



How CockroachDB is helping us do the impossible with global cloud storage

Jacob Willoughby, CTO

What do I mean by impossible?



Agenda

- Who we are and what we do
- Our database requirements
- How does CockroachDB stack up?
- Success stories
- Wrap up



Agenda

- Who we are and what we do**
- Our database requirements
- How does CockroachDB stack up?
- Success stories
- Wrap up

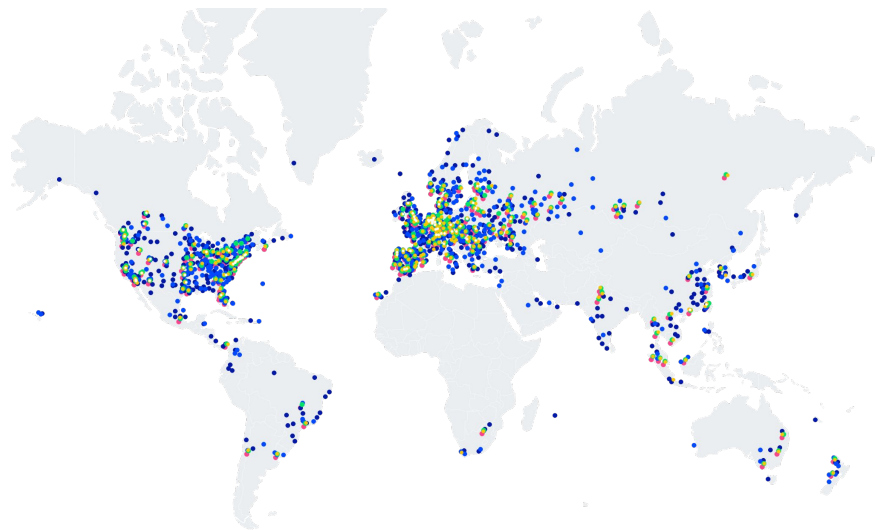


A new model for cloud storage



Storj object storage is fundamentally different

- ✓ Unlocks stranded capacity across the enterprise
- ✓ Security/privacy-first architecture
- ✓ Works with existing heterogeneous hardware
- ✓ High durability, availability
- ✓ CDN-like performance
- ✓ Minimal carbon footprint
- ✓ Simple, scalable pricing



Storj has 3 key elements

Nodes

Supply

Tens of thousands of shared hard drives store data on the network, without access to any complete file or usable data

Applications

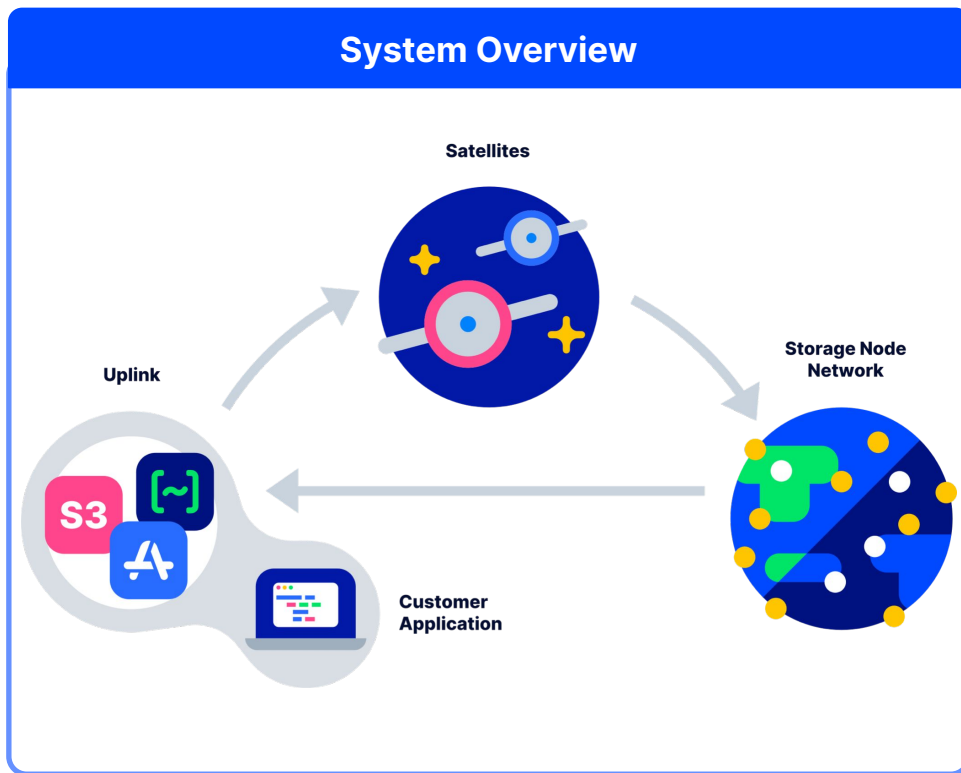
Demand

Client applications store encrypted and encoded files split into pieces and stored across a distributed network

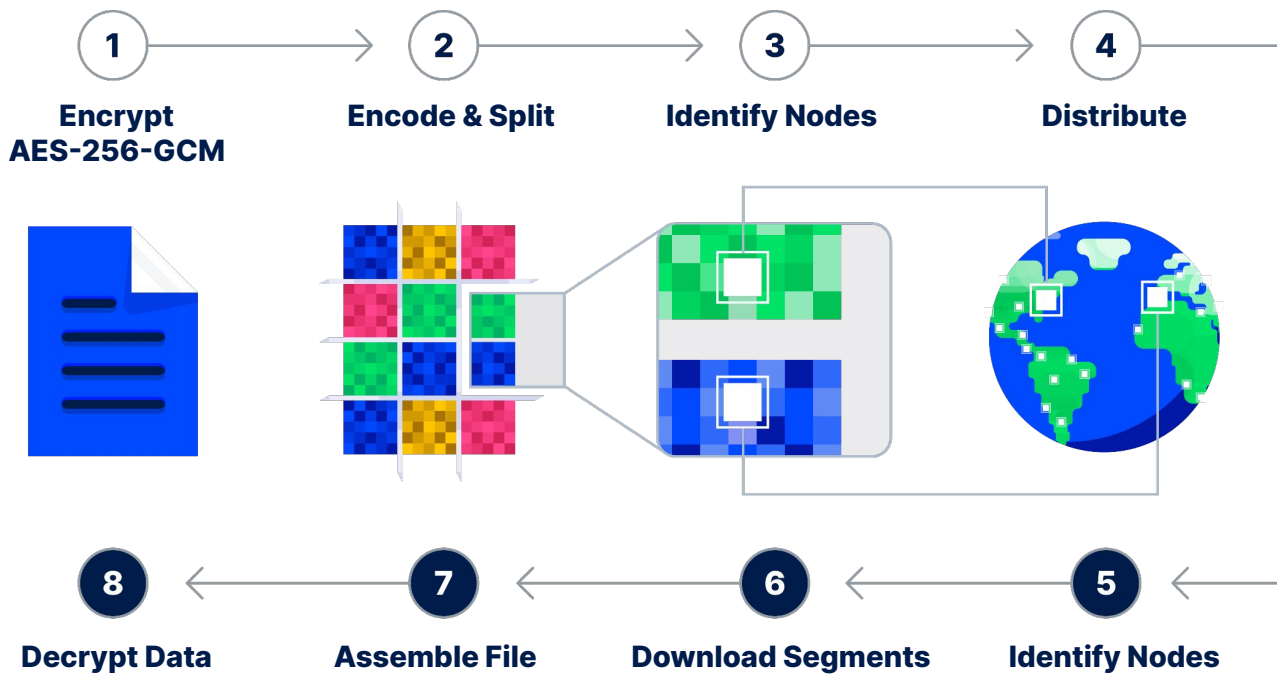
Satellites

Network

The network enables applications to store data, ensures data reliability, manages access controls and data audit/repair



Storj uses encryption and erasure encoding to distribute data keeping it durable and available



Live Upload Demo

- Upload every 6 seconds
- See locations of where files are stored in real-time!

<https://demo.storj.dev/>



Agenda

- | Who we are and what we do
- | **Our database requirements**
- | How does CockroachDB stack up?
- | Success stories
- | Wrap up





Designing our metadata orchestrator: **early goals**



**Horizontal
scalability**



Performance



**Strong
consistency**



Designing our metadata orchestrator: What ended up making the decision

**Postgres wire
compatibility**

**Open source and
on prem option**

**How established
is it?**




Agenda

- Who we are and what we do
- Our database requirements
- How does CockroachDB stack up?**
- Success stories
- Wrap up



How did CockroachDB stack up (as of 2019)?

	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



Agenda

- | Who we are and what we do
- | Our database requirements
- | How does CockroachDB stack up?
- | **Success stories**
- | Wrap up



Success stories


storage: Make Put handle WriteTooOldError more like CPut #38668

 Merged **craig** merged 2 commits into `cockroachdb:master` from `bdarnell:defer-write-too-old`  on Jul 9, 2019

 Conversation (18)

 Commits (2)

 Checks (0)

 Files changed (18)



bdarnell commented on Jul 3, 2019

Member



Reviewers



andreima



knz



nvanbens

kvserver: deadlock in protectedts error message #66759

 Closed

tbg opened this issue on Jul 14, 2021



tbg mentioned this issue on Jul 14, 2021

sql/kv: implement lock_timeout session variable #67513

 Closed



tbg mentioned this issue on Mar 16, 2022

kvserver: add generalized circuit-breaking to catch, e.g., mutex deadlocks #77366

 Open



erikgrinaker mentioned this issue on Jan 18

kvserver: Raft stalls during remote node startup #95433

 Closed



tbg commented on Jun 23, 2022

Describe the problem

[#61610](#) introduced a deadlock while already holding (*Replica)

Success stories



Success stories



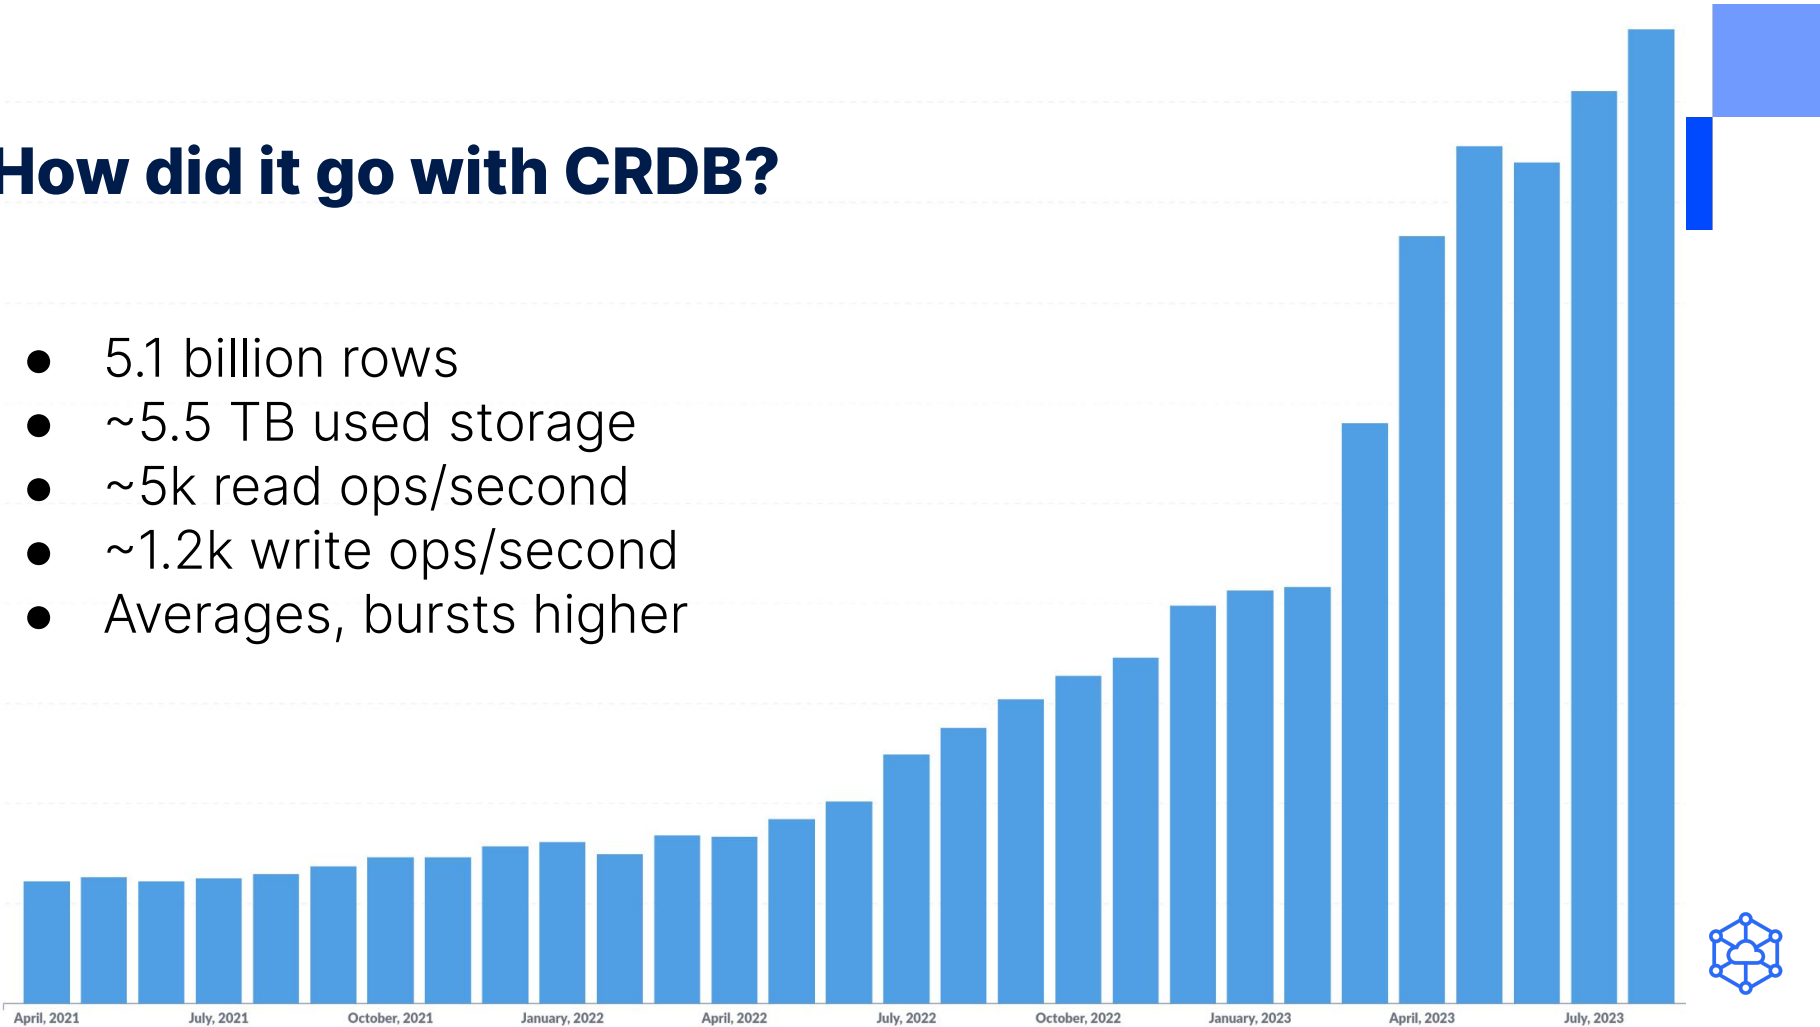
Agenda

- | Who we are and what we do
- | Our database requirements
- | How does CockroachDB stack up?
- | Success stories
- | **Wrap up**

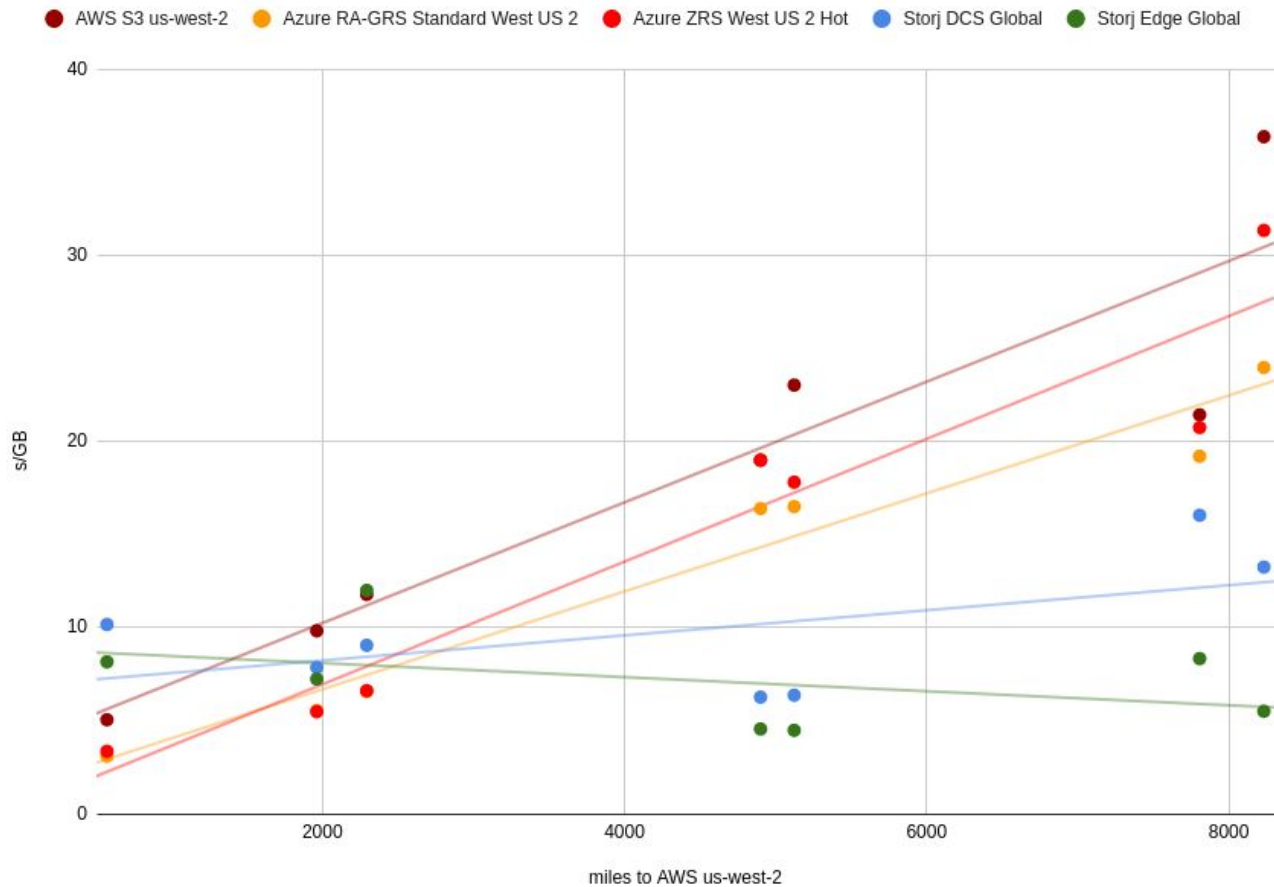


How did it go with CRDB?

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



This IS the impossible





CockroachDB met and exceeded our early goals

**Horizontal
scalability**

Performance

**Strong
consistency**





CockroachDB met and exceeded our early goals



**Horizontal
scalability**

Performance

**Strong
consistency**





CockroachDB met and exceeded our early goals



**Horizontal
scalability**



Performance

**Strong
consistency**





CockroachDB met and exceeded our early goals



Horizontal scalability




Performance



Strong consistency





Final thoughts





STORJ

