data unusable – this solves the issue of a loss of the primary production site.

But, the very nature of this replication process, while solving the site loss issue, introduces a new risk: the remote copies are typically not point-in-time copies but ones that are constantly updated.

The very nature of this replication process introduces a new risk

This means that if corruption or an accidental or malicious update or deletion takes place on the source system, it is quite likely going to immediately propagate to the off-site target, making the copy of no value in a recovery effort.

Where's Your Third Copy?

Every organization should have a NAS protection strategy that includes not only local and in-family replication options but also a third copy of data. There are two distinct replication strategies to be considered, and potentially both implemented, depending on your requirements:

Replicating to a NAS platform at the tertiary site which exactly mirrors the production system. This allows the replicated environment to be used to 'fill in' for the production system should the need arise. Ideally the tertiary site would exist in a bunker or behind an air-gap (a break in the network that only allows full access during a replication window).

Syncing a copy of data is a point-in-time Golden Copy of the data that may never be updated or is updated very infrequently. This allows for data to be restored at a point time before corruption or unwanted updates took place. Syncing can use another NAS or the Cloud as a target.



The Datadobi Solution

Datadobi has created two unique solutions to address heterogenous off-site NAS data copy needs. In addition to many other benefits, both allow for controlled granular copy and restore, giving companies full control over what data they replicate.
















DobiReplicate

Replicating all data as well as file system structure, shares and exports required for end user and application access, and all permissions that control file access to another NAS typically behind an air-gapped environment in a bunker site. This enables users to be 'failed-over' to their third copy environment.

DobiSync

A Golden Copy of data on another NAS or in the Cloud. This keeps a point-in-time version of the data so a restore can be done at a time before corruption can take place.

The Datadobi Solution

Product Benefits	 DobiSync	 DobiReplicate
Purpose-built to successfully carry out the data copy tasks with the greatest flexibility, including heterogeneous targets and scheduled replication		
Ability to plan and design the data copy process within the same solution that will execute it		
All aspects of data copy execution and recovery is managed from a single pane of glass		
Provide real-time monitoring and easily consumed, on-demand and automated reports		
Can perform discovery and analysis to provide a clear picture of the source and target systems' data		
Natively store and retrieve/restore data from a cloud repository		
Execute a failover to a target system automatically with full management and outcome tracking of the process		
Execute a failback to a target system automatically with full management and outcome tracking of the process		

[Compare All Product Benefits](#)

Conclusion

When considering implementing the use of off-site data copies as part of your overall data protection and availability strategies, it is crucial to understand that while a secondary copy of data eliminates some risk, the need for at least a tertiary copy is vital to protect an organizations critical data and to maintain business continuity.

Implementing a tertiary copy should be simple, fast, cost-effective, and above all reliable. However, due to the disparity between NAS platforms and the complexity of keeping that data synchronized, without proper tools and methodology, you will likely encounter numerous challenges.

Datadobi has created purpose-built solutions for the needs of heterogeneous data replication and syncing.

Protect Your Unstructured Data Now

Contact Datadobi today to learn more about how we can help with your company's data copy and off-site protection efforts: info@datadobi.com