

[CASE STUDY]



STORJ

How Storj delivers global storage solutions with unlimited scale





How Storj is using CockroachDB to dramatically scale global cloud storage

INDUSTRY :

SaaS

CHALLENGES :

Building a cloud object storage system that can scale to trillions of objects.

SOLUTION :

A horizontally scalable, highly available solution built on CockroachDB.

About Storj

What if there was an Airbnb for cloud storage? There are millions of storage drives online around the world, and most are less than half full. – what if that space could be utilized? It would cost a lot less than what AWS and GCP charge for cloud storage, and with data encrypted and split across drives, it could offer strong security and performance. Sounds great, right?

That's exactly why [Storj](#) was created. The company started 10 years ago, focused on making cloud object storage better, faster, more secure and sustainable. Since then, Storj has found product market fit focused on the enterprise. Now they are scaling-up dramatically. Storj is building on top of CockroachDB, to achieve global, horizontal scale.

Watch : [How Storj is using CockroachDB to dramatically scale global cloud storage](#)

“ ”

“Other databases claimed to deliver high availability, consistency, and great performance. But if you look under the hood and if you really test them, they simply don’t have it. We looked deeply at what CockroachDB was doing and at the documentation on the consistency model. We knew that we could use it and trust it because there were experts behind it building on a scientific foundation.”

– **Jacob Willoughby**
CTO, Storj



Scaling cloud storage

Storj started building a prototype of their product on PostgreSQL since it's a battle-tested relational database that many developers are familiar with.

However, they quickly realized that this new object storage system was going to need to scale to billions – if not trillions – of objects, and weren't sure that PostgreSQL could accommodate this type of growth while maintaining availability.

High availability was also a priority for Storj, in order to meet their customer SLAs around downtime. Additionally, they wanted to ensure that transactions would always be consistent regardless of where they scaled.

In summary, their database requirements included:

- ✓ Horizontal scale for trillions of objects and exabytes of capacity
- ✓ Performant and able to retrieve global metadata
- ✓ Strong transactional consistency guarantees
- ✓ A reliable foundation with built-in high-availability

Storj also wanted to use a managed solution and a database that would make the transition off of PostgreSQL easy for developers. They did due diligence, evaluating several databases and, CockroachDB was the only solution that met all their requirements.

5.1

BILLION ROWS

9


CLOUD REGIONS

5.5

TB OF STORAGE



In 2019 Storj evaluated several databases to see if they would meet all of their requirements.

	Traditional RDBMS (Postgres)	Sharded RDBMS (Citus)	NoSQL (Cassandra)	Spanner	 CockroachDB	Yugabyte + YSQL	TiDB
Horizontally scalable	NO	KINDA	YES	YES	YES	YES	YES
Good performance	YES	YES	YES	YES	YES	YES	YES
Strong consistency	Failover?	Failover?	NO	YES	YES	YES	YES
PG wire compatible	YES	YES	NO	NEW	YES	YES	MySQL
Open source / on prem	YES	YES	YES	NO	YES	YES	YES
Established history	YES	YES	YES	YES	MATURING	NEW	NEW

A multi-continent solution

In 2019, Storj migrated from PostgreSQL to CockroachDB and started to use that as their primary platform. Because CockroachDB is PostgreSQL wire-compatible, the transition was easy, allowing the team to focus on building the product instead of learning a new technology.

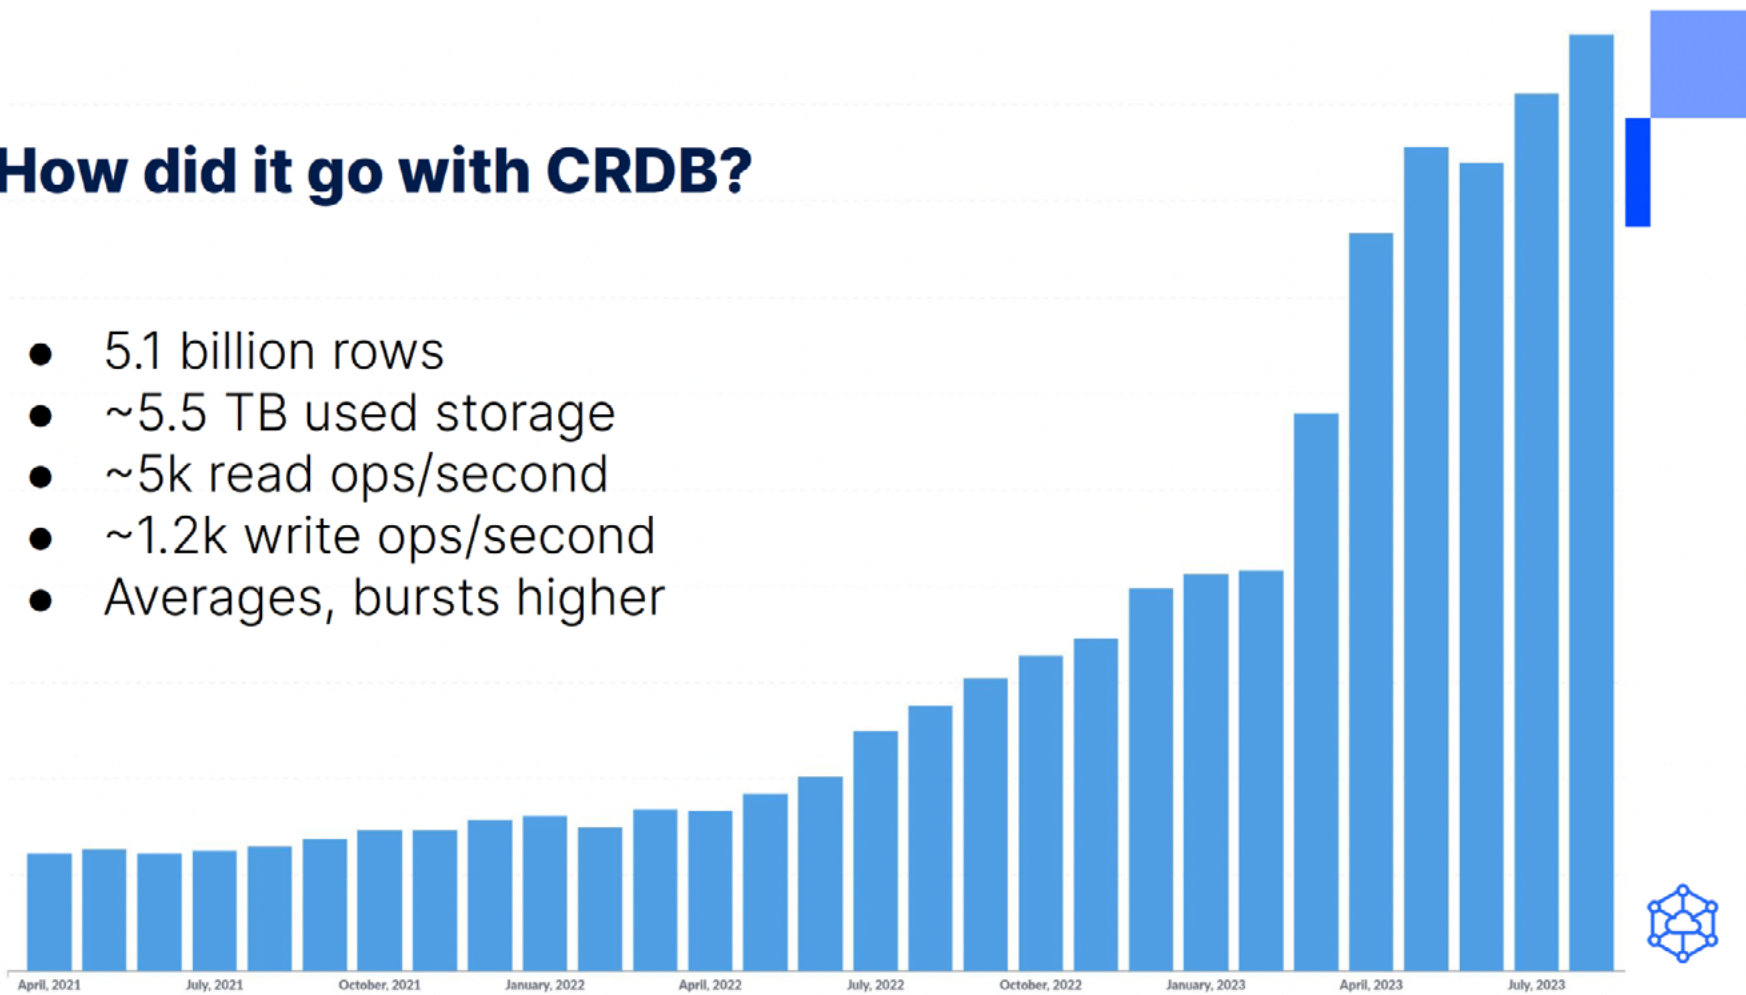
Storj uses CockroachDB for several use cases, including billing, account management, and metadata storage. Storj's object storage metadata is a key component of their system and requires tremendous scale.



The chart below demonstrates how Storj scaled one of their CockroachDB clusters to 5.1 billion rows, using ~5.5 TB of storage. This workload averages around 5K read operations per second and 1.2K write operations per second but gets higher during peak periods.

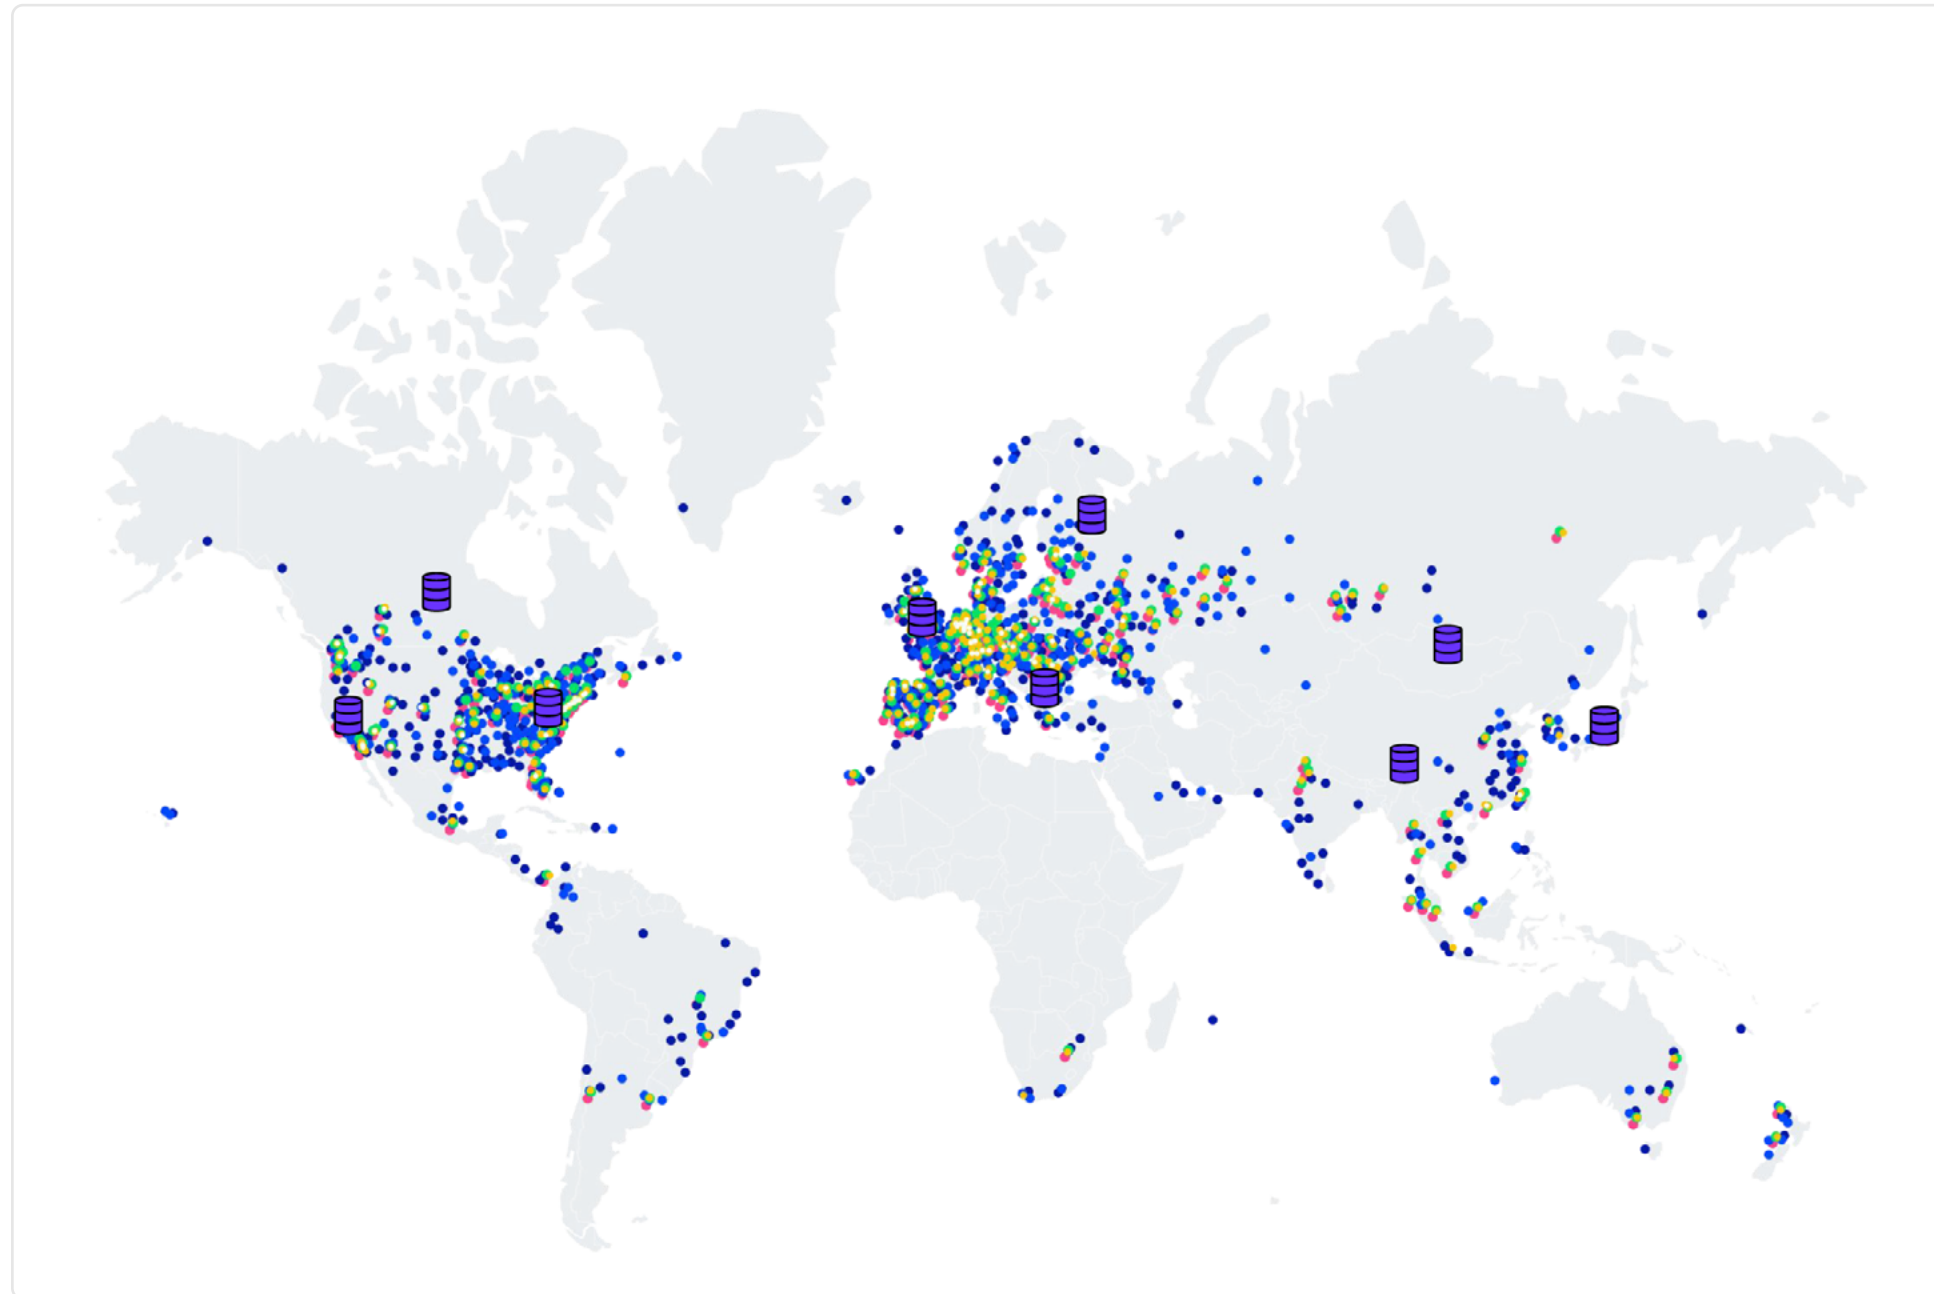
How did it go with CRDB?

- 5.1 billion rows
- ~5.5 TB used storage
- ~5k read ops/second
- ~1.2k write ops/second
- Averages, bursts higher



In terms of a deployment model, they use a Kubernetes engine on Google Cloud and also have CockroachDB deployed in the same regions where they run their API servers. In total, Storj is running in 9 regions. They have multi-region CockroachDB clusters in 3 different continents (North America, Europe, and Asia) with 3 regions in each of those continents.

This deployment is solely based on customer demand, as seen in the diagram below.



Since Storj uses a managed version of CockroachDB, spinning up a new region is easy. It's a "matter of pushing a few buttons in the web UI, and we have a new cluster up and running in minutes."

Looking forward, Storj plans to build a globally performant CockroachDB cluster that spans multiple continents, knowing they have the option to run across multiple clouds.


Using the best tools

There are a few best practices Jacob (CTO) recommends when considering CockroachDB. He says you shouldn't try to prematurely optimize for scaling because it will create a lot of unnecessary work." This is different from other legacy relational systems where capacity planning is an important factor.

For example, a few years ago, they did not have 5.1 billion rows. But when they were ready to scale the load they were able to "just make some easy adjustments to the management console to add additional capacity to the database".

Finally, Jacob says that Cockroach Labs' premium support is essential for Storj. He says that whenever they've contacted support, the Cockroach Labs team quickly dove in and provided immediate support in order to resolve questions or potential issues.

Having this level of engagement is "very comforting to us and essential to build the availability we need."



What's Next

Storj has plans to expand to almost every continent with available clusters. Today, they operate around 23K+ active nodes and are helping customers cut storage costs and carbon footprints by over 80%.

If you are interested in checking out Storj's solutions, visit their [website](#).

Ready to get started?

Go hands-on with 100% free CockroachDB Serverless. Spin up your first cluster in just a few clicks.

“”

“If you can simply make the right choice for a database early on, you won't experience negative impacts when building the product. With CockroachDB, we were able to develop and build just as fast as we would've on top of PostgreSQL or any other database. And when we needed it, we scaled dramatically.”

– **Jacob Willoughby**
CTO, Storj

