



If your data's not available, you're out of business.

Data availability and reliability are the lifelines of every business. How do you ensure your critical data will be there when you need it?



Traditional solutions can be limited.

Today, many organizations rely on a centralized and expensive cloud storage solution to store their data; the traditional footprint includes a single-location data center filled with racks of servers with arrays of hard drives. There are a variety of methods to configure the building, the racks, and the servers to be highly available, reliable, and durable, but vulnerabilities are always present. If there is an unforeseen power or network outage or something as simple as a misconfigured router, and there isn't another copy of the data or another location housing it, access to your data is gone. In reality, such events are rare – but they can happen – and then you're out of luck.



What if I increase my storage locations?

An easy fix would seem to be replicating your data across multiple regions and physical facilities. This spreads the risk of lack of availability much more broadly; the probability that multiple locations will go down at the same time is extremely minimal. However, in achieving this level of availability, you've now multiplied your storage costs – which also increases your management complexity and the associated security footprint. There's a better way to ensure the availability and reliability of your data cost-efficiently.



Here's Where Storj DCS Comes In

Storj DCS was built with an innovative architecture and array of behind-the-scenes tools to help ensure the availability and reliability of your most important asset. This layer of managed services comes standard with every object stored at no additional cost.

When a file is uploaded to Storj DCS, the file is split up into 80 or more pieces and distributed over a range of 11,000+ diverse Storj Nodes worldwide. Our metadata hubs are geographically distributed and facilitate a set of services that handle access management, metadata management, Storj Node reputation audits, and any data repair to maximize integrity and availability.

Additional Peace of Mind With File Repair

Hard drives don't last forever, and we expect and plan for Nodes to fail, as well as voluntarily leave the network. With Storj DCS, file availability isn't impacted by the loss of one piece of a distributed file because of a built-in file repair function that rebuilds missing pieces and then distributes those new pieces for storage on healthy Storj Nodes.

File repair works in conjunction with the audit and uptime checker functions that constantly sample the network to monitor the health and availability of Storj Nodes. If through the failure, loss, or unavailability of any Nodes, the number of available pieces of any object reaches the repair threshold, the file repair function automatically downloads any 29 pieces of that object, re-encodes the object to regenerate the missing pieces, then uploads the replacement pieces to healthy Nodes. This functionality ensures sufficient pieces are available to guarantee the availability of the object.

Storj DCS Maximizes Availability & Reliability

Globally Distributed

- ▶ Multi-region by default
- ▶ No single point of failure
- ▶ Strong consistency model

Automated & Validated Assurance

- ▶ Available globally anytime, anywhere
- ▶ Self-healing storage architecture
- ▶ 99.95% availability
- ▶ 99.999999999% durability



Start building on the decentralized cloud.

www.storj.io

 @storjproject

 github.com/Storj

 storj.io/blog

© 2021 Storj Inc.