STORJ

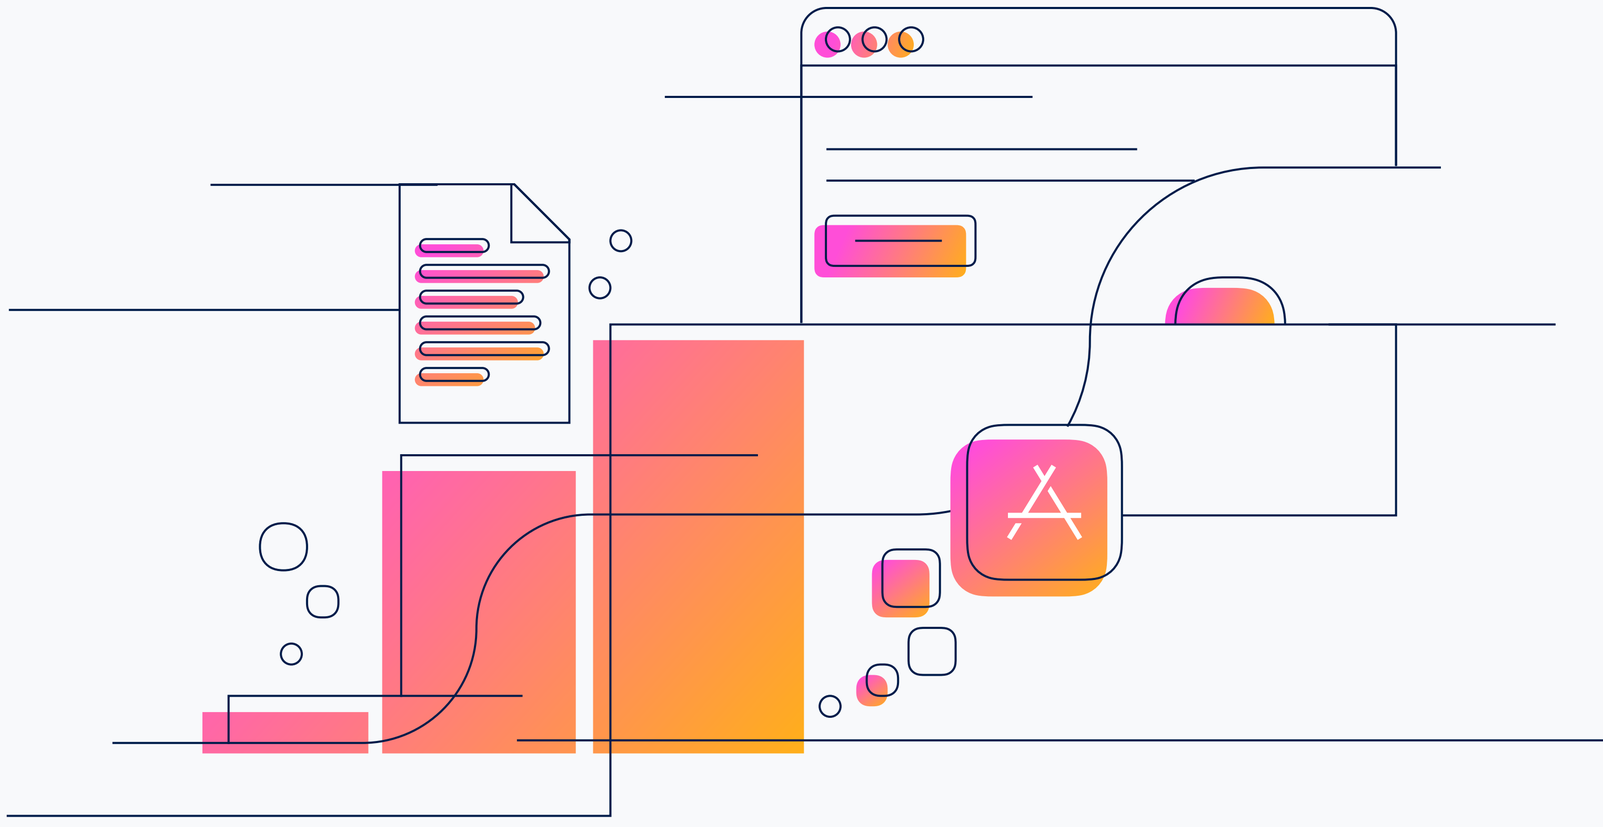

The VAR’s Guide to Distributed Cloud Storage.

The problem:

Supporting growing customer cloud storage needs.

Businesses are amassing data at unprecedented rates. According to Statista, the volume of data created, captured, copied, and consumed worldwide will reach 181 zettabytes in 2025. That’s zettabytes—and some estimate even more.

As a value-added reseller (VAR), you sell a lot of storage. And while you have a lot of products to choose from, all of those options are just variations of the same thing. Cloud object storage hasn't improved architecturally since its inception in 2006, but the world has changed—dramatically. Consumers expect instant response, applications have become extremely complex, bad actors have gotten organized, data protection laws are in effect, and eCommerce is exploding. Legacy cloud storage models simply can’t keep up.

“We expect the data universe to grow more than 10 times from 2020 to 2030, reaching 660 zettabytes—equivalent to 610 iPhones (128GB) per person.”

How the data universe could grow more than 10 times from 2020 to 2030
UBS, July 2023

And companies are starting to push back. A late 2023 Information Week article cited a survey where 94% of respondents said their cloud storage costs were escalating, and 54% said they were rising faster than other cloud costs.

You can solve this problem for them.

Traditional storage is not distributed.

There are a lot of cloud storage vendors who offer “distributed” storage solutions. What they mean is that they can copy complete data sets to several different data centers so it’s safe in case one data center goes down and is more quickly available to regional users. This is duplicated storage, not distributed storage.

Truly distributed cloud object storage encrypts, splits, and distributes data across tens of thousands of existing locations worldwide. Only a small portion of those pieces are needed to reconstitute the file. This provides incomparable durability, ensuring that your customers’ data is accessible anytime and anywhere—even when an outage or disaster occurs.

Your customers get improved performance, durability, scalability and a lower monthly bill. You get a differentiated cloud storage solution that is easy to sell with clear customer ROI.

The distributed model brings quantifiable value to your customers.

- Dramatically reduces costs by taking advantage of existing excess capacity.
- Eliminates capital expense (and carbon footprint) of data center hardware and maintenance.
- Lowers the risk of data loss from data center outages.
- Enables more consistent and faster data transfers from anywhere in the world.

Distributed storage gives you a truly differentiated offering to sell.

- Improve your profit margins while reducing your customers’ cost.
- Expand your product portfolio with a low-risk, high-return storage alternative.
- Differentiate yourself by providing disruptive technology.
- Ultimately build customer loyalty, trust, and retention.

Why you need to offer distributed storage for backups.

“Decentralized storage solutions have exposed the flaws in the current centralized model, emphasizing the importance of data privacy, security, and ownership. With continued development and adoption, disruption in the storage industry is on the horizon.”

A Guide to Decentralised Storage and its Relevance
LinkedIn/DroomDroom, April, 2023

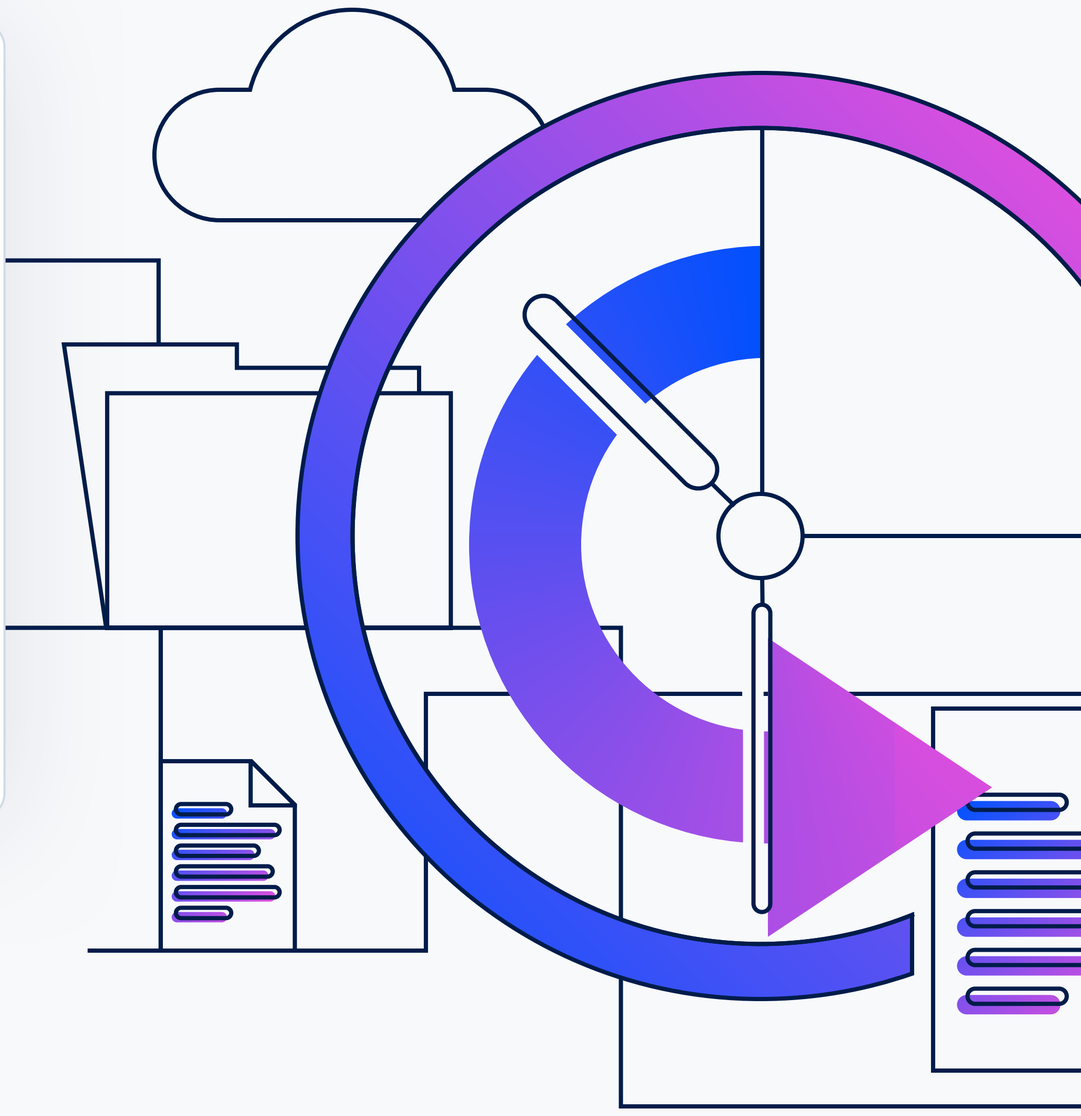

There are clear advantages to a truly distributed storage model, making it easy for you to position and sell. Let’s take a closer look.

Advantage 1

Cost

You can save your customer’s money while significantly improving your margins. And, you can sell more products by helping your customers reallocate their savings.

Legacy storage averages more than $1,000 per month for 50 TB of basic storage.

Distributed storage will save up to 90% for the same 50 TB.

For that much lower price, distributed storage includes enhanced security, durability, and global performance. The data is distributed across excess capacity on existing hardware, eliminating data center costs. Which also dramatically improves your customers’ carbon footprint, supporting their environmental sustainability initiatives (more on that later).

Advantage 2

Performance

Provide a product that natively and automatically addresses speed and availability at significantly less cost.

Avoiding network performance and latency issues with legacy storage requires duplicating the data in different regions.

Distributed storage eliminates single points of failure and network congestion because it’s multi-region by default.

Uploads are segmented and sent to a number of storage locations at the same time. Download requests are sent in parallel to everyplace the data is stored and the fastest responses are used to reconstitute the objects, making availability incredibly fast and efficient. Unused global storage space becomes a massive global network.

Advantage 3

Management and flexibility

Offering a product that gives your customers complete control with simplified management is a competitive differentiator.

Legacy storage provides zero flexibility and little control—your customers are locked in. Proprietary services and APIs make it difficult to move data and applications to another provider.

Distributed storage gives your customers complete flexibility and control.

With S3 compatibility, it’s easily used with backup and data protection software. It automatically leverages the most appropriate space available at any time, regardless of provider, reducing your customers’ operational complexity. Your customer has exclusive access to their data and there are no proprietary APIs to contend with. Distributed storage systems can work seamlessly with many existing data protection processes and tools as well, so there’s no learning curve.

Advantage 4

Security & risk management

With the enhanced security, durability, and availability of distributed storage, you can improve customer value and long-term retention.

"Most ransomware attacks target unstructured datasets, making centralized storage solutions an attractive target for encryption and/or data exfiltration of large amounts of data. Traditional storage systems are not equipped to prevent data exfiltration or manipulation, as they rely on solutions outside of the storage domain."

Gartner Magic Quadrant for Distributed File Systems and Object Storage.
November 2023

Storing complete datasets in one place dramatically increases your customers' threat surface, and even temporary outages can be catastrophic.

Distributed storage dramatically reduces your customers' threat surface by encrypting the data, splitting data into segments, automatically adding a replication factor, then distributing these segments across storage nodes around the world.

This distributed model significantly lowers the risk from all kinds of disasters. Very few of these pieces are needed to reconstitute the data and the system can retrieve those pieces from the fastest nodes, providing higher performance no matter where you are in the world.

It also enhances data security and privacy by encrypting data and metadata in transit and at rest. Advanced security features, such as identity management and multi-factor authentication, boost defenses against unauthorized access and cyber threats. With distributed storage, even the storage provider has no access to the data.

Distributed storage improves data durability as well. Because data is broken up and stored in many locations, it's more reliably and consistently preserved over time, regardless of hardware

Advantage 5

Compliance management

Complying with data protection regulations increases legacy cloud storage costs as this often means additional storage fees for each region and transfer fees between regions, plus a lot of extra configuration and testing work for security teams.

Helping your customers simplify their compliance effort is a compelling selling point to the security team.

Distributed storage meets backup rule redundancy and data protection regulations without added cost.

Mitigating the risk of data loss or halted business operations is expensive, regardless of which regulatory jurisdictions your customers operate in. Distributed storage automatically does the replication and distribution to tens of thousands of storage locations. This is more durable than what legacy storage calls “multi-region” and it comes standard. No added configuration and no added cost.

Future Advantage

Sustainability

Selling distributed data storage aligns you with your current customers’ Environmental, Social, and Governance (ESG) initiatives and will become more of a differentiator as pressure to reduce carbon emissions rises.

A McKinsey report on data center inefficiency states that most data centers are using only an average of 30% of their hard disk capacity because they need to build for future demand. Distributed storage leverages this existing underutilized disk space. This model drives more effective use of disks that are already spinning. Additionally, distributed storage can leverage existing hardware for much longer as there is less risk of data loss from older drives. Research shows that a major creator of carbon emissions is the mining and manufacturing of hard drives. Using existing drives for longer avoids 83% of the carbon emissions of cloud storage.

Across Europe, regulations that require businesses to report on their carbon footprint are going into effect in 2025. IT leaders in the US know that these regulations will eventually include them and are looking at ways to easily track carbon emissions. They also realize that once

these emissions are reported, there will be incredible pressure to reduce them.

Distributed storage actually reports carbon emissions as well as carbon avoided for every user on every invoice. Forward-thinking IT leaders will appreciate this benefit as well as the ability to show corporate leadership how much carbon they have proactively avoided.

The Storj advantage for VARs.

Storj is the only enterprise-grade distributed cloud object storage provider.

Storj leverages zero-trust/zero-knowledge fundamentals and provides end-to-end customer-side data and metadata encryption. Together, these models enable full security. Unlike legacy storage vendors, your customers have complete control of your data, no one else has access—not even the Storj team.

Based on recognized industry standards, Storj is completely AWS S3 compatible, making transferring or duplicating data (as a test) as simple as 15 minutes and a few lines of code. It provides integrations for 20 data management and recovery solutions, including Acronis, Veeam, Splunk, and MSP360, as well as dozens of other technology integrations.

Storj eliminates the extra costs of multi-region replication and storage for backups. Combined with the already lower cost of storage, this drives incredible savings for customers.

Add distributed storage to your portfolio.

As a VAR, how do you sell an increasingly costly “commodity” in a way that drives value for customers? With legacy storage, you’re getting squeezed by increasing costs and narrower margins. This model is already starting to break.

Leading VARs focus on delivering customer value with their product portfolios. Leveraging innovative partners enables you to deliver more value at less cost. Storj is the visionary in distributed object cloud storage for backups. Storj has a channel first approach to sales and actively works to achieve big wins for partners. It’s simple to register deals and get the information you need to convey the value of distributed storage to your customers.