



Image credit: metamorworks

**GIGA  
OM** Enrico Signoretti  
Jul 12, 2022

# GigaOm Use Case Scenario for Decentralized Object Storage for Video v1.0

Best Practices and a Practical Example for Implementation

Cloud & Infrastructure

# GigaOm Use Case Scenario for Decentralized Object Storage for Video

Best Practices and a Practical Example for Implementation

## Table of Contents

- 1 Summary
- 2 Report Methodology
- 3 Design Criteria
- 4 Additional Considerations
- 5 Solution Profile
- 6 Analyst's Take
- 7 About Enrico Signoretti
- 8 About GigaOm
- 9 Copyright

# 1. Summary

In the media and entertainment (M&E) industry, video productions invest heavily in large-scale infrastructures to store vast amounts of data, both in the cloud and on-premises. Video and other media assets coming from different sources have to be edited and rendered into the final product by teams that are often globally distributed. Depending on how the videos will be consumed, the final product is rendered in multiple versions and kept for long periods of time afterward. This challenge is even bigger now with videos that are shot at 4K and 8K resolutions.

No matter the size of the company, media-rich content requires a lot of storage capacity. It has to be reliable, fast, and, at the same time, reasonably priced. Object storage is considered one of the best options for storing unstructured data due to its scalability, cost, simplicity, and accessibility, but it also poses challenges, especially when data needs to be globally accessed and distributed. This GigaOm Use Case Scenario report explores the application of decentralized object storage in video production and collaboration.

## ABOUT THE GIGAOM USE CASE SCENARIO REPORT

This GigaOm Use Case Scenario report focuses on specific scenarios and best practices to improve adoption of technologies, exploring both use case design criteria and a viable technical solution. In this context, a particularly demanding use case for decentralized object storage can be found in video production and collaboration.

Use of video is growing in every industry for a variety of reasons: surveillance, training, marketing, conference calls archiving, and so on. These are generic use cases found in organizations of every size, but when we focus on the M&E industry, and specifically on video production, we find that:

- Videos are recorded and edited in different locations.
- Users take advantage of compute resources from different providers to render videos that must be centralized in a single location.
- All video archives are now nearline.

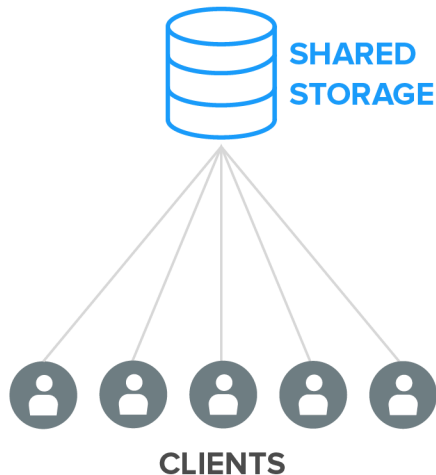
- Content has to be distributed efficiently on different channels and platforms.
- Users don't use S3 protocol directly, but want familiar file interfaces (SMB or NFS) and media asset management tools (MAM) to simplify their workflows.

These requirements create a formidable problem. Meeting high-level standards in terms of availability and resiliency can be challenging, especially when the total cost of the infrastructure must be accounted for. In this regard, the user should take into account some important aspects:

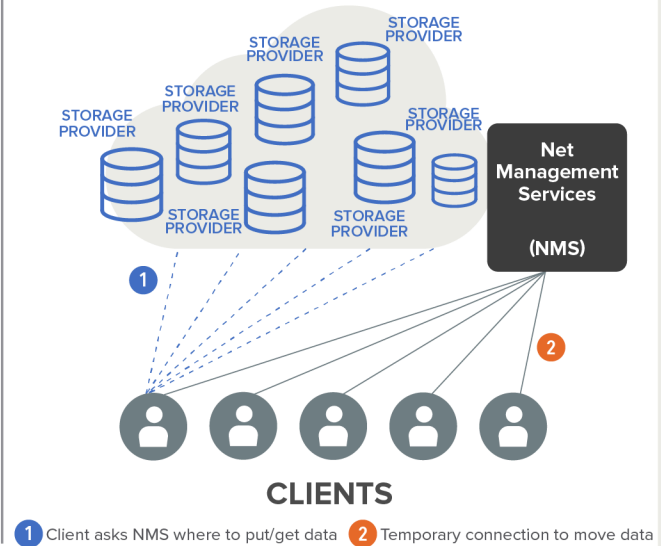
- **Infrastructure resiliency:** Object storage is usually resilient, but it is crucial to consider business continuity and disaster recovery for on-premises/hybrid infrastructures.
- **Data accessibility and availability:** Even in the public cloud, having multiple copies of data can be a requirement in case of a zone or region failure. Keeping data synchronized is very expensive and creates additional synchronization issues if data needs to be accessed concurrently from multiple locations or for backup reasons.
- **Performance:** Even though performance is not usually associated with object storage, parallelism and throughput are important characteristics to consider, especially when video is involved. Content delivery networks (CDN) are a solution, but they are expensive and complicate the infrastructure topology.
- **Scalability:** This can be an issue for on-premises deployments, especially for large systems installed in locations with limited space.
- **Cost and TCO:** Cost can be one of the biggest issues when video is involved, especially in hybrid and public cloud environments because of complex billing mechanisms and egress fees.

Decentralized storage is a solution to these challenges. A decentralized storage system is based on a peer-to-peer (P2P) network, a type of architecture that has found some success for data distribution and file sharing. Instead of storing data in a centralized system made up of data centers, it is chunked, distributed, and stored on thousands of nodes in a global network or the internet. **Figure 1** compares traditional shared storage with decentralized storage. This latest version of decentralized cloud storage has evolved greatly and is now enterprise-grade and considerably more secure, performant, private, and durable than a centralized cloud provider. It is also a fraction of the cost.

## TRADITIONAL SHARED STORAGE



## DECENTRALIZED STORAGE



Source: GigaOm 2022

Figure 1. Traditional and Decentralized Storage

We have seen several attempts over the past decade—largely unsuccessful—to build a decentralized, or P2P, network infrastructure. But risks are generally mitigated by the large number of unused commodity resources across the internet, better security, and blockchain technology that ensure data immutability and consistency. It is now easier to take advantage of the abundance of unused, and sometimes unreliable, resources to build performant and secure storage infrastructures. More so, the interest in web3 and decentralized internet technologies has attracted large investments, accelerating product development and the growth of a solution ecosystem.

## 2. Report Methodology

A GigaOm Use Case Scenario report analyzes new technology to provide decision-makers the information they need to take advantage of new technologies for highly rewarding IT strategies, while considering potential risks associated with its deployment.

This report discusses common design criteria for implementing new technology and how it may interact with existing processes and infrastructure components. The report aims to work with field-validated designs and practices to simplify adoption and minimize risks. Our analysis highlights the use case while providing the necessary information and examples to speed up the adoption process.

The idea is to define what the user should expect from this technology while providing the necessary view and basic design considerations to start the evaluation process with partners and technology vendors.

In this regard, readers will find similarities with the GigaOm Key Criteria and Radar reports. The Use Case Scenario report can be considered a follow-up to this report, going deeper into the use case evaluation and eventual technology adoption. The objective is to provide complementary and additional information to the decision-makers and inform them on how to proceed further in their technology, solution and vendor evaluation process.

### 3. Design Criteria

Decades ago, decentralized storage and peer-to-peer networks were not common in the traditional enterprise, and were usually viewed with skepticism due to their complexity, potential risks, and other challenges. The latest solutions based on this technology, however, can hide the complexity while providing a user experience similar to traditional public cloud storage, with the same interfaces, but dramatically better cost models and security. Thanks to the cloud and the rise of microservices, users are now more familiar with highly distributed applications and more open to evaluating decentralized infrastructure solutions. Here is a list of the most important characteristics of a modern decentralized storage system for video workloads:

- **Cost:** One big issue with centralized public cloud infrastructures is cost and billing complexity. \$/GB metrics can be attractive, but the total cost can quickly become a problem when it comes to egress fees, especially for data distribution, sharing and collaboration.
- **Performance:** Most object stores are not fast enough to deliver the performance needed by demanding applications like video streaming. Such applications need data close to the edge, where it is consumed, and the solution is a CDN. The architecture of decentralized storage may help reduce the need of CDNs, simplifying several aspects of the infrastructure.
- **Integration:** Decentralized storage usually offers open or proprietary API sets to access all its functionalities and characteristics. The industry is now standardized on S3 interface and it is mandatory to have this type of interface to simplify the integration with commercial applications. However, the native API approach is preferable when available because it provides better performance and additional features not usually available in the S3 protocol.
- **Solid back-end network:** Although the front end looks like any other object store thanks to the S3 interface, the key differentiator of decentralized storage is its back-end architecture and the network of active storage providers. This back-end network is central to the solution, and the user should always consider how it has been implemented alongside its size in terms of storage providers.
- **Parallelism:** If the decentralized storage system supports erasure coding, every single file is split in multiple segments and each segment is sent to a different node. This is done primarily to improve data resiliency and availability, but the positive side effect is that all communications between client and back end are parallelized, improving overall upload and download performance.
- **Data protection and integrity:** Object storage has high standards for data availability and most providers offer extensive data protection and multiple checks on data integrity. The decentralized



storage solution must at least match these standards.

- **Geo-replication and disaster recovery:** Public cloud services often allow the copying of data in different regions. This reduces latency and increases data availability in case of access problems in one of the regions. Many decentralized storage systems eliminate the need for multiple copies since data is distributed geographically.
- **Security:** Both on-premises and cloud object stores offer features to increase data security, including support for WORM APIs (ObjectLock), versioning, and end-to-end encryption. Distributed storage, thanks to blockchains and other mechanisms, provide strong versioning, data integrity, and WORM capabilities. The user should always check if these back-end characteristics are properly mapped to the S3 front end.
- **Ease of use:** For users, the most challenging aspects of decentralized storage come from the architecture's complexity, how resources are managed in the back end, and billing mechanisms (usually linked to crypto tokens of some sort). The most successful decentralized providers remove all of these by providing standard billing systems that have a simple user interface and hide the rest.

When it comes to data decentralization in the cloud, it is important to evaluate how the service is presented and delivered. Many decentralized storage solutions offer a standalone S3 gateway that can be installed and managed on-premises, which presents several challenges for a regular user. The most effective solutions are based on SaaS offers with simple pay-as-you-go subscription models. In this case it is the decentralized network provider that manages access through the S3 interface. This model is widely used for cloud-storage solutions and simplifies the procurement process.

## Benefits

Major benefits from the adoption of a decentralized storage solution for video storage include:

- **Cost and flexibility:** Many decentralized storage solutions in the market offer simple payment solutions based on pay-as-you-go models and more sophisticated options based on prepaid plans. In many cases, the user can be both storage consumer and provider at the same time, an option that adds a level of complexity, but also contributes to improving the overall \$/GB. TCO is significantly less expensive compared to traditional options, especially when multi-region data availability is included.
- **Productivity improvements:** Sharing videos becomes easier and data transfers can be globally



optimized. S3 interface can be used by file sharing and video streaming tools while the decentralized back end offers more parallelism and performance than traditional solutions. More so, modern decentralized networks can prioritize access to data segments closer to the user and reduce latency while improving speed of data movements. This enables users to more efficiently share and collaborate on the same video files globally.

- **Infrastructure simplification:** With decentralized storage there is less need for data replication to the edge. In some cases, CDNs become useless since they are designed to bring data closer to the user and reduce latency to improve performance. When the decentralized network is designed to optimize data retrieval, favoring closer and faster data storage nodes in the back end, CDN becomes redundant unless the user requires the lowest possible latency and network traffic minimization. A good decentralized network approach improves data access performance without additional caching components for several workloads.
- **Familiar user experience:** Modern decentralized storage solutions offer both native and S3 API options. The latter enables the user to take full advantage of the capabilities of decentralized storage, but the former is the best choice for most enterprises. S3 gives access to a vast solutions ecosystem, simplifies operations, and minimizes the learning curve and initial adoption. It is worth noting that most users adopt S3 object stores as a back end of other storage or data management systems such as NAS gateways, collaboration software, or MAM solutions.
- **Resiliency and availability:** Depending on the implementation, regions can just be access points to the same data. This greatly simplifies several aspects of data accessibility and removes the need for mechanisms that add latency and limitations on how data is accessed.

## Challenges

The most common challenges with the adoption of a decentralized storage solution for video storage include:

- **Size of the provider:** Most decentralized storage vendors are small and many are startups, which may be a problem for risk-averse organizations. In this case the user should start with a secondary application and use case.
- **Partner ecosystem:** Even though the S3 API is a de facto standard, work is needed to verify that the service is compatible with the apps in production. Many providers should be compatible with popular MAMs, storage, collaboration solutions, and backup solutions, and their integration certified.

- **Limited API set:** Some S3 implementations are incomplete, limiting applicability and use cases. Even when there is a list of certified products, it is important to check compatibility and list of supported APIs.
- **Billing system:** The user should always check the availability of traditional billing and payment systems. Some vendors provide complicated payment mechanisms based on crypto tokens acquired on a crypto exchange, making it difficult to use the service, while others also offer traditional payment systems such as credit cards, for example.
- **Data import:** The amount of data some organizations need to migrate to the cloud requires dedicated storage appliances because it is time consuming and not cost-effective otherwise. Many decentralized service providers don't provide this service, so availability of migration services must be verified.

For these reasons, the user should always keep an eye on the product roadmap to see how the decentralized network vendor is reacting to emerging business needs.

## Timeline

Decentralized storage with an S3 API can be adopted quickly, like any S3-based service. If the user already has data stored in other private or public storage systems, the data migration can affect the timeline and the migration process may add a considerable cost due to egress fees. For this reason the user should consider a gradual adoption of decentralized storage, starting with new projects and moving old video archives only when necessary.

## 4. Additional Considerations

In addition to what we discussed in the previous section of this document, the user should consider that the decentralized storage solution may be used for common object storage use cases such as backup and disaster recovery. S3 repositories are becoming a standard backup target for many data protection solutions, mainly as a second tier for long-term retention. In this case, a decentralized data store can provide better availability at a lower cost.

### **Emerging Tech to Consider**

Decentralized storage with a standard S3 interface enables distributed organizations to access data globally, avoiding silos while keeping access latency consistent. S3 is the simplest way to integrate this type of storage with existing applications, but by using native APIs, decentralized storage offers additional features that can increase efficiency, performance, and application portability. In this context it is likely that some software vendors will take advantage of native APIs to improve integration with the decentralized storage platform.

## 5. Solution Profile

The GigaOm Use Case Scenario report associates the theory behind prominent and emerging use cases for technology with a practical example driven by a solution available in the market. The solution is usually selected among those already proven to be effective in this situation and is already present in recent GigaOm Key Criteria, Radar, or Sonar reports.

### Storj

Storj provides an end-to-end solution composed of a decentralized storage network, development tools to upload and access storage with native APIs, a private S3 gateway, and public S3-compatible service. The Storj platform is open source, secure, and private. It also provides a way for developers to reduce complexity and save time and costs when storing and accessing growing volumes of data. Users can take advantage of the open-source code to build their own private infrastructure.

Storj's public cloud object storage service is a full-fledged object store supporting some of the latest Amazon AWS S3 APIs. It includes multipart uploads, which are in great demand among the many commercial applications in the video industry, for example.

The network was designed for privacy, security, reliability, and efficiency. It already has over 15,000 active storage nodes distributed across more than 90 countries, providing network capacity, which should attract the interest of enterprise users. Every file ingested by the system is encrypted, erasure-coded, and distributed to different locations for maximal security and availability.

To ensure high data reliability, availability, and parallelism, the erasure coding scheme has up to 80 segments, with only 29 necessary to reconstitute any file. This increases the overall security of the system and improves parallelism and read performance for read-intensive workloads. The Storj network protocol continuously checks node availability and service level to guarantee data availability and discover potential network issues, eliminating bad actors on the network and rebuilding data when necessary.

Storj capacity providers (nodes) are well incentivized to keep the service running thanks to the ability to monetize efforts quickly and transparently using traditional payment methods. The Storj partner program offers options to third-party developers and solution providers to be an active part of the ecosystem. The list of Storj partners for video management and handling includes NAS gateways and MAM solutions, among others. It is worth noting that Storj offers a service level agreement (SLA)

comparable with that of AWS S3 in terms of data durability and availability.

Storj is an end-to-end solution that provides both developer- and enterprise-grade experiences, including a large and growing network of crowd-sourced storage nodes with a choice of upload options (including an S3 compatible gateway, browser, and CLI), and an expanding solution ecosystem. The Storj platform is open source, extra secure and private, and provides an exceptional way for IT departments and developers to reduce complexity and save time and costs when storing and accessing growing volumes of data.

## 6. Analyst's Take

Storing and operating large amounts of data is challenging, even more so for video content, and many users are discovering its associated costs too late. Traditional on-premises solutions may be low on initial acquisition cost, but they do not provide the availability and resiliency required by many users. More so, video production workflows are becoming increasingly complex with video recorded in one location, edited somewhere else by multiple operators, sometimes rendered in the cloud, and finally distributed on several channels. Hybrid and public cloud approaches are also expensive and often complex to manage, requiring considerable work (admins-per-petabyte). Decentralized storage can be an alternative that can be cost conscious, agile, and efficient.

With increasing interest in web3 and a more decentralized internet, decentralized storage can be a good first step in this direction. Decentralization removes some of the security, privacy, cost and lock-in risks/concerns associated with centralized services. At the same time, thanks to the standard interfaces, even the most conservative of users can find a familiar interface and safely investigate this technology and its benefits. The goal is to be ready when there are a large number of solutions taking advantage of this paradigm and architecture. This is especially true for video storage, where organizations of all sizes must keep up with the increasing amount of data, improve distributed workflows, and share it safely for consumption.

In this scenario, Storj offers a good combination of features, is easy to use, and focuses on efficiency and total cost of ownership. It is working with several solution providers in this space to integrate and certify its solution, offering a simple yet video-friendly fee structure, payable via credit card or invoicing. The benefits of the Storj solution for users in the video industry are clearly visible in terms of total cost of ownership when the reasonable \$/GB is associated with ease of use and the productivity improvements coming from decentralized network technology.

## 7. About Enrico Signoretti

Enrico Signoretti has more than 25 years in technical product strategy and management roles. He has advised mid-market and large enterprises across numerous industries, and worked with a range of software companies from small ISVs to global providers.

Enrico is an internationally renowned expert on data storage—and a visionary, author, blogger, and speaker on the topic. He has tracked the evolution of the storage industry for years, as a Gigaom Research Analyst, an independent analyst, and as a contributor to the Register.



## 8. About GigaOm

GigaOm provides technical, operational, and business advice for IT's strategic digital enterprise and business initiatives. Enterprise business leaders, CIOs, and technology organizations partner with GigaOm for practical, actionable, strategic, and visionary advice for modernizing and transforming their business. GigaOm's advice empowers enterprises to successfully compete in an increasingly complicated business atmosphere that requires a solid understanding of constantly changing customer demands.

GigaOm works directly with enterprises both inside and outside of the IT organization to apply proven research and methodologies designed to avoid pitfalls and roadblocks while balancing risk and innovation. Research methodologies include but are not limited to adoption and benchmarking surveys, use cases, interviews, ROI/TCO, market landscapes, strategic trends, and technical benchmarks. Our analysts possess 20+ years of experience advising a spectrum of clients from early adopters to mainstream enterprises.

GigaOm's perspective is that of the unbiased enterprise practitioner. Through this perspective, GigaOm connects with engaged and loyal subscribers on a deep and meaningful level.

## 9. Copyright

© [Knowingly, Inc.](#) 2022 "*GigaOm Use Case Scenario for Decentralized Object Storage for Video*" is a trademark of [Knowingly, Inc.](#). For permission to reproduce this report, please contact [sales@gigaom.com](mailto:sales@gigaom.com).