GIGAOM

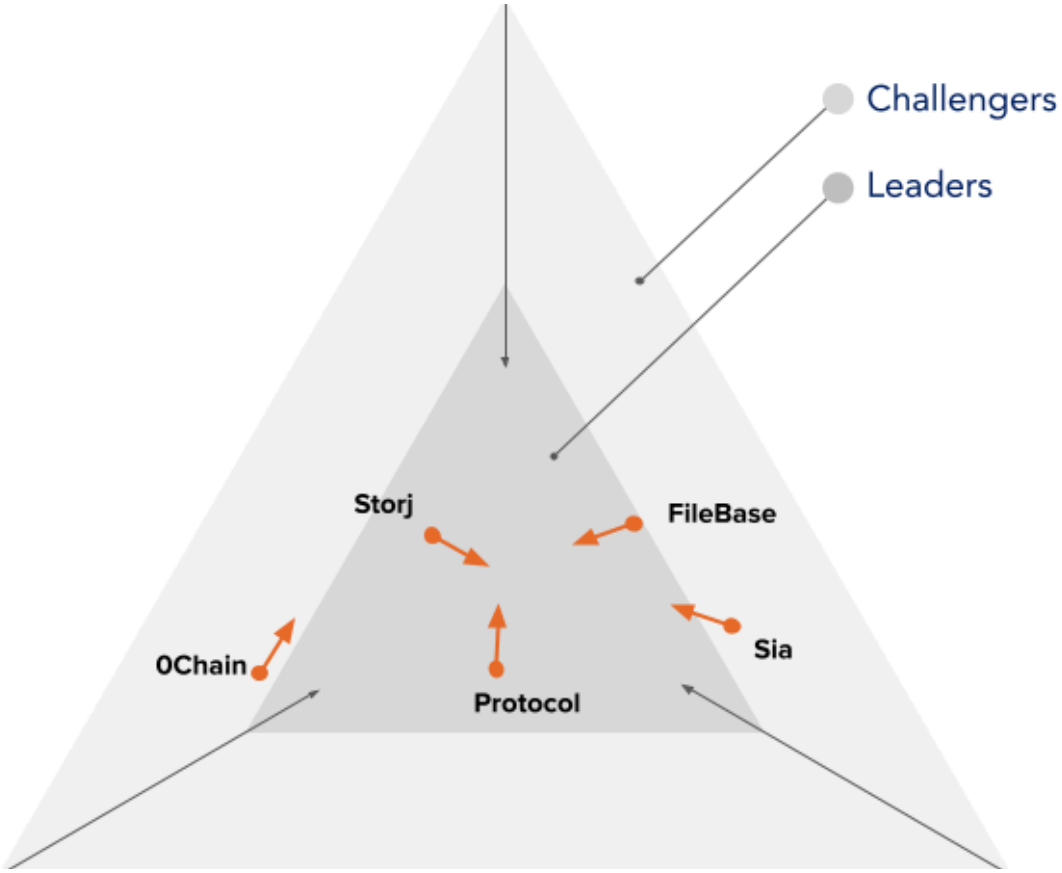

Enrico Signoretti
Jul 12, 2021

GigaOm Sonar for Decentralized Storage

v1.0

An Emerging Technology Insight Report

GigaOm Sonar for Decentralized Storage

An Emerging Technology Insight Report

Table of Contents

1. Summary

A decentralized storage system is based on a peer-to-peer (P2P) network, a type of architecture that found some success in the past for data distribution and file sharing. Instead of storing all the data in a centralized system, data is chunked, distributed, and stored on many nodes in a local network or the internet. **Figure 1** compares traditional shared storage with decentralized storage.

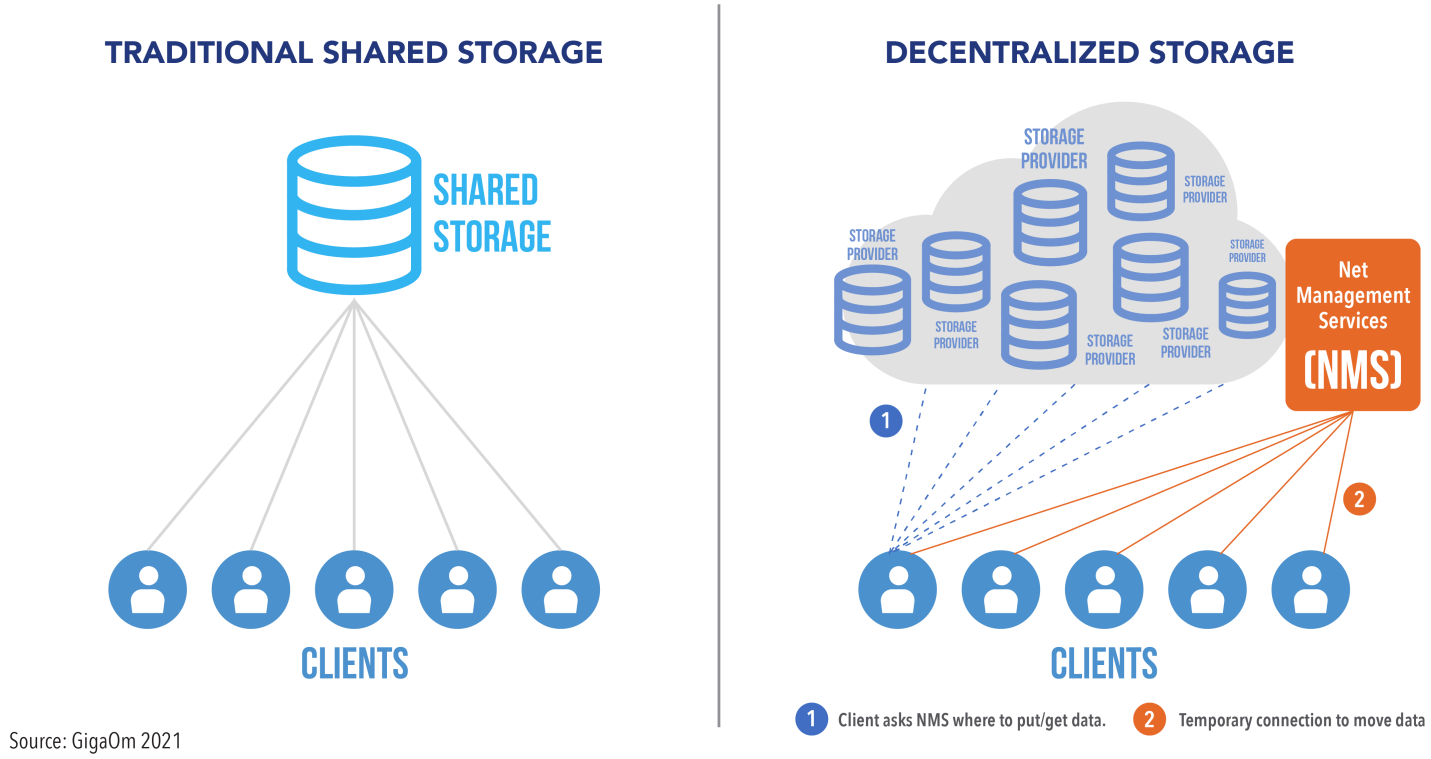

Figure 1. Comparison Between Traditional and Decentralized Storage Systems

We have seen attempts—largely unsuccessful ones so far—to build a decentralized, or P2P, network infrastructure over the years, but the risks are now mitigated by recent developments, such as the large number of unused commodity resources across the internet, better security, and blockchain technology to ensure data immutability and consistency. In the end, it is now easier to take advantage of the abundance of unused, and sometimes unreliable, resources to build strong and secure storage infrastructures.

Decentralized storage and peer-to-peer networks are not common in the traditional enterprise and they are usually viewed with a lot of skepticism due to overall complexity, potential risks, and other challenges. That said, the latest generation of solutions based on this technology can hide most of the complexity while providing a user experience similar to traditional public cloud storage, with better cost models and security. Moreover, thanks to the cloud and the rise of microservices, enterprises are more familiar with highly distributed applications and more open to evaluating decentralized infrastructure solutions.

How We Got Here

The public cloud as we know it has a tremendous impact on IT agility and flexibility, but it can also lead to increasing costs and potential lock-ins. Particularly with growing business needs that lead to the adoption of hybrid and multi-cloud solutions, users are finding it more and more difficult to consolidate data on a single public cloud or on-premises. Additionally, applications are becoming more mobile and data needs to follow along, but this may not be possible due to data gravity, time, and cost of migrations. This situation creates inefficiencies that can limit an organization's ability to build effective multi-cloud strategies.

In order to maintain flexibility, minimize risks of lock-in, and keep costs down, users are looking to build data repositories that are cost effective and accessible from everywhere, with good performance and with standard access interfaces. Data consolidation also simplifies data management and security. In this regard, decentralized cloud storage is a great alternative to traditional cloud storage providers.

Blockchain technology plays a fundamental role in modern decentralized storage. Not only can it be used to create an internal currency to simplify raw storage resources and data services trading, it can also improve peer-to-peer network security, information integrity, and data reliability.

Even though the basic concepts behind decentralized storage technology solutions are similar, implementations differ greatly. There are several players already focused on different use cases and market segments, with many of them targeting developers and enterprises looking for relatively inexpensive and highly secure storage.

About the GigaOm Emerging Technology Impact Report

This GigaOm report is focused on emerging technologies and market segments. It helps organizations of all sizes to understand a technology, its strengths and weaknesses, and its fit within an overall IT strategy. The report is organized into four sections:

- **Technology Overview:** An outline of the technology, its major benefits, possible use cases, and the relevant characteristics of different product implementations in the market.
- **Considerations for Adoption:** An analysis of the potential risks and benefits of introducing products based on this technology into an enterprise IT scenario, including table stakes and key differentiating features, as well as consideration of how to integrate the new product into the existing environment.
- **GigaOm Sonar:** A graphical representation of the market and its most important players focused on their value proposition and their roadmaps for the future. This section also includes a breakdown of each vendor's offering in the sector.
- **Near-Term Roadmap:** A 12-18 month forecast of the future development of the technology, its

ecosystem, and major players in this market segment.

2. Report Methodology

A GigaOm Emerging Technology Insight report analyzes new technology trends and aims to provide decision makers with the information to build forward-looking, and potentially highly rewarding, IT strategies while taking into account the risks of adopting products that are not yet fully validated by the market or available from established players.

In talking about bleeding edge technology and without clear categorizations of the market players, the goal of this report is to eliminate hype and educate our readers on the technology before helping them to navigate through different product implementations. The analysis here is focused on highlighting core technology, use cases, and differentiating features, rather than drawing feature comparisons. This approach is taken mostly because the technology area is new and the overlap between solutions is minimal at this point. In fact, even if core technology is based on the same principles, the product implementations are still differentiated and focused on particular use cases.

The goal is to define the basic features the user should expect from products that correctly implement the technology, while keeping an eye on characteristics that will have a role in building the differentiating value over time.

In this regard, the user will find similarities with the Key Criteria and Radar reports. The Sonar report can be considered an early assessment for recently introduced technologies and market segments. The evaluation method for emerging technology is based on:

- **Core Technology:** Emerging table stakes
- **Differentiating features:** Potential value and key criteria

Over the years, depending on technology maturation and user adoption, emerging technology may remain a niche or may evolve to become mainstream, as shown in **Figure 2**. The objective of this report is to intercept new technology trends before they become mainstream and provide the tools to understand their value for potential early adoption and the highest ROI.

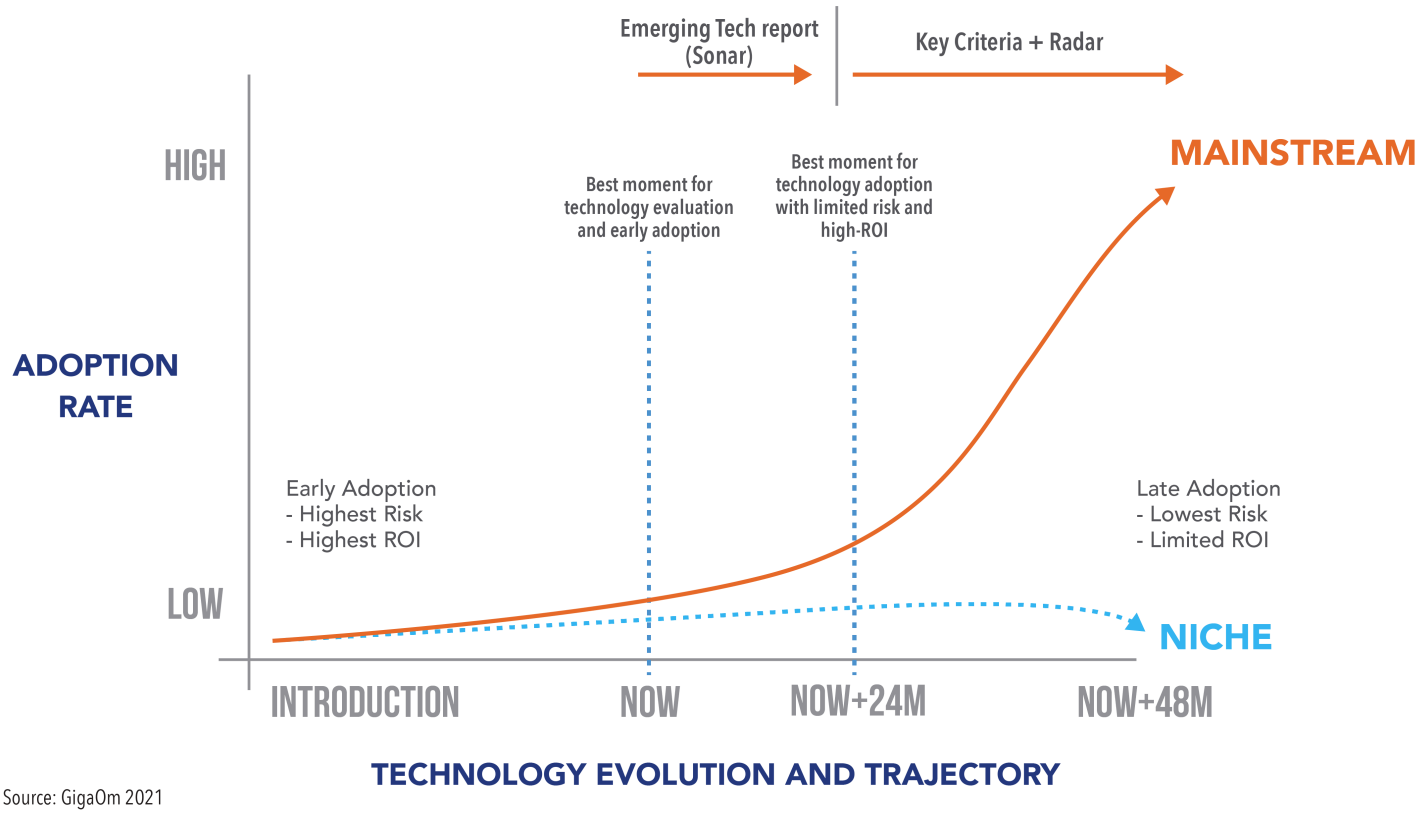

Source: GigaOm 2021

Figure 2. Evolution of Technology

3. Technology Overview

One of the biggest challenges with decentralized storage is to build a global peer-to-peer network that's reliable and extensive and has mechanisms to keep data healthy and always accessible.

A decentralized storage infrastructure is built out of several components that work together to provide the required level of data integrity, resiliency, and availability. Although the specific implementation may vary, the main components of this type of infrastructure will consist of:

- **Storage providers:** These are storage nodes on the internet with available raw storage capacity and bandwidth to share. Depending on the implementation, these nodes can be either small or large, from a few hundred gigabytes to hundreds of terabytes. Some vendors provide software for all the major platforms (Linux, Windows, and macOS) and the user needs only to allocate free space for this service. Anyone with some free capacity, a stable internet connection, and a computer with good uptime can be a candidate.
- **Network management services:** This is a set of services for monitoring the health of the network, the availability of storage nodes and their performance, metadata management, and data protection. These services also check data integrity, restore data if necessary, and indicate where data is actually located. Encryption, authentication, and logging are key features as well. These services are usually managed by a network operator, typically the service provider. They can be centralized or distributed depending on the implementation and are fundamental for the correct functioning of the entire infrastructure. In fact, they can be part of the blockchain infrastructure for improved decentralization and security.
- **Storage consumers:** These are the actual clients, and might be an application using native APIs or a gateway providing compatibility with standard protocols such as S3. For example, some decentralized storage vendors offer both APIs and different kinds of gateways that can be accessed globally, like other cloud storage services, or on-premises, depending on the use case.

In this type of architecture, all data is usually encrypted in transit and at rest, and, to assure adequate protection, it can be replicated on multiple nodes or erasure coded and then distributed to limit its footprint. Furthermore, to minimize the impact on a single storage node and improve overall data availability, each file is usually divided into small chunks that are distributed on as many storage nodes as possible. All of these data protection and optimization mechanisms are designed with the idea that some storage nodes may not be available when necessary. That is why the system creates a number of copies of data, five to 10 in most cases, or very long erasure codes (for example, 30 segments with at least 20 of them necessary to rebuild the information). Some solutions also allow users to select the data protection mechanisms that best suit their needs.

As noted, blockchain technology has a role in decentralized storage systems. Although it was not necessary for building a traditional peer-to-peer network, it is now common among all decentralized storage providers, and it can be useful in two ways:

- **Better security:** Each transaction can be recorded in a distributed ledger to eliminate risk of data tampering and improve overall data integrity. This is particularly useful when the storage system holds sensitive information, and for long-term data retention, compliance, and secured write-once-read-many (WORM) repositories.
- **Internal currency and marketplace:** Building a secure digital currency (or utility tokens) to sell storage resources or purchase services simplifies trading on the platform. At the same time, it creates the foundation for a system to properly remunerate those who offer resources on the network. In fact, one of the most important ways to build a healthy and expansive peer-to-peer network is to create the right incentives to make resource sharing appealing for a large number of users. Having an internal blockchain-based currency allows for the efficient management of micropayments and ensures that all transactions are properly recorded.

Additionally, many decentralized storage systems use smart contracts to simplify resource trading and automate payments between storage providers and consumers. These constructs, common within the cryptocurrency world, are used to define the type of resource shared by the provider (bandwidth, capacity, performance, and service availability) and the cost of the service. The smart contract is also used to verify whether the provider and consumer are effectively working as agreed.

Use Cases

Due to the nature of both peer-to-peer architectures and geographical data distribution, access latency can be inconsistent, limiting the number of use cases for which this technology can be a good fit. We are seeing the following common use cases:

- **Active archives:** The cost of decentralized storage can be lower than other cloud storage options, oftentimes even lower than cold-tier storage from major cloud providers. Decentralized storage can be a cost-effective option while offering faster data retrieval.
- **Backup:** Large decentralized storage networks can provide high parallelism and high and sustained throughput, even though single transactions may be slow due to the amplification of operations in the backend. Thus, decentralized storage can be used as a backup target, especially to store secondary and long-term retention backup copies.
- **Content distribution**: Decentralized storage is a good object storage backend for content distribution. Data is globally distributed and highly available while the large number of copies or erasure-coded chunks ensures good read performance. Additionally, if the decentralized storage system can manage read prioritization from storage providers on the back end, users will receive data consistently and from the nearest locations.
- **Video storage:** Today, enterprises are storing more video data, including surveillance recordings and training and marketing material. The original copies of these videos are often stored for weeks, months, or even years, but are rarely accessed after the first few hours. Decentralized storage can provide the necessary performance to stream the file to a limited number of devices when necessary, while retaining the video over a long period of time.

- **Mobile and cloud-native applications:** Developers can take advantage of the technology to store large amounts of data with secure global access at a relatively low and predictable cost. By using native APIs, they can also optimize data access and further improve security and privacy.
- **Secure collaboration and private data sharing:** Decentralized storage offers the strongest security standards and some vendors are building secure data sharing and collaboration applications with embedded security features. These apps enable users to work in a very familiar and easy-to-use environment without the risk of malicious attacks, such as data thefts and tampering.

4. Considerations for Adoption

Before adopting a decentralized storage solution, the user should clearly understand the benefits and risks associated with this still-emerging technology. Decentralized storage solutions based on peer-to-peer networks have been successful in the consumer space, thanks to solutions like BitTorrent, but their adoption in the enterprise, and for data storage applications, has been limited. A few attempts were made in the past, but the technology wasn't mature, internet bandwidth was limited, and the risks were too great compared to the benefits. Now internet bandwidth is less of an issue and the technology has matured enough to become viable in the enterprise. That said, it still may be difficult for a traditional enterprise to adopt modern decentralized storage.

Decentralized storage can be implemented on-premises or in the cloud. Unfortunately, to be effective on-premises, decentralized storage requires a large quantity of resources geographically distributed. This limits on-premises deployment in small and medium-sized organizations, for which the only option is the public cloud. Large enterprises and service providers could consider adopting this technology on-premises because of the vast and far flung quantity of unused resources at their disposal, but there are a few important considerations to make before starting, including:

- **Resource management:** One of the main benefits of decentralized storage is the ability to take advantage of unused resources on the network and use them to build a consolidated storage pool. The user can tap into either dedicated storage nodes or capacity available on servers doing other jobs, or a combination of both. In the first case, the cost is not very different from traditional object stores. With the latter, keep in mind that there will be some contention over resources eventually.

  For example, consider a primary application running out of space because of capacity now allocated to decentralized storage. Reclaiming disk space can be easy, but monitoring activity and moving data to other repositories can be challenging. The solution chosen for the decentralized storage system should include the tools to monitor resource consumption and avoid this type of risk.

- **Infrastructure management:** Storage providers are software components that need to be installed on each node of the network. In a large network there will be several thousand nodes working together, all of which need to be installed, configured, managed, and updated over time. The time this will take can be substantial, and even with the right automation tools it can be challenging to keep the network running at full efficiency while keeping it secure and highly available.

Because of the complexity in an on-premises deployment, public cloud implementations of decentralized storage have more chances of success. In fact, infrastructure management is also decentralized and assigned to each single storage provider. Storage nodes are kept efficient and accessible by their owners, who are incentivized to do so because of the utility tokens they receive if the service is up, efficient, and at available capacity. Tokens can be converted into money or used to purchase other services on the same network, which both mitigates and distributes risk. Storage consumers do not see any of the complexity at the backend. The biggest risk in this type of scenario

is that the organization that manages the network services, including the token and money conversion, won't be able to provide the right incentives or quality of service. Transparency is key in this case. The user should always have the ability to check the status of the network and its history and trends before storing data on it.

In the end, once the complexities and risks are eliminated, decentralized storage can be considered a good option for storing secondary data securely and at a very reasonable cost. There are several options available now and they differ both in the technology used in the back end and in the potential use cases they can serve, ranging from long-term cold data storage to content distribution and more. API support is another important aspect. In fact, a growing number of services offer S3 compatibility, which simplifies adoption and avoids any form of lock-in.

Key Characteristics for Enterprise Adoption

The most important characteristics of decentralized storage to evaluate before adoption include:

- **Technology:** The user should evaluate the technology behind this type of service. In general, every component should be decentralized and highly redundant to minimize single points of failure, risk of service disruption, scalability issues, and exposure of the network to malicious attack.
- **Incentive mechanisms:** The incentives to keep the network stable and growing should be easily identifiable and quantifiable. Utility tokens should be available and should be associated with conversion methods that monetize the tokens and make them available for purchase with traditional payment systems.
- **Security:** This is crucial. The user should always check on encryption methods for data in transit and at rest, how data is distributed over the network, and data integrity mechanisms, including the use of blockchain to ensure that data has not been tampered with.
- **Access protocols:** The service provider should make the service available via standard APIs like S3, as well as provide native APIs that take full advantage of the features available in the decentralized storage system.
- **Ease of use:** The product should be easy to use, provide a consistent user experience, and feature tools to minimize impact on daily operations. In addition, the solution provider should offer SDKs, ready-to-go integrations, best practices, and certifications for common products usually present in enterprise infrastructures.
- **Roadmap:** The vendors in this space are relatively small and new to the market. Knowing their strategy and future development plans enables decision-makers to evaluate whether a solution has the potential to meet their needs going forward. In this regard, open-source products should be favored because it is easier to check on adoption trends and community and development activity.

Table 1 shows how well the solutions discussed in this report did in each of these areas.

5. GigaOm Sonar

The GigaOm Sonar provides a forward-looking analysis of vendor solutions in a nascent or emerging technology sector, based on each vendor's strategy, technology, and roadmap. The GigaOm Sonar chart plots the current position of each solution against these three criteria across a series of concentric triangles, with those set closer to the center judged to be of higher overall value. The forward-looking progress of vendors is further depicted by arrows that show the expected direction of travel over a period of 12 to 18 months.

The GigaOm Sonar chart (**Figure 3**) is defined by three axes. They are:

- **Roadmap:** When assessing emerging technologies, it is important to take a forward-looking approach and to not only describe the necessary requirements for initial adoption, but to also understand the expected future development of the technology. This is particularly important for organizations that seek to expand beyond an initial targeted use case to maximize employment of the solution in a way that can enhance return on investment (ROI).
- **Technology:** It is critical to understand whether the vendor has developed a solution with the necessary differentiation to stand out in the crowd. Is the product architecture solid and ready to support the growing number of features and capabilities that users will require over time?
- **Strategy:** This metric takes into account the vendor's go-to-market strategy and its ability to create a solution ecosystem around its product. Strategy also reflects the company's ability to articulate its vision and accomplish the goals on its roadmap.

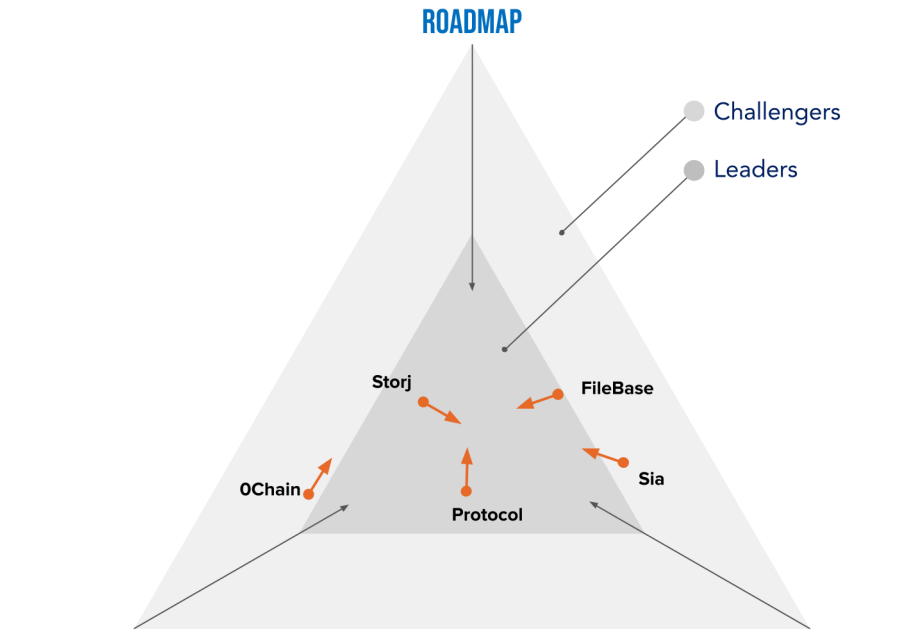

Source: GigaOm 2021

Figure 3. GigaOm Sonar for Decentralized Storage

As you can see in the Sonar chart in **Figure 3**, Storj and Filecoin are in the leading positions. Storj is a good fit for developers and SMEs, while Filecoin has a solid architecture that may be more familiar to larger enterprise users. Filebase is a simple solution and does not own its network, but can take advantage of Sia and Storj, giving small organizations an easy adoption path to this technology. Sia, with its Skynet, and 0Chain both have compelling solutions on paper but they still need to prove themselves in large-scale storage applications.

Table 1: Key Characteristics for Enterprise Adoption

	KEY CHARACTERISTICS					
	Technology	Incentive Mechanisms	Security	Access Protocols	Ease of Use	Roadmap
0Chain	+++	+++	+++	++	++	+++
Filebase	++	++	+++	+++	+++	+++
Filecoin	+++	+++	+++	+	++	+++
Sia	+++	++	+++	+	+	++
StorJ	+++	+++	+++	+++	++	+++

+++ Exceptional: Outstanding focus and execution
++ Capable: Good but with room for improvement
+ Limited: Lacking in execution and use cases
– Not applicable or absent

Source: GigaOm 2021

6. Market Landscape

0Chain

0Chain is one of the most recent startups in this area. Its main architectural differentiator is performance, which enables users to expand the number of potential use cases for decentralized storage. In fact, the product is designed to allow a complete separation of all the services that are part of the network into different nodes. The technology is not yet available on the mainnet (currently in Betanet). However, the company is working on several products that will take advantage of the core technology, including a private version of the network for large organizations (0Lake) that is already available on Amazon AWS and Oracle Cloud.

The flagship solution developed by 0Chain is 0Box, which provides a familiar user experience with all the benefits of decentralized storage, including infinite storage on mobile, desktop, and Web platforms. With one phone number, the user can have multiple wallets, and each wallet can have multiple allocations, which is ideal for personal users and also allows organizations to scale their storage needs across departments and employees. Both users and organizations will be able to purchase storage in USD, ETH, or ZCN and have a seamless experience.

The user has very granular control over data stored on the network, including the erasure code scheme, actual data location, required performance, and capacity. This enables the user to create several storage tiers for different applications.

Ochain has already launched its cryptocurrency, ZCN, which is traded on some of the major exchange cryptocurrency markets. ZCN is used to remunerate storage providers as well as to purchase services. The experience can be simplified through regular credit card transactions for applications like 0Box.

The company is also working on S3 and IPFS interfaces for developers who are already using those protocols in their applications and want to be able to transition to 0Chain storage without any code changes. In addition, 0Chain is developing S3 migration tools, an S3-compatible server, and a Kubernetes storage class for microservices to send data to 0Chain storage. And it is working on an exciting feature—a virtual network file server (NFS) that can be directly mounted onto a Linux system.

0Chain has incorporated a free storage protocol that will allow 0Box subscribers and any corporation to issue free storage coupons to utilize its dStorage platform. This creates an incentive mechanism for the consumption of storage and enables earnings for blobbers (storage providers).

0Chain is also offering features to support non-fungible tokens (NFTs), which includes data permanency using multi-payer pools, allocation for immutable NFTs, enabling multiple data assets tied to an NFT, provisions for live-streamed NFTs, and atomic transfer of data assets upon sale of an NFT.

Strengths: 0Chain's fully decentralized and very efficient protocol and architecture can potentially include a wider variety of use cases. The ecosystem already includes some interesting applications

and this capability can be a winning factor in the medium and long term.

Challenges: Still in Betanet, there is no proof yet that the network can scale without issues. Even though the roadmap is very compelling, there is no official S3 interface available at the moment (it should be available later this year) and this could have an impact on enterprise adoption.

Filebase

Filebase provides an S3-compatible object storage service and it will soon launch an on-premises version of its product as well. Data ingested into the system is stored across multiple decentralized storage networks such as Storj and Sia. From this perspective, the solution can be seen as a broker for different secure storage services, minimizing the risk of lock-in for the customer. The solution is very cost-effective and the user can choose which network to store data on, with the ability to switch to a different one at any time. The pricing policy is very attractive and comparable to cold storage options from mainstream providers, while the service provides geo-redundancy and performance.

Filebase also offers a caching technology to further improve access performance and reduce costs. The aggressive caching and increased performance enable the solution to serve a broader range of use cases, including active archiving and large file transfers to backends for mobile applications. The Filebase user interface is easy to use and includes a basic object browser to help the user with basic tasks such as manual data uploads and downloads. Buckets can be private or public, simplifying data sharing when necessary.

S3 compatibility still needs improvement but it includes most of the common methods needed by commercial applications.

Strengths: Filebase uses multiple storage networks to minimize risk and give more options to users. The service is easy to use, offers good performance, and is cost effective.

Challenges: S3 compatibility needs further development to expand the number of possible use cases. A more sophisticated mechanism for data placement and replication in the backend could improve efficiency and data availability of the system.

Filecoin

Protocol Labs has been among the many development contributors to the open source IPFS and Filecoin projects. Filecoin is a fully decentralized storage network while IPFS is a distributed peer-to-peer network that allows distributed apps to address data via a content hash instead of according to where it's located on the network (its IP address). IPFS and Filecoin share a very large ecosystem with more than 220 applications and clients that have already adopted this solution for both technical reasons and its attractive cost. In fact, the network has already passed seven exabytes in total capacity with a large and growing community of storage providers (miners)—more than 2,500. The software stack was built by more than 7000 Github contributors who actively worked on the open

source projects.

Unlike many solutions in this space, most of the storage providers in the Filecoin network have committed to offering datacenter resources and have made an important initial investment in both hardware and collateral to ensure the quality of service, data availability, and long-term data reliability. Data reliability is guaranteed through the Proof-of-Spacetime algorithm that cryptographically challenges miners to verify data on a daily basis. In fact, storage providers are highly incentivized to store important data for end users by being offered a 10x reward multiplier through the Filecoin+ program. Storage providers participating in the network are required to pledge Filecoin collateral and can lose that money if they are not reliable enough. At the same time, end users gain full control of their data by stipulating smart contracts to get access to the best providers and protection over time.

Most common use cases range from a backend for content delivery networks to immutable storage for large archives and NFTs. Other solutions use the Filecoin network for storage that is frequently accessed around the globe.

Strengths: Filecoin is a large and growing network and ecosystem, with strong incentive mechanisms for large data center storage providers. This can be viewed positively by enterprises and may help boost initial adoption.

Challenges: Even if prices are competitive when compared to major cloud storage providers, the lack of erasure coding could be a challenge in terms of overall system efficiency.

Sia Foundation

Sia Foundation offers an open source platform that provides services to store and retrieve data in a decentralized fashion, along with a marketplace where utility tokens, Siacoins, can be exchanged and mined. The Sia platform makes heavy use of blockchain technology, and the primary goal is to create a decentralized, peer-to-peer network to offer services for developers and users needing to implement and access different types of storage services. Sia Foundation develops the core network services while Skynet Labs works on Skynet (discussed below).

With Sia, files are always erasure-coded and encrypted before being uploaded to the network, and each segment is stored on a different node. Sia also provides mechanisms for checking data integrity alongside all the necessary network checks to verify and rebuild data availability on the network. Other mechanisms are in place to remunerate storage providers (hosts) with Siacoins.

Last year, Skynet Labs launched Skynet, a new service dedicated to content hosting and distribution based on the same core technology. Skynet is aimed at simplifying the adoption of Sia technology for developers who need to store and access data across the globe in a secure fashion, with the data always available and censorship-resistant.

A growing number of third-party solutions can now take advantage of the backend network

infrastructure provided by Sia. Most of these solutions aim to provide a storage service that combines the strengths of Sia with a much more user-friendly interface, such as S3 or FUSE volumes.

Strengths: Sia is a blockchain-based network with a growing ecosystem and some end-to-end applications that serve as proof of concept for the core technology while making the platform more accessible to a broader audience. The network and its ecosystem are still small but growing at a steady rate.

Challenges: Even though they are available from third parties, there is no direct support of standard APIs like S3 yet. The network seems limited in the total amount of capacity under management at the moment and has yet to prove its scalability. Skynet is a promising solution but is not user friendly, and monetization mechanisms are still to be implemented.

Storj

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

Storj DCS, Storj's public cloud object storage service, is a full-fledged object store supporting some of the latest Amazon AWS S3 APIs, including multipart uploads, features in great demand among the many commercial applications in the enterprise space.

The network was designed for privacy, security, reliability, and efficiency and can already count more than 13,000 active storage nodes distributed across more than 90 countries, providing capacity to the network, 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 is made of up to 80 segments, with only 29 necessary to reconstitute any file. This also increases the overall security of the system and can be of help to improve parallelism and to better read performance for read-intensive workloads. The Storj network protocol continuously checks node availability and service level to guarantee data availability and to 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 their efforts quickly and transparently, while enterprise users can decide to purchase storage resources by acquiring STORJ tokens or by using traditional payment methods. The Storj partner program offers additional options to third-party developers and solution providers to be an active part of the ecosystem. It is worth noting that Storj offers a service level agreement (SLA) that is comparable with that of AWS S3 in terms of data durability and availability.

Strengths: An end-to-end solution that provides both developer- and enterprise-friendly experiences which include 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 developers to reduce complexity and save time and costs when storing and accessing growing volumes of data.

Challenges: Even though Storj provides an S3-compatible service, the number of commercial certified solutions is still limited and this slows enterprise adoption. The system is not fully decentralized, with Storj keeping control over some of its components. This could be seen as an advantage by some enterprise users, while others might argue that the solution is not fully decentralized, which comes with its own concerns. In this regard, large organizations can always decide to implement a private infrastructure based on the same technology.

7. Near-Term Roadmap

Most of the vendors in this area are working to expand their networks and improve the usability of their solutions. The most attractive services are those that offer interfaces that are familiar to enterprise and commercial applications and this will continue for the foreseeable future. Other solutions will remain in niches and will be unlikely to reach the level of growth acceleration necessary to attract investment and customers. How these vendors will be able to attract more storage providers and expand their backend is an important consideration and a metric that users should always take into account in their evaluations.

In the near future, SMBs are more likely than large enterprises to adopt this type of service. Large organizations are likely to stay at the window and wait to see if they will be able to take advantage of decentralized storage for their secondary or tertiary cold data. That said, some solutions could be very effective for storing certain types of data and might be able to provide resiliency at a low cost.

Developers are another target for decentralized storage, as it gives them a cost-effective option while providing better control over data with embedded security, geo-distribution, and other features that are very expensive and more difficult to manage in traditional environments. That said, developers also prefer to use standard APIs instead of risking too much with an API or a service that might not be available in the future. Clearly, the ecosystem is very important and decentralized storage networks that can provide standard API gateways such as S3 are more appealing.

The most competitive vendors will work closely with partners to build a solutions ecosystem that includes a broader set of commercial applications. At the same time, most of these vendors do not yet have a clear go-to-market strategy, and this deficiency could hamper their growth. Vendors building a partner network that includes system integrators, MSPs, and VARs will be more successful than others.

Some vendors will also present new products aimed at on-premises infrastructures. Although the idea is compelling on paper, implementation of the products and their manageability in large enterprise environments will be key to success.

8. Analyst's Take

This is not the first wave of decentralized, peer-to-peer storage solutions that we have seen but, unlike in the past, this is the first time we have all the conditions necessary for success. They include the right monetization mechanisms, security and a more efficient way to store data that also leads to a greener approach to data storage. In fact, everyone is looking for better infrastructure efficiency and solutions to save money. Decentralized storage can provide all of this, along with the most rigorous security standards.

This market segment is really new and solutions are different in terms of both features and target use cases. Companies like Storj and Filebase can be more attractive for small to medium enterprise organizations while others, like Filecoin, are more interesting for larger organizations, even with the challenges posed by dynamic pricing and SLAs managed on a per-node basis. Sia and Skynet offer a solution that might find a good fit in the media industry, for example, while 0Chain's architecture looks appealing for active data.

It will be very interesting to keep an eye on all of these solutions in the next 18 to 24 months. Most are facing new challenges, including support for NFTs, demand for increased scalability and performance, and new features aimed at simplifying adoption.

9 About Enrico Signoretti

Enrico 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.

10. 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.

11. Copyright