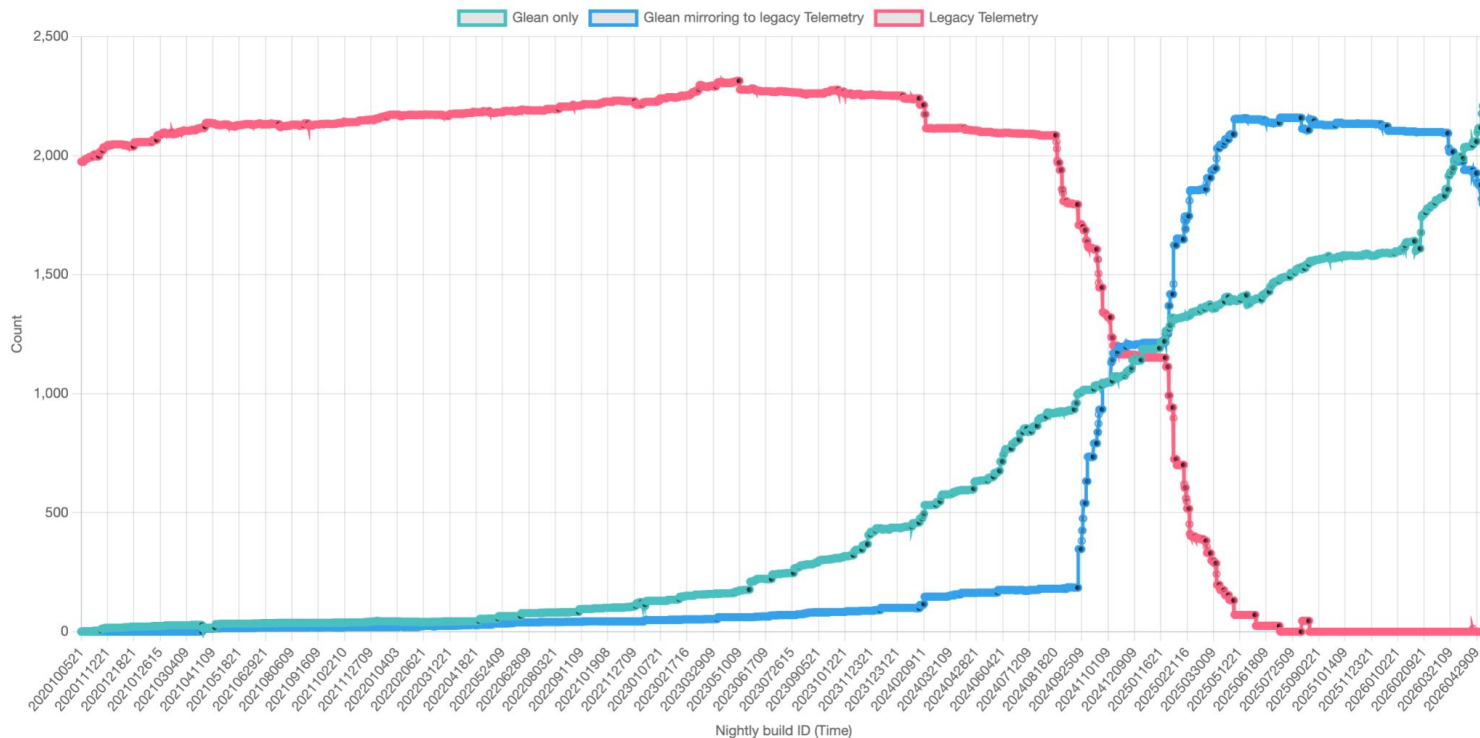


# A new storage backend for Glean

Jan-Erik Rediger – Glean Team  
Thundertalks – 2026-05-13



# Glean is now everywhere



<https://arewegleanyet.com>

Glean promises:  
recorded data will be stored and send eventually

`Glean.tabs.count.add(1)`



Glean SDK



[a file on disk]

## 2019 - 2022: Glean's database: rkv → LMDB

rkv: A simple, humane, typed key-value storage solution

LMDB: Lightning Memory-Mapped Database

# 2022-2026: Glean's database: rkv safe-mode

rkv: A simple, humane, typed key-value storage solution

safe-mode: serialize a hashmap to disk

# rkv safe-mode issues

- No data consistency guarantees
- Write amplification
- Very bare-bones key-value

We don't want to be database engineers\*!

\* maybe secretly we want to be, but that's not what we get paid for

Let's use a real database



## 2026-\*: Glean's database: SQLite

- A database for database things
- Better data consistency guarantees
- Vast production experience at Mozilla

## 2026-\*: Glean's database: SQLite

- Early micro benchmarks: 1.2x - 5x faster than rkv
- Slightly higher memory usage for the same dataset
- +10-13% disk storage
  - 501 kb vs 568k

2026-\*: Glean's database: SQLite



154

tentative, timeline might change

## 2026-\*: Glean's database: SQLite

- Transparent for all consumers
- Automatic migration of data
- Possibility for a rollback if issues arise during rollout

