



## Backups

# A Case for Decentralized Cloud Storage

### A Better Way to Backup Data

The way organizations back up their data has evolved significantly over the years. As early as the 1960's magnetic tape became the prevalent backup media and remained so for several decades. Hard disk drives (HDDs) started to become a more popular backup option in the 1990's, but it wasn't until the 2000's when storage densities started to rapidly increase alongside significant drops in price that HDDs started to become the norm. Then as accessibility and the ubiquitous nature of cloud storage offered a secure, resilient, simple, and cost-effective way to back up data, organizations quickly began to shift to backing up to the cloud.

Speed, resiliency, security, cost, and simplicity have always been key factors in driving the shift and adoption from one backup technology to another. The same is true today. While centralized cloud storage has been recently leading the way in the destination of organizations' backups, the dramatic advantages that decentralized storage offers improvements in resiliency, immutability and security - with less cost and complexity - has begun to trigger a new shift in data backup evolution.

# Backup Cloud Storage Challenges

Centralized cloud storage lacks the desired resiliency due to unreliable redundancy and vulnerability to significant scale outages that can impact entire regions.

The insufficient granularity of access controls and single point of failures inherent to centralized cloud storage can make stored backups more susceptible to risk.

Major hyperscale centralized cloud providers charge excessive storage, egress, and operational fees for large backup files, while budget cloud vendors tack on hidden fees and penalize frequent egress.

Budget backup vendors sacrifice speed and reliability in order to offer low cost services

The complexity of setting up backups on multiple availability zones and regions in centralized hyperscale cloud storage environments significantly drive up engineering and administrative time, effort, and costs.

## Solution

The unique and massive parallelism of Storj Decentralized Cloud Storage (DCS) that delivers unmatched security, immutability and performance - with 99.5% multi-region availability, but 80% of the storage and egress cost of today's hyperscalers.

## Outcomes

Delivers unmatched resiliency and parallelism to help achieve desired Recovery Time Objectives (RTOs) and Recovery Point Objectives (RPOs) through built-in global redundancy combined with high-performance multi-threaded file transfers.

Lowers storage and egress costs by at least 80% compared to hyperscalers and eliminates cost complexity and lock-in with transparent, predictable flat-free pricing for storage and egress with no added fees for multi-region redundancy or surprise fees for frequent backup egress.

Reduces the complexity associated with multi-region redundancy with built-in global distribution and redundancy, automated data orchestration, S3 compatibility, simple upload and download processes, and more.

## Improving Backup Resiliency

Backing up your data is about more than just protecting and preserving it against unforeseen problems. It's about helping you ensure that your business has the resiliency to continue running even in the face of equipment failures, outages, or any kind of incident that could disrupt your operations. While there are a variety of factors that contribute to the resiliency of a backup solution, two of the most fundamental aspects are availability and reliability.

Data redundancy plays a major role in addressing the needs of availability and reliability. In fact, major hyperscalers tout their ability to deliver high redundancy by providing several geographically distinct availability zones within individual regions. However, the multiple [December 2021 outages in AWS' East region](#) demonstrate the ineffectiveness of redundancy within such centralized infrastructures, given that those so-called redundant availability zones in the East region weren't accessible for hours during some of those outages.

The single-point-of-failure nature of centralized infrastructures makes them inherently vulnerable to such large-scale outages. You can attempt to protect against such outages by further increasing redundancy by backing up your data to multiple regions, but that doubles your storage, security, and administration costs while increasing your engineering complexity.

"The latest AWS outage is a prime example of the danger of centralized network infrastructure."



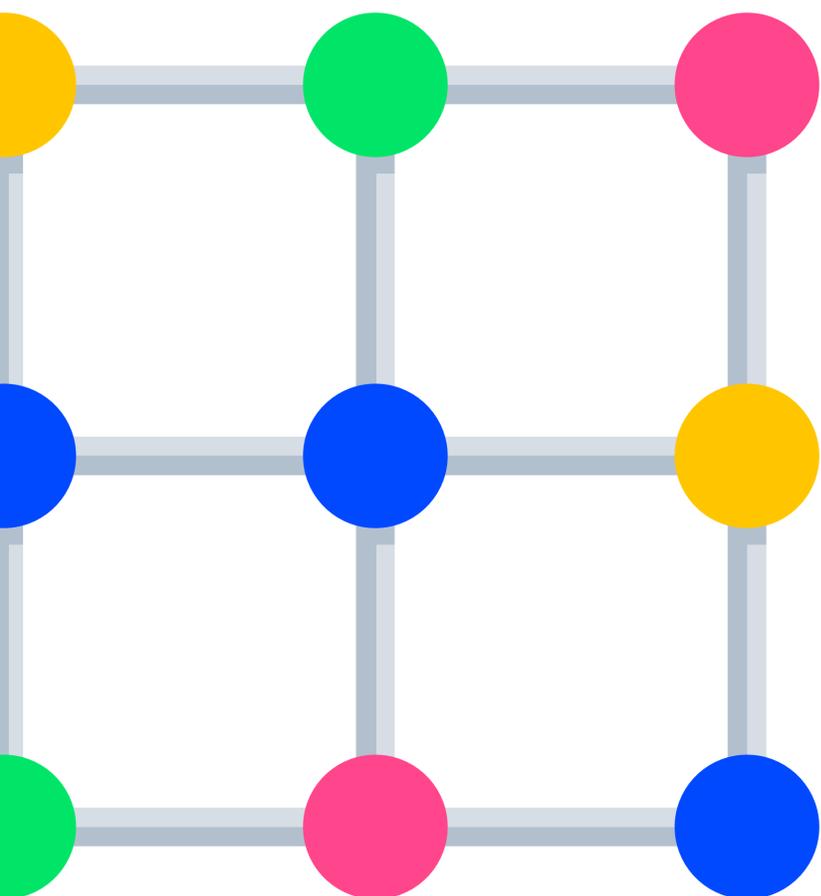
**Sean O'Brien**

Visiting Lecturer in Cybersecurity at Yale Law School

The problems with availability and reliability typically increase even more if you turn to storage-only providers that focus primarily on offering a low-cost alternative. These vendors offer even less redundancy than larger vendors. Plus, you still have to pay extra for what little redundancy they do offer.

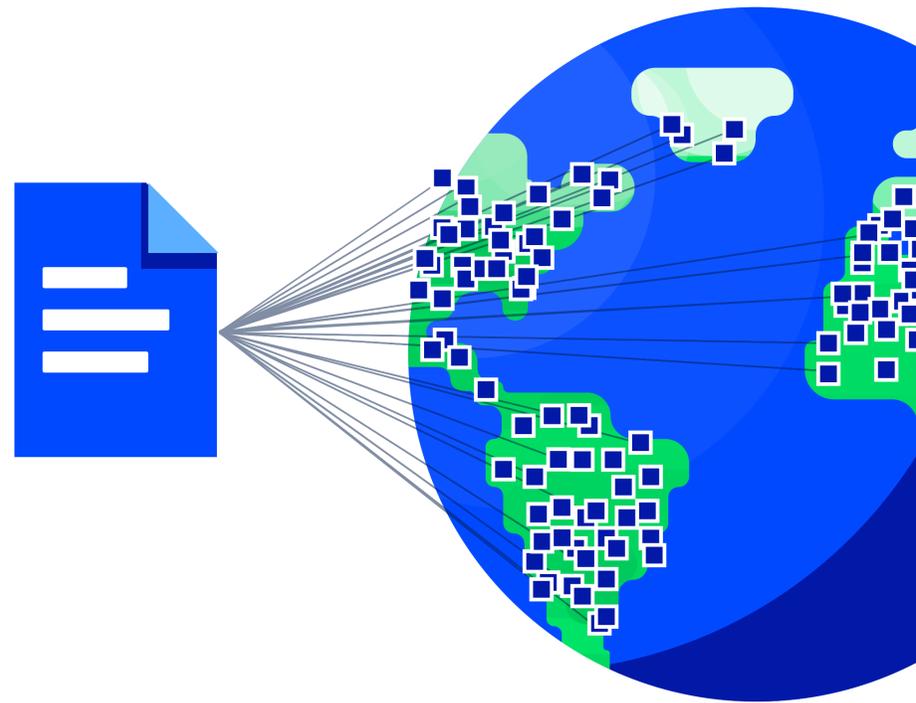
But perhaps of greater concern is the reliability or quality of storage that these low-cost vendors offer. In order to compete on price, they often buy large quantities of inexpensive, very high-capacity hard drives that have slower performance. This lets them optimize their data center rack space while lowering their electricity bills. But the slower performance of those inexpensive, high capacity drives means that if you ever have to recover data or applications in the event of a disaster, you might have trouble achieving your Recovery Time Objectives (RTOs).

The bottom line is that you need to be careful with budget backup vendors. A disaster could push the ability of their infrastructures to their limits, denying you the resiliency you want when you need it most.



The highly distributed nature and performant design of Storj DCS offers the resiliency and throughput you need to achieve the RTO your backup strategy requires. For less than half the price of a single availability zone from a centralized hyperscaler, Storj DCS provides highly redundant, multi-region distributed cloud storage with ultra-high availability for your backups. When you back up to Storj DCS, it splits your backup object files into 80 or more erasure-coded pieces that are stored across 13,500 geographically diverse nodes and ISPs in more than 90 countries across the globe.

The built-in redundancy and global distribution of Storj DCS helps ensure your backups are always available. Additionally, to deliver super-fast throughput when uploading and downloading your backups, Storj DCS implements multi-layered parallelism. For example, when you need to access a backup file, Storj DCS only requires 29 of those 80-plus erasure-coded pieces to reconstruct the file. However, it actually requests 39 pieces to be downloaded simultaneously, using only the fastest 29 pieces that reach you first.



In addition to base parallelism provided by Storj DCS, you can also achieve even faster download speeds by taking advantage of segment parallelism. As the name suggests, segment parallelism transfers a file's segments in parallel. As an example, a 512MB file is made up of eight 64MB Segments ( $512/64=8$ ). So, when you download that 512MB file, not only will Storj DCS use the fastest 29 pieces in parallel, it will download 232 pieces of the file in parallel (8 segments X 29 pieces = 232 total pieces). This should enable file downloads to complete eight times faster than downloads where segments transfer serially. Plus, unlike some centralized cloud providers that make you pay extra to boost your download speeds, Storj DCS doesn't charge more to give you faster downloads.

“Storj really has delivered an industry first—enabling multi-GB speed and providing us with unmatched performance, increased parallelism, redundancy, and resiliency—and all cost-effective.”



**Dr. Antonin Portelli**

University of Edinburgh Research Explorer

## Ensuring Immutability and Security of Backup Files

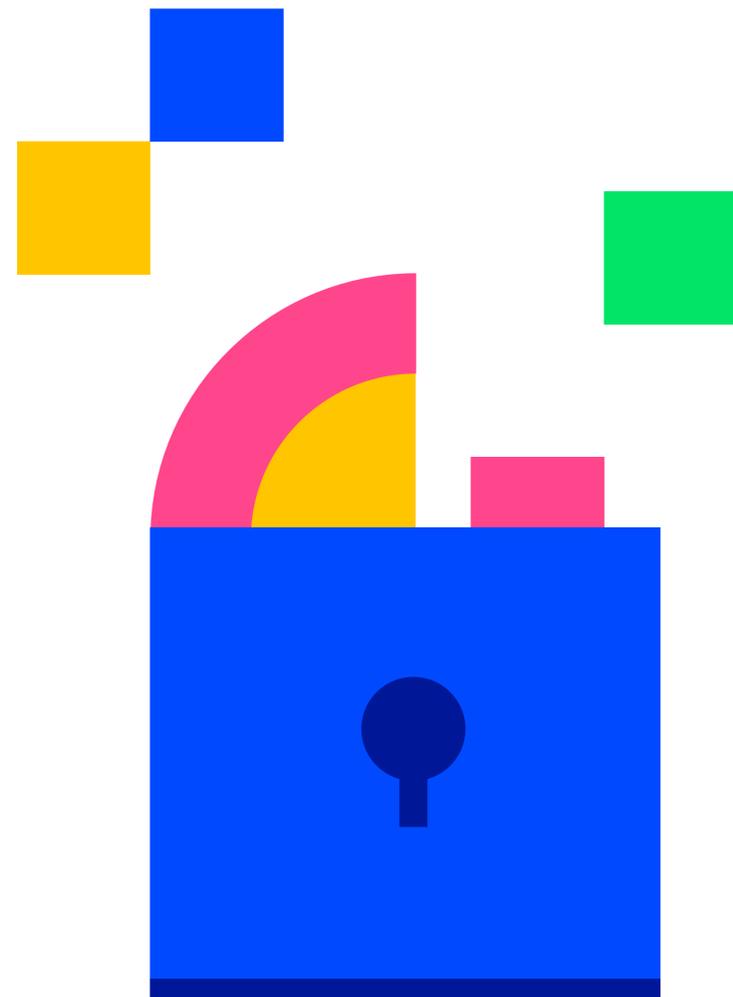
When you back up your data, you want to know that it remains secure and unchanged. That includes protecting it against data loss, data corruption, and accidental or intentional deletions or edits. The key to ensuring the security and immutability of your backups is having a provider with a zero-trust architecture. However, hyperscale centralized cloud infrastructures are not zero trust by default. Not only do they rely on single point of failure architectures, but mistakes or problems with permission grants can occur too easily with their identity and access management (IAM) frameworks.

Centralized cloud storage vendors assign access controls at a bucket or folder level. So, if you accidentally misconfigure those folders or you somehow end up with a leaky bucket, all the backup files inside are put at risk. Organizations that want to ensure zero trust with a centralized cloud vendor typically have to hire a consultant to do continual audits to ensure compliance and make sure nobody has access to anything they shouldn't have access to.

To help mitigate these risks and provide greater file immutability, Storj DCS provides granular access control at the object file level that can be programmatically automated. That means when you set up a backup to be stored on Storj DCS, you can designate a specific individual or individuals to be the only ones to have access rights to that backup. By giving you granularly restrictive grants, you're able to more easily ensure only the right people have access to individual backups - mitigating intentional or accidental deletions.

And to further fortify your access controls, Storj DCS embeds a macaroon inside each object file's URL as part of our native link sharing mechanism. And since we separate the encryption function from our macaroon-based access management capabilities, you can manage both 100% from the client-side.

Overall, Storj DCS delivers an edge-based defense-in-depth strategy combining a zero-trust approach, zero-knowledge infrastructure, end-to-end encryption, and credential level access control. It's a multilayered approach that provides indisputably stronger security than the standard public cloud approach. Plus, Storj DCS continuously performs cryptographic audits of the storage nodes and data associated with your backup files that ensure an enterprise-grade SLA of 11 9s of durability and 99.95% availability.



## Lowering Backup Cost and Complexity

One of the main drivers that pushed many organizations to turn to the cloud for backups was the desire to eliminate the capital expenses (CapEX) associated with investing in large racks of local hard drives, as well as the cost to maintain those hardware infrastructures. Even though the move to the cloud eliminated that hardware CapEX, organizations still pay a significant amount in operating expenses (OpEx) with major hyperscale centralized cloud vendors. Those high expenses come in the form of significant storage fees, egress fees, and operational fees. Plus, storage and egress costs at least doubles if you want the redundancy that comes from storing your backups in more than one region.

Due to the high costs associated with mainstream centralized cloud vendors, some organizations turn to budget backup cloud vendors. Not only can this lead to the resiliency issues previously discussed, but such solutions aren't as cost-effective as they appear on the surface. For example, budget backup cloud vendors often tack on large added fees if you egress more data than you upload. To make sure they can stay profitable, some vendors charge three months' worth of storage for any uploaded data (including overwrites) regardless if you only keep the data for a month, week, or shorter.

When you back up with Storj DCS, you don't get hit with hidden fees. Storj DCS gives you predictable pricing with flat fees for capacity and bandwidth. Plus, Storj doesn't lock you in with complex cost structures or exorbitant egress fees. Storj takes an innovatively different approach to cloud object storage for backups, providing you dramatic savings compared to what centralized cloud solutions offer. Due to the globally distributed network of independent node operators, it doesn't have the cost of building and operating its own centralized data centers.

As a result, Storj can pass significant savings onto its customers. For example, for less than half the price of a single availability zone from a centralized hyperscale cloud storage vendor, Storj offers organizations multi-region distributed cloud storage for their backups with ultra-high availability through more than 13,500 geographically diverse points of presence spread across the world. So, instead of having to multiply your backup costs by the number of each geographic region where you want backup redundancy, Storj DCS automatically gives you global redundancy without any additional costs by distributing your files among its storage nodes in multiple countries on multiple continents. In addition, Storj DCS offers geofencing so you can choose the specific region or regions your data will reside in - like the United States or the European Union - helping you meet compliance requirements.

Even more significant are the savings that Storj DCS provides on egress fees. For example, centralized hyperscale cloud providers typically charge about \$90 or more per terabyte (TB) of egress in their North America distribution areas. Those prices go up even higher outside of North America, and in some cases to a significant degree. But Storj DCS charges a flat fee of \$7 per TB of egress regardless of geographic location. Additionally, Storj doesn't charge you extra for frequent backups or deletions of your backups. That saves you money, while giving you the flexibility to create backup policies that meet your needs, not the dictates of your provider.

“Storj saved the day for us. We could have used S3, but that would have cost \$20k plus. Also, using the Storj DCS uplink direct improved our performance by at least 10x.”

**Baris Aksu**

Founder & CEO at Cod3r

## Guidance for Achieving the Best Backup Results

To ensure the greatest backup success, we always recommend that you take advantage of a proven backup utility to facilitate your backup processes and the management of your actual backups. This will save you time and facilitate your ability to implement your desired backup strategies, such as retention policies. A good backup utility can also deduplicate and compress your files. A backup utility that uses snapshots or can pack and compress your files into archives will significantly lower your cost of ownership. Additionally, since Storj DCS is optimized for large object files of 64MBs or greater, such a utility will help you ensure greater overall backup success. Check out some of our backup partners for a perfectly integrated backup solution - [Duplicati](#), [Rclone](#) and [Restic](#).

# Optimizing Disaster Recovery

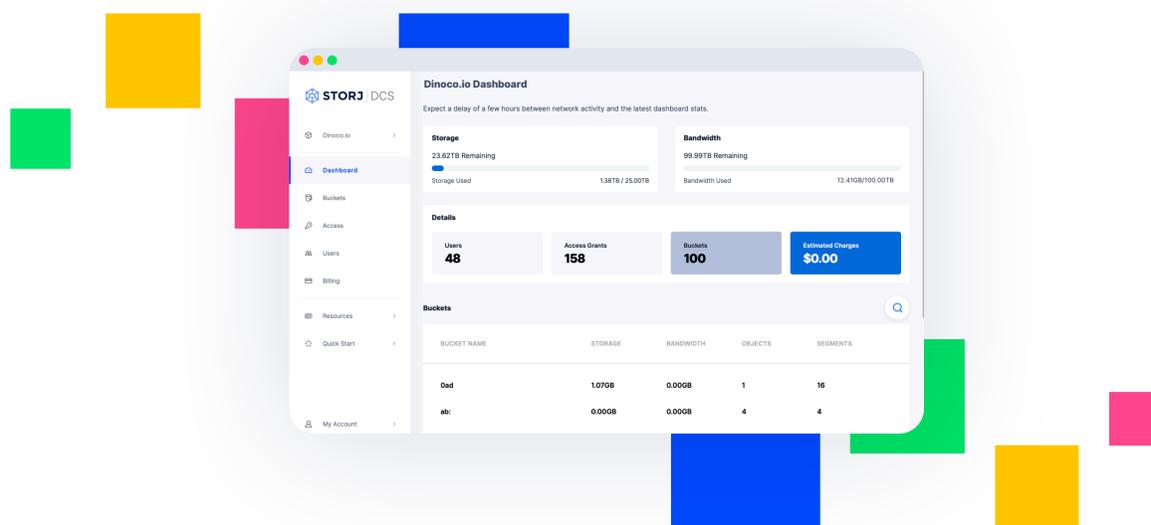
The innovative structure and design of the Storj DCS globally distributed network and massive high-speed parallelism also makes it the ideal foundation for any organization's disaster recovery strategy. It offers the fastest option to recover your data regardless of location. This is especially true for dynamic organizations that might want to move their data centers between regions due to geo-political events, disasters, or other geographic concerns.

Storj DCS delivers the resiliency, performance, immutability, and value you want in a backup and disaster recovery solution. It delivers the performance of hot backups with multi-region high availability at cold backup costs. The massive parallelism inherent in its decentralized and distributed network gives you high performance backup transfers to make sure you can achieve fast RTOs.

For more information on how to take advantage of the innovation of decentralization in Storj DCS to more easily, efficiently, and cost-effectively improve data backup speed, durability and immutability, visit [www.storj.io](http://www.storj.io).

## Experience Storj DCS today.

Visit [www.storj.io](http://www.storj.io) to get started and see for yourself how decentralized storage is more secure, private, and cost-effective than traditional cloud storage.



**Start building on the decentralized cloud.**

[www.storj.io](http://www.storj.io)

