STORJ

PERFORMANCE ANALYSIS REPORT

An AI and Human Analysis on the IDS Benchmarking Report of Object Storage Performance.

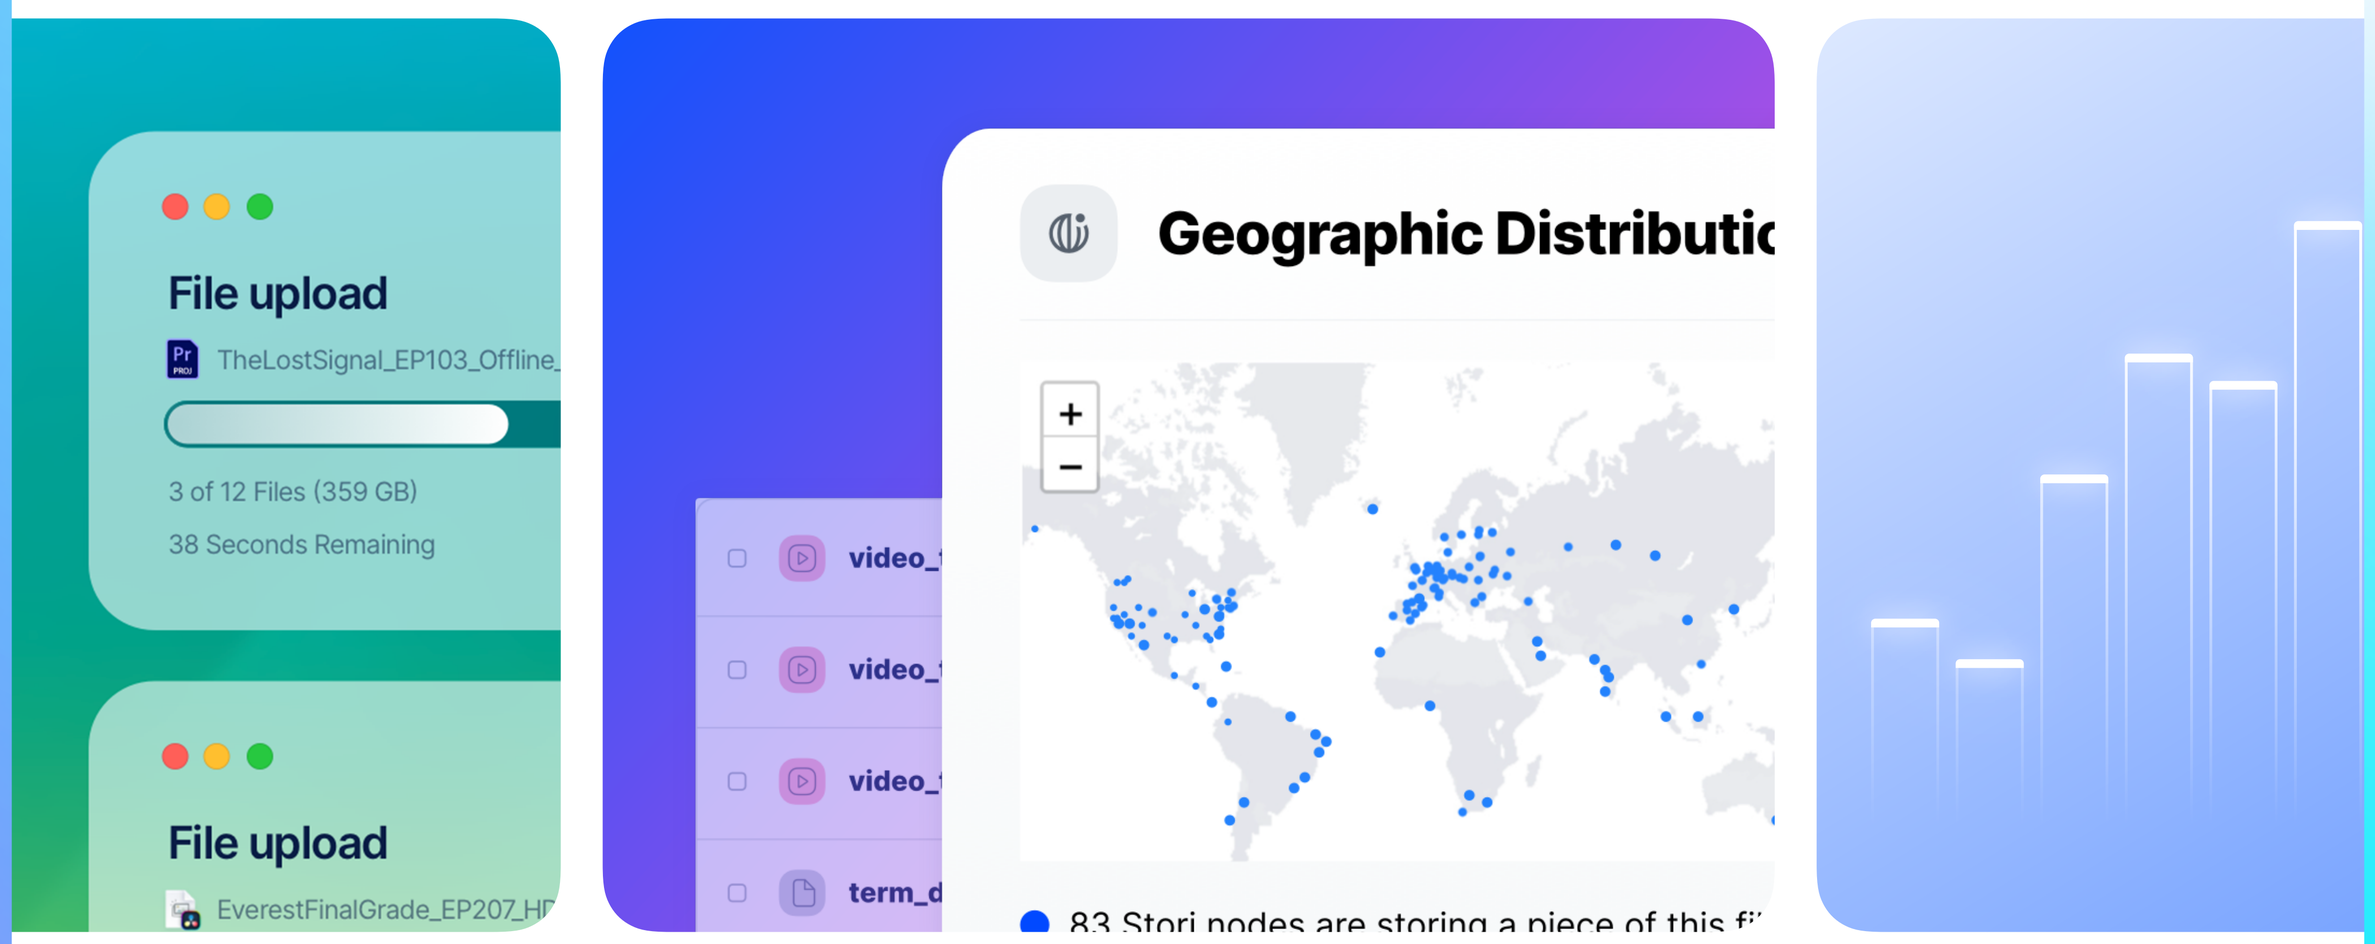

Introduction

Storj has consistently improved its global upload and download performance on its distributed object storage over the past year. So, a few months ago, Storj asked Integrated Digital Systems (IDS) to benchmark Storj distributed object storage performance against AWS S3, Wasabi, and Backblaze using real world conditions. Readers are welcome to delve into the **original report** for detailed technical review.

Test Objective from the IDS report:

This test aims to benchmark the performance of Storj Object Storage, AWS S3, Backblaze B2, and Wasabi Hot Storage. The primary goal is to quantify the performance differences among these services across various regions and parameters, including node location, bucket location, file size, and time of day. Additionally, the test seeks to determine how cloud storage performance can be expected to vary based on its location and access point for a production environment with globally distributed resources.

The benchmarking utilized compute nodes in Los Angeles, London, New York, and Sydney that were configured to run an hourly cron job to test the performance of every S3 compliant bucket, for each provider, in each region, with various file sizes. The results were extremely comprehensive with an incredibly large spreadsheet of data points to consider.

IDS wrote conclusions and analysis showing the following key findings on performance variability.

Storj

Storj maintains consistent performance across different time periods with less volatility than other providers.

Wasabi

Wasabi demonstrates the most pronounced time-based performance fluctuations, with notable degradation during certain hours, likely correlating with increased backup and batch processing workloads from other customers.

Backblaze

Backblaze shows performance instability with extreme outliers that can dramatically skew average performance for specific time periods, making it less predictable for consistent workloads.

AWS

AWS demonstrates relatively stable performance throughout the day with minimal time-based variations.

Storj key takeaways from performance benchmarking and AI analysis.

First off, IDS did an incredible job benchmarking object storage performance in the way real media productions operate. The rigor and extensiveness of the testing were impressive.

The clear takeaway from this report is that while Storj has claimed it is “globally fast”, perhaps the more important differentiator for businesses is the performance consistency that stands out for large file workloads.

If you are evaluating options for cloud storage, here is where Storj is the clear choice.

When your data needs to be accessible to globally distributed teams.

Proof point from report: Storj provides a predictable and usable performance baseline for all production teams, regardless of their geographic location relative to the data. Remote production teams require consistent, predictable access to their content, regardless of location for seamless collaboration

When the files you need to share are large.

Proof point from report: Storj exhibits a unique and counter-intuitive architectural advantage for media ingest, where upload (PUT) performance demonstrably increases with file size. This accelerated ingest capability stands in stark contrast to the significant performance degradation observed across all tested competitors—AWS, Backblaze, and Wasabi—under the same conditions.

When consistently stable performance is critical.

Proof point from report: Storj delivers consistent and stable performance over time, mitigating the business risk and production downtime associated with the volatile, outlier-prone performance profiles of competitors like Wasabi and Backblaze. This reliability is a critical, enterprise-grade feature for mission-critical post, news, and sports demanding workflows.

AI Analysis: The Object Storage Competitive Performance Landscape: A Granular Analysis

This result is great for Storj, but, with such a comprehensive chunk of data, it seemed prudent to ask AI what else it could find.

Of course, AI outputs can’t be fully trusted yet, so Storj experts combed through the analysis. The outcome was consensus that AI had some good points, despite being overly dramatic at times. So, this report lets AI have its say, while humans reviewed and commented on the AI interpretation of the report. The goal of this report is to provide a balance in bias to thoughtfully determine which object storage solution is right for your workloads and workflows. Here’s what AI found and Storj’s human response.

A. The Predictability & Consistency Index: Quantifying Operational Risk

The time-series tests, conducted over a full week, provide a clear view into the performance consistency of each provider. The data reveals that not all throughput is created equal; the
variability of performance is a critical factor that translates directly into operational and business risk.
The time-series charts for Wasabi and Backblaze expose significant performance volatility. For example, the chart for Wasabi's London bucket shows GET throughput fluctuating wildly, while the Backblaze London chart displays a massive, anomalous spike in GET performance around 13:00 local time. The report explicitly notes that Wasabi demonstrates "pronounced time-based performance fluctuations" and Backblaze is "prone to extreme outliers". In stark contrast, the performance graphs for Storj and AWS are visibly flatter, with the report concluding that Storj maintains "consistent performance... with less volatility" and AWS is "relatively stable".

This performance volatility is more than a technical metric; it represents a tangible business risk. For an organization running a critical, time-sensitive process—such as a nightly database backup to meet a recovery point objective (RPO), a media file transfer under a strict broadcast deadline, or a batch processing job that gates subsequent workflows—relying on "average" performance is untenable. These processes must be designed for worst-case or at least predictable scenarios. The volatility shown by providers like Wasabi and Backblaze introduces profound uncertainty. A critical operation might fail to meet its Service Level Agreement (SLA) or miss a crucial window not because the provider is slow on average, but because it was unacceptably slow at the specific time the job ran. This uncertainty complicates operational planning and forces engineers to build in larger, more expensive buffers, or else accept a higher risk of failure.

Table 1: Performance Predictability Scorecard

Provider	Predictability score (1=best)	Key characteristic	Supporting evidence	Best-fit workload
aws **AWS**	1	Most predictable	"Relatively stable performance" noted in report; flat time-series charts.	Mission-critical, region-locked apps and workflows.
Storj	2	Consistent baseline	"Maintains consistent performance with less volatility" noted in report.	Mission-critical, globally distributed apps and workflows.
Wasabi	3	High volatility	"Pronounced time-based performance fluctuations"; spiky throughput in charts.	Non-critical, price sensitive.
Backblaze	4	Prone to outliers	"Prone to extreme outliers that can dramatically skew average performance".	Best-effort, archival.

Storj response on section A.

Storj landed at number two behind AWS, but it is a strong showing for a newer entrant in the cloud storage space. The volatility noted in Wasabi and Backblaze's performance isn't surprising. Their pricing has long attracted attention, but low cost doesn't guarantee reliability when performance matters. What's more interesting is how AI workloads are evolving. Inference, semantic search, and media-driven model training depend on fast, predictable access to large datasets. Storj's architecture is built for high-throughput retrieval, and it is an excellent foundation for these AI media use cases. AI sitting on top of reliable, globally performant storage, teams can move faster and focus on their creative storytelling.

B. The Tyranny of Distance: Deconstructing the "Latency Tax" of Centralized Clouds

The comprehensive test suite, which measured performance from four compute nodes to various global bucket locations, starkly illustrates the performance penalty of geographic distance in traditional, centralized cloud architectures. When a compute node accesses a storage bucket in a distant region, the performance of AWS, Wasabi, and Backblaze collapses.

A clear example is found in the 200MB file tests: a compute node in Sydney accessing an AWS EU bucket achieves a meager 25.3 MiB/s GET throughput with a high latency of 513.6 ms. Under the exact same test conditions, Storj Global delivers 49.3 MiB/s GET throughput—a 95% improvement. While Storj Global is not the absolute leader in every single local test, it consistently provides a significantly higher performance floor for these critical cross-region scenarios, effectively "eliminating catastrophic performance degradation for distributed teams". This pattern of severe degradation for traditional providers is repeated across the entire dataset for all file sizes.

This phenomenon forces a global organization using a traditional provider like AWS into a difficult position. It cannot simply place all its data in a single region like us-east-1 and expect its teams in Sydney or London to remain productive. The standard architectural solution is to implement a complex and costly multi-region strategy. This involves creating and managing distinct storage buckets in multiple geographic regions (e.g., us-east-1, eu-west-2, ap-southeast-2), which immediately introduces significant operational overhead. Engineers must manage separate Identity and Access Management (IAM) policies, configure and monitor data replication services (like AWS S3 Cross-Region Replication), and handle different regional endpoints within their application logic.

This workaround imposes a "Latency Tax" that is paid in both direct financial costs and indirect operational complexity. The direct costs include data egress fees for every byte replicated across regions, as well as the cost of storing duplicate data sets. The indirect costs, which are often larger, stem from the man-hours required to design, implement, and maintain this fragile, distributed infrastructure.

Storj's "Global" bucket architecture inherently solves this problem without forcing the user to become a distributed systems expert. It provides strong global performance from a single bucket, with a single endpoint and a unified policy set. This means Storj is not just offering better cross-region performance; it is offering a fundamentally simpler, more elegant, and more cost-effective architecture for any application with a global user base.

Storj response on section B.

After years of centralized cloud thinking, it’s no surprise that many teams struggle to grasp how Storj delivers one global region of storage at a lower cost than a single AWS S3 region. Storj’s performance advantage isn’t just about geography,it’s rooted in how data is encoded and retrieved. Unlike traditional replication or region-bound architectures, Storj uses erasure coding, specifically Reed-Solomon, to split files into multiple segments with built-in redundancy. This enables the system to reconstruct the original file from a subset of the fastest available segments, dramatically improving upload and download performance, especially for large media files.

While competitors degrade under load or large file sizes, Storj accelerates. Reed-Solomon encoding enables this by tolerating segment loss without compromising integrity, making it ideal for distributed ingest and global access. Yes, there are additional fees and workarounds required to make the legacy centralized storage model work for today’s more global workflows. Storj avoids all of that.

C. Workload-Specific Performance Deep Dive: The Large-File Upload Advantage

The most striking finding from the benchmark report is a unique and powerful performance characteristic demonstrated by Storj: PUT (upload) throughput improves as file sizes increase. This is a fundamental architectural divergence from every other provider tested.

When accessed from the New York compute node, Storj's PUT throughput for 200MB files was 689.1 MiB/s. For 5GB files, this throughput increased to 768.4 MiB/s, an improvement of 11.5%. This scaling advantage was also observed from the Los Angeles node, where PUT throughput for 5GB files was a formidable 473.3 MiB/s.

This behavior stands in sharp contrast to the traditional providers, all of which experienced significant performance degradation with larger files. From the same New York node and testing the same file sizes, AWS US East's throughput dropped by 20.5%, Backblaze US East's dropped by 33.7%, and Wasabi US East's dropped by 23.4%. This is not a minor variance; it signals a core architectural difference. The parallel, decentralized nature of Storj's network appears to become more efficient as file size increases, likely by engaging more storage nodes concurrently for a sustained transfer. Conversely, centralized providers appear to hit ingress bandwidth bottlenecks or processing overhead that throttles performance on larger, longer-running uploads.

The industries most acutely pained by slow large-file uploads include Media & Entertainment (transferring raw 4K/8K video footage), Life Sciences (uploading multi-gigabyte genomic sequencing data), and Big Data/AI (ingesting massive training datasets for model development). For these high-value verticals, upload speed is not a matter of convenience; it is a direct driver of productivity, discovery, and time-to-market. A faster upload means a video editor can begin a project sooner, a VFX studio can hit a deadline, or a researcher can begin a critical analysis earlier. This makes Storj a "killer application" for these specific, demanding workflows.

Table 2: Large-File PUT Performance Scaling Comparison (New York Node)

Provider	200MB PUT Throughput (MiB/s)	5GB PUT Throughput	Performance change (%)
Storj US	689.1	768.4	+11.5
Wasabi US East	721.6	552.9	-23.4
AWS US East	672.2	534.5	-20.5
Backblaze US East	444.6	294.8	-33.7

Data sourced from benchmark report tables for 200MB and 5GB files.

Storj response on section C.

Storj is an extremely high-value solution for the world's most data-intensive and geographically demanding production workflows. From raw 8K ingest to collaborative color grading and VFX rendering, these workflows require fast, reliable access to massive files across distributed teams. Storj's architecture delivers performance at scale without the constraints of traditional region-bound storage. The analysis validates that M&E is a great fit for Storj, which is why Storj has been developing a **solution to specifically solve post-production challenges.**

Next steps for Storj.

The performance benchmarking performed by IDS highlighted the benefits of globally distributed storage for media workflows. That said, Storj isn't stopping in its journey to continue to be incredibly fast and consistent globally. Here are some of the plans Storj has in the works.

Intelligent tiering for workflow optimization.

Storj is implementing tiering logic that responds to how data is actually used in production settings, taking into account ingest, media access frequency, and file sizes. Intelligent tiering assists the production to ensure the proper performance, helping to manage budgets, whether you're in editorial, VFX, or finishing, when you need it.

Event hooks for automating pipelines.

Storj is adding native support for storage-level events, such as file arrival, completion, or access. These events automatically trigger downstream workflows. This unlocks orchestration for M&E and AI teams who rely on fast, accelerated pipelines.

Accelerated optimizations.

Storj is continuing to optimize PUT/GET concurrency and segment selection to push performance to the edge, especially for globally distributed teams. The architecture already scales with file size, and Storj will continue to scale as workflow intensity pushes the limits of cloud storage.

Workflow-native integrations with telemetry

Storj will provide deeper integration with the tools teams rely on, including media asset products, editing platforms, and VFX/color render pipelines. These integrations will reduce production friction and shorten the time from ingest to creative output. Alongside this, Storj will introduce telemetry, which gives teams real-time visibility into performance and system behavior, so they can iterate faster, troubleshoot confidently, and optimize workflows with data-driven insights.

About Storj

Storj delivers a fast unified cloud solution designed for global media, AI and creative workflows. Combining distributed storage, advanced compute, and cutting-edge file access, Storj delivers a unique platform that powers efficiency and innovation to operate on a global scale.

storj.io