



STORAGE DEVELOPER CONFERENCE



BY Developers FOR Developers

AI's Environmental Storage Problem

Pioneering Architectures that Achieve a 90%
Carbon Footprint Reduction

Presented by Damein Morgan
Senior Software Engineer, Storj

Agenda

Storj DCS: The Basics

Costs of AI's Transformative Explosion

Cutting Carbon with Storj DCS

AI Workload Case Studies

Questions



Damein Morgan

Senior Engineer at Storj

https://www.linkedin.com/in/damein_morgan


LOCKHEED MARTIN 

**NORTHROP
GRUMMAN** 

FRESNO STATE 

UCSB 

How It Works



Storj Nodes

Thousands of shared hard drives store pieces of data on the network, without access to any complete file or usable data. Node operators fairly (and profitably) compensated.

Applications

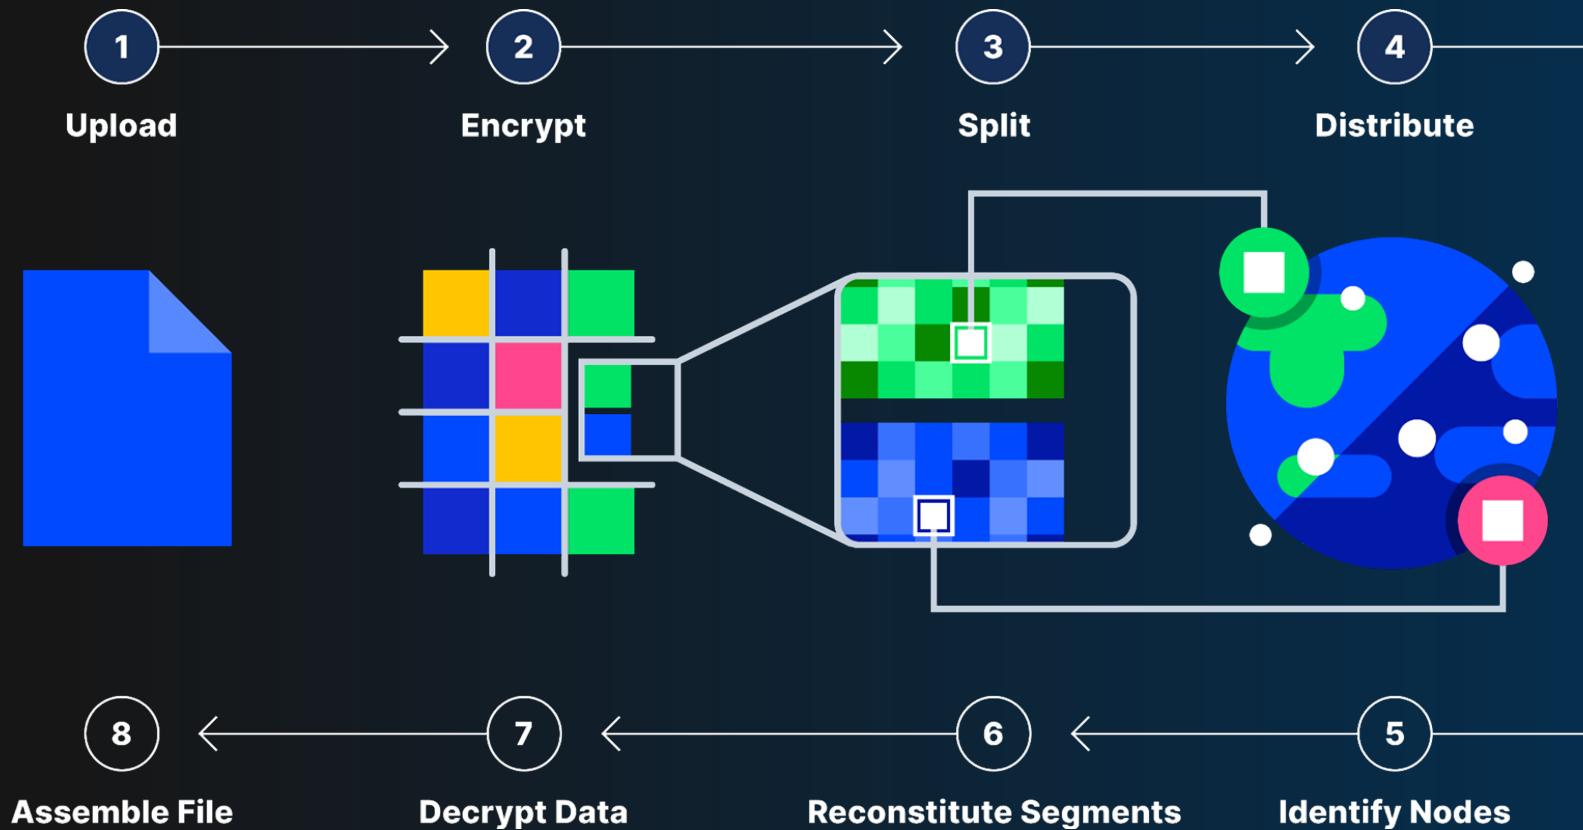
Client applications store encrypted and encoded files split into fragments and stored across the distributed storage network.

Satellites

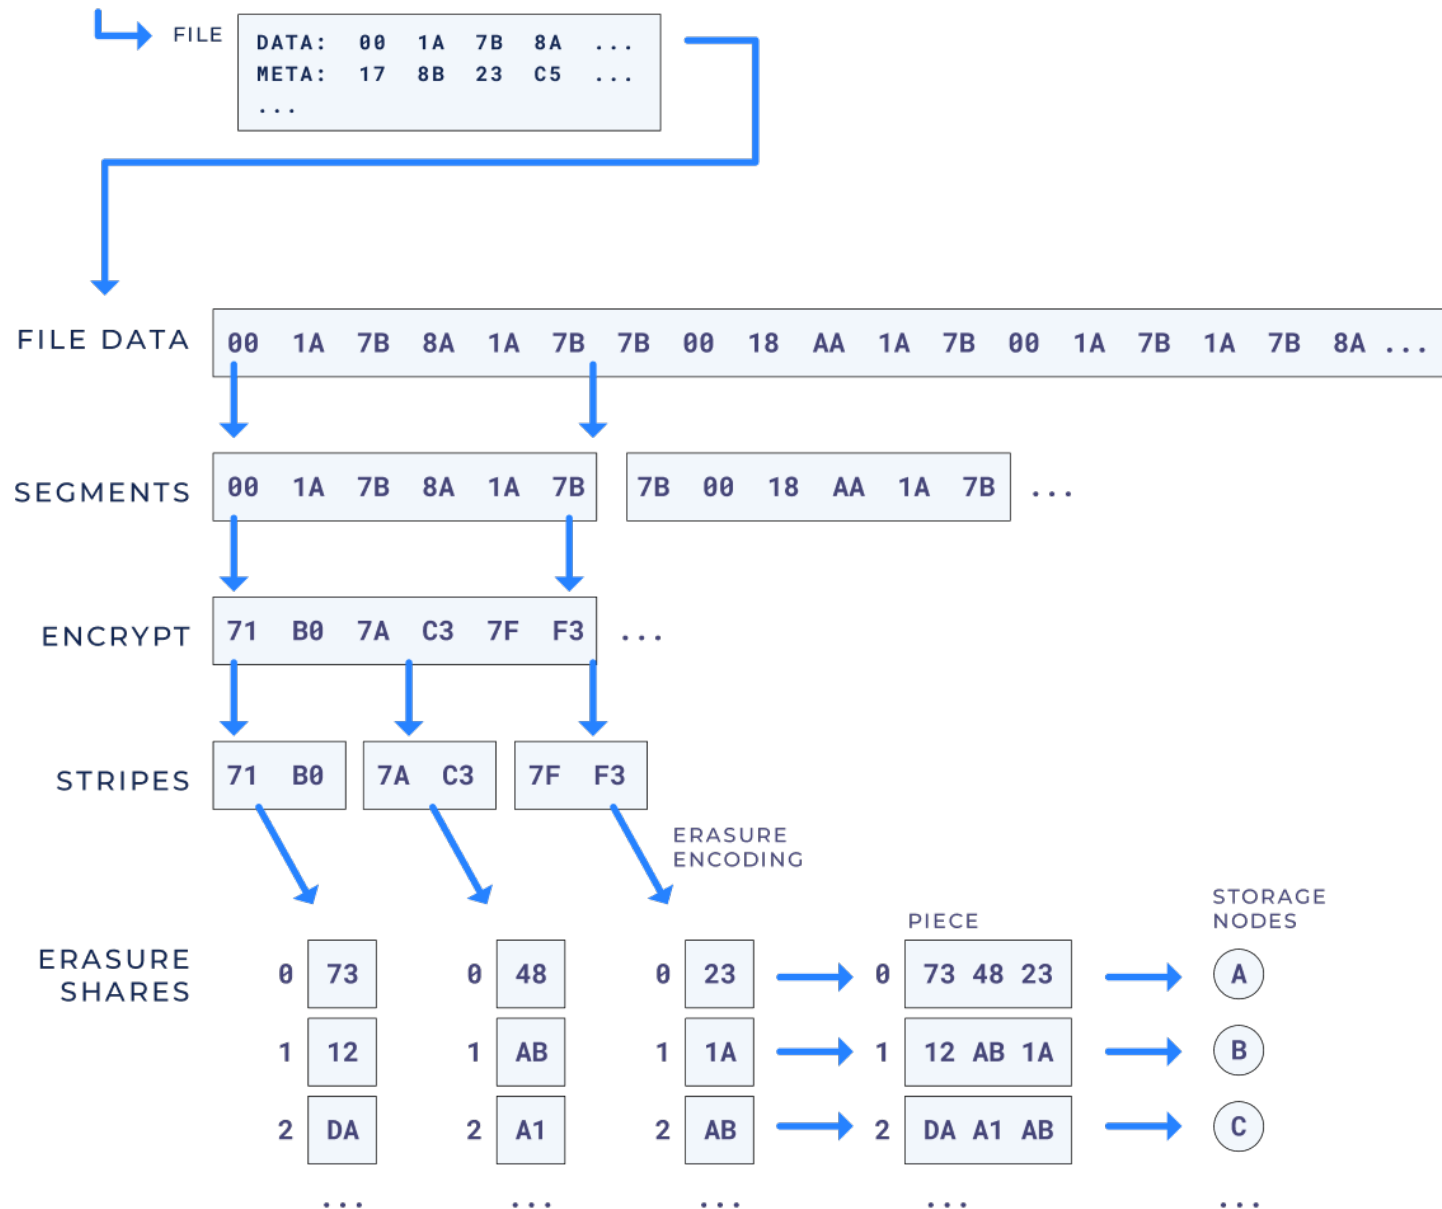
The Storj network enables applications to store data, ensures data reliability, manages access controls, and pays storage nodes.

How It Works

What Happens to Objects?




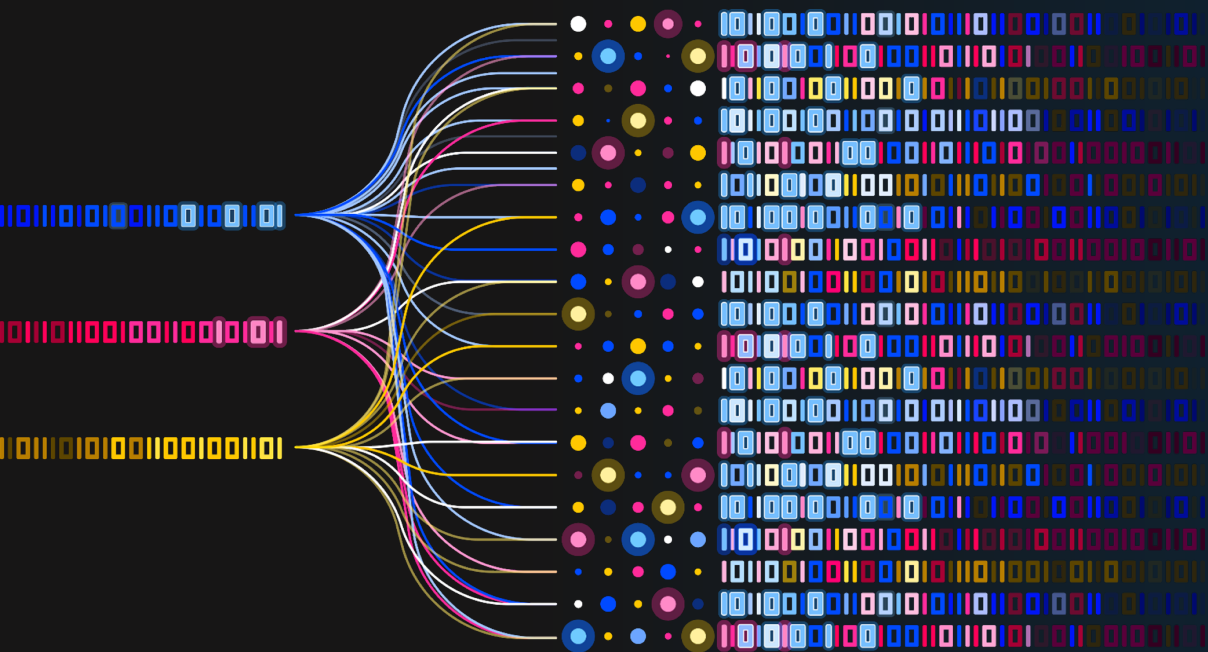
BUCKET



How It Works


What Happens to Objects?





AI's Transformative Explosion

Global Data Generated Annually




Source: Statista.com

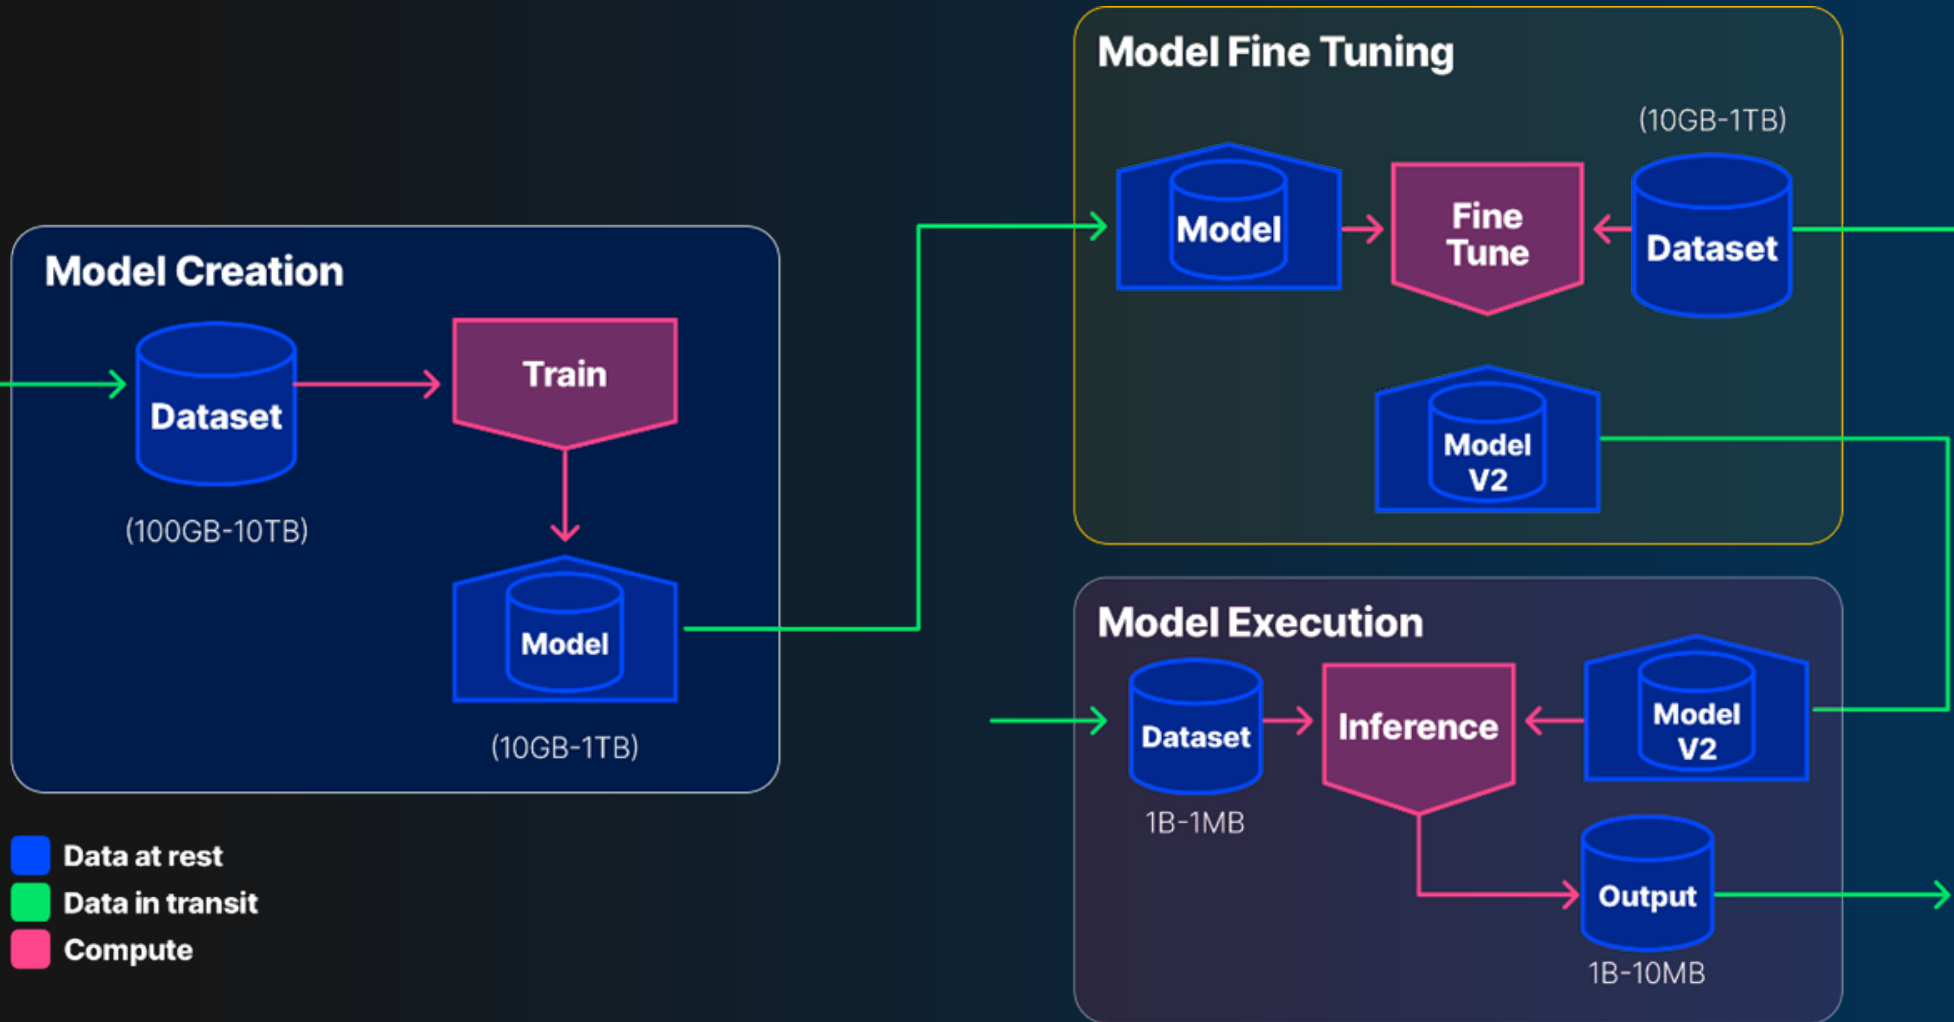
181 zettabytes on DVDs would circle the earth 222 times.



On CDs, it would get you halfway to mars.




AI Model Development Lifecycle






Cutting Carbon with Storj

U.S. Data Center Electricity Use By End Use



Carbon Cost of a 1TB Drive



Storj puts unused capacity to work.



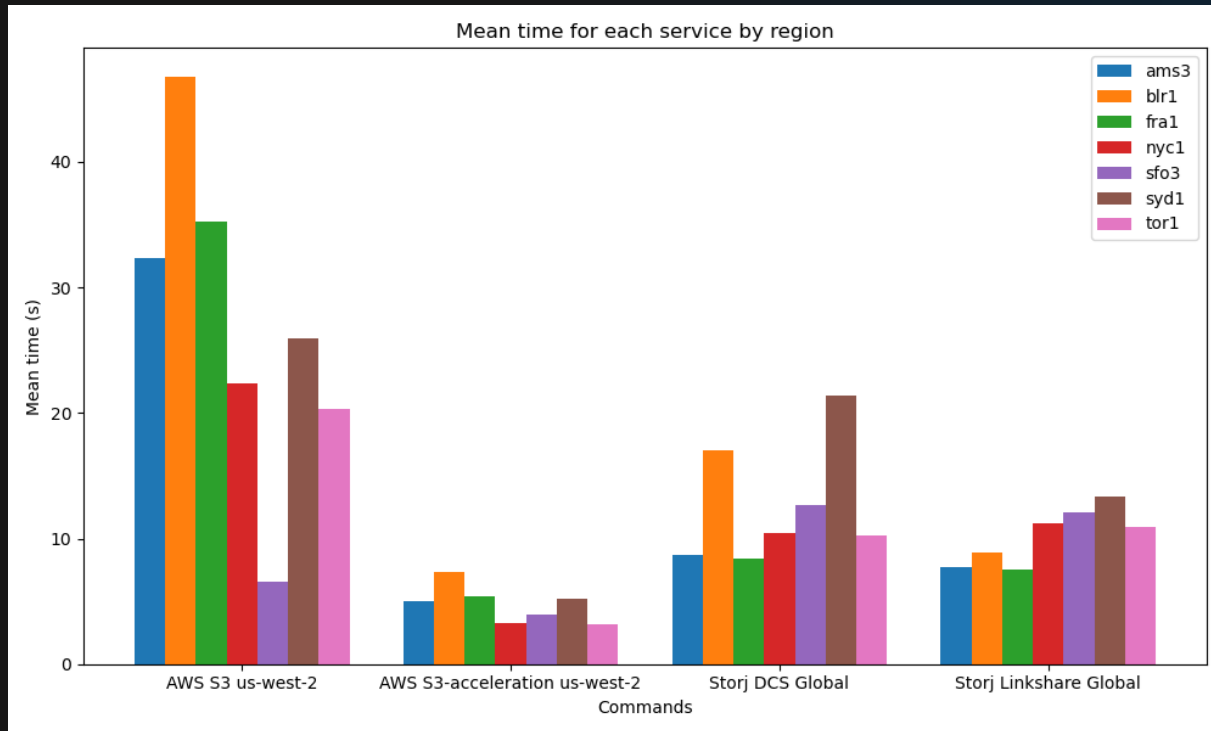
AI Workload Case Studies



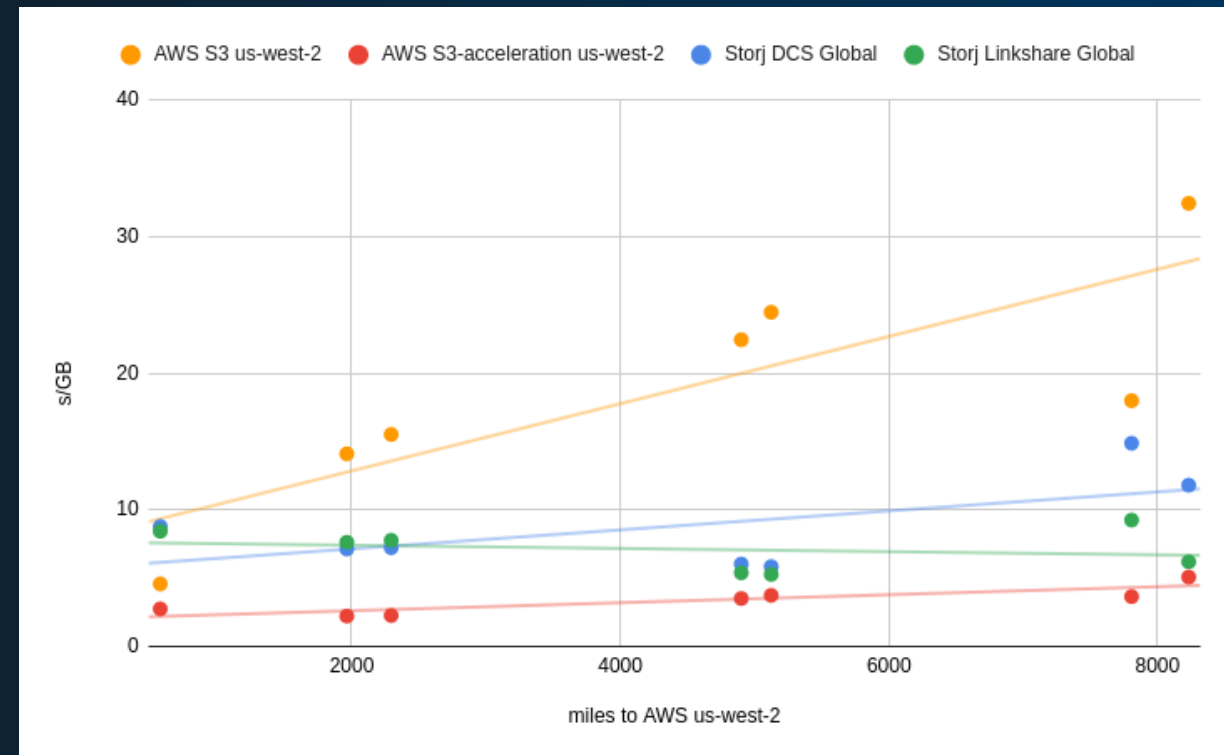
LAION-5B Dataset

1,760 encrypted erasure - encoded pieces of a 1.44GB object, stored around the world





What did we find?















Standard Storj performance in general is much better than standard S3 performance, considering global distribution

Storj HuggingFace monkey patch

Monkey patch for HuggingFace Hub to download Git-LFS blobs from Storj

This patch aims to demonstrate the transfer speed that can be achieved with `huggingface_hub` Python library when utilizing the power of the [Storj Decentralized Cloud Storage](#).










HuggingFace Hub stores all large files in Git-LFS.

 <code>pytorch_model-00001-of-00007.bin</code>	9.9 GB	 LFS	
 <code>pytorch_model-00002-of-00007.bin</code>	9.86 GB	 LFS	
 <code>pytorch_model-00003-of-00007.bin</code>	9.85 GB	 LFS	
 <code>pytorch_model-00004-of-00007.bin</code>	9.86 GB	 LFS	

When the `huggingface_hub` Python library requests to download such a file, the download request is redirected to the Git-LFS CDN hosted at `cdn-lfs.huggingface.co`.

This monkey patch modifies the `huggingface_hub` library to redirect Git-LFS downloads to the Storj Linksharing service hosted at `link.storjshare.io`.

Storj Decentralized Storage

	Multiregion	Included
	Performance	Up to 5,000 Gbps
	Availability	99.95%
	Durability	11 9s
	Privacy	AES 256 Encryption
	Security	Edge Delegated
	API	Unlimited
	Storage	\$0.004 /GB/Month
	Bandwidth	\$0.007 /GB

Questions?



Please take a moment to rate this session.

Your feedback is important to us.