




[← Back](#)

Blog post

# Implementing "seen by" functionality with Postgres

2022-07-18 • 33 minute read

Victor  
Guest Author



**tl;dr: Use HyperLogLog, it's a reasonable approach with great trade-offs and no large architectural liabilities. For a quick & dirty prototype, use `hstore`, which also performs the best with integer IDs.**

The year is 2022. You're head DBA at the hot new social site, SupaBook... Your startup is seeing eye-boggling growth because everyone loves fitting their hot-takes in posts restricted to `VARCHAR(256)`.

Why `VARCHAR(256)`? No particular reason, but you don't have time to get hung up on that or ask why -- **you just found out that the priority this quarter is tracking content views across all posts in the app.**

"It sounds pretty simple" a colleague at the meeting remarks -- "just an increment here and an increment there and we'll know which posts are seen the most on our platform". You start to explain why it will be non-trivial, but the meeting ends before you can finish.

Well, it's time to figure out how you're going to do it. There's been a complexity freeze at the company, so you're not allowed to bring in any new technology, but you don't mind that because for v1 you would have picked Postgres anyway. Postgres's open source pedigree, robust suite of features, stable internals, and awesome mascot Slonik make it a strong choice, and it's what you're already running.

*(insert record scratch here)*

Sure, this scenario isn't real, but it could be - that last part about Postgres definitely is. Let's see how you might solve this problem, as that imaginary DBA.

## Experiment setup

We've got the following simple table layout:

In SQL migration form:

```
hideCopy 📄
1 CREATE EXTENSION IF NOT EXISTS uuid-oss;
2 CREATE EXTENSION IF NOT EXISTS citext;
3
4 -- Create a email domain to represent and con
5 CREATE DOMAIN email
6 AS citext
7 CHECK ( LENGTH(VALUE) <= 255 AND value ~ '^[a
8
9 COMMENT ON DOMAIN email is 'lightly validated
10
11 -- Create the users table
12 CREATE TABLE users (
13     id bigserial PRIMARY KEY GENERATED BY DEF
14     uuid uuid NOT NULL DEFAULT uuid_nonmc_v1(
15
16     email email NOT NULL,
17     name text,
18     about_html text,
19
20     created_at timestamptz NOT NULL DEFAULT M
21 );
22
23 -- Create the posts table
24 CREATE TABLE posts (
25     id bigserial PRIMARY KEY GENERATED BY DEF
26     uuid uuid NOT NULL DEFAULT uuid_nonmc_v1(
27
28     title text,
29     content text,
30     main_image_src text,
31     main_link_src text,
32
33     created_by bigint REFERENCES users(id),
34
35     last_hidden_at timestamptz,
36     last_updated_at timestamptz,
37     created_at timestamptz NOT NULL DEFAULT M
38 );
```

This basic setup has taken the (imaginary) company quite far -- even though the `posts` table has millions and millions of entries, Postgres chugs along and serves our queries with impressive speed and reliability. Scaling up is the new (and old) scaling out.

## How should we do it?

Well we can't pat ourselves for our miraculous and suspiciously simple DB architecture all day, let's move on to the task at hand.

Like any good tinkerer we'll start with the simplest solutions and work our way up in complexity to try and get to something outstanding, testing our numbers as we go.

## Try #1: The naive way, a simple counter on every Post

The easiest obvious way to do this is to maintain a counter on every tuple in the `posts` table. It's obvious, and it's almost guaranteed to work -- but maybe not *work well*.

The migration to make it happen isn't too difficult:

hideCopy

```
1 BEGIN;
2
3 ALTER TABLE posts ADD COLUMN seen_by_count;
4
5 COMMENT ON COLUMN posts.seen_by_count
6 IS 'simple count of users who have seen the
7
8 COMMIT;
```

There's one *obvious* glaring issue here -- what if someone sees the same post twice? Every page reload would cause inflated counts in the `seen_by_count` column, not to mention a lot of concurrent database updates (which isn't necessarily Postgres's forte to begin with).

Clearly there's a better way to do things but before that...

# Writing a test suite before the CPUs get hot and heavy

How will we know which approach is better without numbers?! Measuring complexity and feeling can only get us so far -- we need to get some numbers that tell us the performance of the solution at the stated tasks - - we need benchmarks.

Before we can declare any solution the best, in particular we need a *baseline!*. The simplest possible incorrect solution (simply incrementing a counter on the Post) is probably a reasonable thing to use as a benchmark, so let's take a moment to write our testing suite.

Let's do this the simplest one might imagine:

Generate a large amount of users

Lets model for 1000, 10k, 100K, 1MM, and 10MM users

Generate an even larger amount of fake posts attributed to those users

This is a bit harder -- we need to define a general distribution for our users that's somewhat informed by real life...

An average/normalized distribution doesn't quite work here -- on sites like twitter 10% of users create 80% of the tweets!

Generate a *description* of "events" that describe which post was seen by whom, which we can replay.

We want the equivalent of an effect system or monadic computation, which is easier than it sounds -- we want to generate an encoding (JSON, probably) of *what to do*, without actually doing it

We'll just do consistent "as fast as we can" execution (more complicated analysis would burst traffic to be ab it closer to real life)

OK, let's roll our hands up and get it done:

## Script: User seeding

Here's what that looks like:

```
1  /**
2   * Generate a list of synthetic users to be l
3   *
4   * @param {object} args
5   * @param {number} [args.count] number of use
6   * @param {number} [args.aboutHTMLWordCount]
7   * @param {string} [args.outputFilePath] outp
8   * @returns {any[][]} List of generated synth
9   */
10 export async function generateUsers(args) {
11   const count = args.count || DEFAULT_USER_CC
12   const aboutHTMLWordCount = args.aboutHTMLWc
13
14   const outputFilePath = args.outputFilePath
15   if (!outputFilePath) {
16     throw new Error('output file path must be
17   }
18
19   for (var id = 0; id < count; id++) {
20     const user = {
21       id,
22       email: `user${id}@example.com`,
23       name: `user ${id}`,
24       about_html: fastLoremIpsum(aboutHTMLWor
25     }
26
27     // Write the entries to disk (returning r
28     if (args.outputFilePath) {
29       await appendFile(outputFilePath, `${JSC
30     }
31   }
32 }
```

Nothing too crazy in there -- we generate a bunch of JSON, and force it out to disk. It's best to avoid trying to keep it in memory so we can handle much larger volumes than we might be able to fit in memory.

If you'd like to see the code, check out [scripts/generate/users.js](#) in the repo.

## Script: Post seeding

Along with users, we need to generate posts that they can view. We'll keep it simple and take an amount of posts to make, generating from 0 to `count` of those.

It's very similar to the user generation code, with the caveat that we can take into account the 80/20 lurker/poster rule. here's what that looks like:

It's a bit long so if you'd like to see the code, check out [scripts/generate/posts.js](#) in the repo.

## Script: action (API call) seeding/generation

This script is a bit tricky -- we need to inject some randomness in the performing of the following actions:

- Record a new view of a post
- Retrieve just the count of a single post
- Retrieve all the users who saw a post

I've chosen to use [autocannon](#) so I needed to write a request generation script which looks like this:

```
1  const process = require('process')
2
3  const POST_COUNT = process.env.TEST_POST_COUNT
4    ? parseInt(process.env.TEST_POST_COUNT, 10)
5    : undefined
6  const USER_COUNT = process.env.TEST_USER_COUNT
7    ? parseInt(process.env.TEST_USER_COUNT, 10)
8    : undefined
9
10 /**
11  * Request setup function for use with autocannon
12  *
13  * @param {Request} request
14  * @returns {Request}
15  */
16 function setupRequest(request) {
```



```

6 function setupRequest(request) {
17 // ENSure we have counts to go off of
18 if (!POST_COUNT || !USER_COUNT) {
19   throw new Error('Cannot setup request with counts')
20 }
21
22 // Pick a random post to do an operation on
23 const postId = Math.floor(Math.random() * POST_COUNT)
24
25 // Choose pseudo-randomly whether to register a user
26 const operationChoice = Math.floor(Math.random() * 10)
27 if (operationChoice < 1) {
28   // 10% of the time, get *all* the users
29   request.method = 'GET'
30   request.path = `/posts/${postId}/seen-by/`
31 } else if (operationChoice < 7) {
32   // 60% of the time, get the count of seen-by users
33   request.method = 'GET'
34   request.path = `/posts/${postId}/seen-by/`
35 } else {
36   // 30% of the time, add a new seen-by entry
37   const userId = Math.floor(Math.random() * USER_COUNT)
38
39   // Most of the time we'll be *setting* seen-by counts
40   // And we'll get the count (so we can show the user)
41   request.method = 'POST'
42   request.path = `/posts/${postId}/seen-by/${userId}`
43 }
44
45 return request
46 }
47
48 module.exports = setupRequest

```

Nothing too crazy here, and some back of the envelope estimations on how often each operation would normally be called. These numbers could be tweaked more, but we *should* see a difference between approaches even if we messed up massively here.

If you'd like to see the code, check out [scripts/setup-request.cjs](#) in the repo.

## Glue it all together

Once we're done we need to glue this all together into one script, with roughly this format:

```

1 // ...

```

```

1 export default async function runBenchmark()
2   // Start the server
3   // Reset before test
4   // Generate & insert users
5   // Generate & insert posts
6   // Generate actions (API Calls) to run
7   // Execute the API calls
8   // Write JSON results to tmpdir
9   // Stop the server
10 }

```

If you want to see what the code *actually* ended up looking like, check out `scripts/bench.js` in the repo.

Along with the benchmark, we'll standardize on the following settings:

```

1 export SEEN_BY_STRATEGY=simple-counter # of:
2 export TEST_USERS_JSON_PATH=/tmp/supabase-see
3 export TEST_POSTS_JSON_PATH=/tmp/supabase-see
4 export TEST_POST_COUNT=1000
5 export TEST_USER_COUNT=100000
6 export TEST_DURATION_SECONDS=60
7
8 ## Use custom postgres image built with hll e
9 ## NOTE: `make db-custom-image` must be run b
10 #export DB_IMAGE=postgres-14.4-alpine-hll
11 #export DB_IMAGE_TAG=latest

```

## Our first run, on the naive solution

Alright, finally we're ready. Let's see what we get on our naive solution. We expect this to be *pretty fast*, because not only is it *wrong*, but it's just about the simplest thing you could do.

On my local machine, here's our baseline (output from `autocannon`):

```

1
2
3
4
5
6
7

```

Stat	2.5%	50%	97.5%	99%	Avg
Latency	0 ms	2 ms	6 ms	6 ms	2.03

```

7
8
9
10
11

```

```
7 Stat | 1% | 2.5% | 50% | 9
8
9 Req/Sec | 297 | 318 | 389 | 5
10
11 Bytes/Sec | 54.1 kB | 57.9 kB | 70.8 kB | 9
12
13
14 Req/Bytes counts sampled once per second.
15 # of samples: 60
16
17
18 Percentile | Latency (ms) |
19
20 0.001 | 0 |
21
22 0.01 | 0 |
23
24 0.1 | 0 |
25
26 1 | 0 |
27
28 2.5 | 0 |
29
30 10 | 0 |
31
32 25 | 0 |
33
34 50 | 2 |
35
36 75 | 3 |
37
38 90 | 5 |
39
40 97.5 | 6 |
41
42 99 | 6 |
43
44 99.9 | 9 |
45
46 99.99 | 16 |
47
48 99.999 | 23 |
49
50
51 23k requests in 60.02s, 4.28 MB read
```

As you might imagine, pretty darn good latency across all the requests.

### Back to trying things out


Now that we've got a basic baseline of our tests, let's continue trying out ideas:

## Try #2: Storing the users who did the "see"ing, with `hstore`

The next obvious thing (and probably a core requirement if we'd asked around), is knowing *who* viewed each post. Well if we need to know who, then we probably need to store some more information!

Postgres has native support for arrays and a data structure called a `hstore`, so let's try those. It's pretty obvious that having hundreds, thousands, or millions of entries in one of these data structures, inside a tuple isn't the *greatest* idea, but let's try it anyway and let the numbers speak for themselves.

Here's what the migration would look like:

```
hideCopy 
1 BEGIN;
2
3 CREATE EXTENSION IF NOT EXISTS hstore;
4
5 ALTER TABLE posts ADD COLUMN seen_count_hstore
6 NOT NULL DEFAULT ''::hstore;
7
8 COMMENT ON COLUMN posts.seen_count_hstore
9 IS 'count of users that have seen the post,'
10
11 COMMIT;
```

`hstore` provides support for both GIST and GIN indices, but after reading the documentation we can conclude that we don't necessarily need those for the current set of functionality.

### Caveats

Well as you might have imagined, this is obviously pretty bad and will eventually be hard to scale. If you expect only 0-50 entries in your column `text[]` is perfectly fine, but thousands or millions is another ballgame.

Thinking of how to scale this, a few ideas pop to mind:

Compress our columns with LZ4 which is newly supported TOAST column compression (I first heard about this thanks to Fujitsu's fantastic blog post)

`PARTITION` our `posts` table

### Performance

OK, time to get on with it, let's see how it performs with an `hstore`:

```
1
2  Stat      | 2.5% | 50% | 97.5% | 99% | Avg
3
4  Latency   | 0 ms | 2 ms | 5 ms  | 6 ms | 2.15
5
6
7  Stat      | 1%   | 2.5% | 50%   | 9
8
9  Req/Sec   | 287  | 305  | 348   | 5
10
11 Bytes/Sec | 53.9 kB | 56.9 kB | 64.5 kB | 9
12
13
14 Req/Bytes counts sampled once per second.
15 # of samples: 60
16
17
18 Percentile | Latency (ms) |
19
20 0.001      | 0             |
21
22 0.01       | 0             |
23
24 0.1        | 0             |
25
26 1          | 0             |
27
28 2.5       | 0             |
29
30 10        | 0             |
31
```

31		
32	25	1
33		
34	50	2
35		
36	75	3
37		
38	90	5
39		
40	97.5	5
41		
42	99	6
43		
44	99.9	9
45		
46	99.99	9
47		
48	99.999	16
49		
50		
51	22k requests in 60.02s, 4.1 MB read	

Not too far off! While we didn't try the pathological case(s) of millions of people liking the *same* post to hit breaking point, a slightly more random distribution seems to have done decently -- we actually have *lower* 99.999th percentile latency versus the simple counter.

An average of 2.15ms versus 2.05ms with the simpler counter is a ~4% increase in the average latency (though of course, the p99.999 is lower!).

### Try #3: An Association table for remembering who liked what

A likely requirement from the original scenario that we've completely ignored is remembering *which* users liked a certain post to. The easiest solution here is an "associative" table like this one:

In SQL:

```
hideCopy 📄  
  
1 begin;  
2  
3 create table posts_seen_by_users (  
4   post_id bigint references posts (id),  
5   user_id bigint references users (id),  
6   seen_count bigint not null default 0 check  
7   primary key (post_id, user_id)  
8 );  
9  
10 commit;
```

### Caveats

In production, you're going to want to do a few things to make this even remotely reasonable long term:

`PARTITION` the table (consider using partition-friendly `pg_partman`)

Move old partitions off to slower/colder storage and maintain snapshots

Summarize older content that might be seen lots

Consider a partitioning key up front -- post IDs are probably a reasonable thing to use if they're sufficiently randomly distributed

These are good initial stop-gaps, but a realistic setup will have many problems and many more solutions to be discovered.

(It will be a recurring theme but this is a spot where *we probably don't necessarily want to use stock Postgres* but instead want to use tools like [Citus Columnar Storage](#), [ZedStore](#), or an external choice like [ClickHouse](#)).

### Performance

Alright, enough dilly dally, let's run our test bench against this setup:

```
1
2
3
4
5
6
7
8
9
10
11
12
13
14
15
16
17
18
19
20
21
22
23
24
25
26
27
28
29
30
31
```

Stat	2.5%	50%	97.5%	99%	Avg
Latency	0 ms	2 ms	8 ms	8 ms	2.5

Stat	1%	2.5%	50%	99%
Req/Sec	238	254	321	400
Bytes/Sec	43.4 kB	46.3 kB	58.5 kB	80.0 kB

Req/Bytes counts sampled once per second.  
# of samples: 60

Percentile	Latency (ms)
0.001	0
0.01	0
0.1	0
1	0
2.5	0
10	0



32	25	1
33		
34	50	2
35		
36	75	4
37		
38	90	7
39		
40	97.5	8
41		
42	99	8
43		
44	99.9	11
45		
46	99.99	25
47		
48	99.999	30
49		
50		
51	20k requests in 60.02s, 3.57 MB read	

A little bit more divergence here -- 99.999%ile latency @ 30 which is almost double what it was for simple-hstore.

Average is coming in at 2.50ms which is 16% slower than simple-hstore and 21% slower than simple-counter.

**Try #4: Getting a bit more serious: bringing out the HyperLogLog**

We'll just draw the rest of the owl now.

What's HyperLogLog you ask? Well it's just a probabilistic data structure! Don't worry if you've never heard of it before, it's a reasonably advanced concept.

You may have heard of Bloom Filters and they're *somewhat* related but they're not quite a great fit for the problem we're solving since we want to know how many people have seen a particular post. Knowing whether one user has seen a particular post is useful too -- but not quite what we're solving for here (and we'd have to double-check our false positives anyway if we wanted to be absolutely sure).

HyperLogLog provides a probabilistic data structure that is good at counting *distinct* entries, so that means that the count *will not* be exact, but be reasonably close (depending on how we tune). We won't have false positives (like with a bloom filter) -- we'll have a degree of error (i.e. the actual count may be 1000, but the HLL reports 1004).

We have to take this into account on the UI side but and maybe retrieve the full count if anyone ever *really* needs to know/view individual users that have seen the content, so we can fall back to our association table there.

Given that every second there are about 6000 tweets on Twitter(!), this is probably one of the only solutions that could actually work at massive scale with the limitations we've placed on ourselves.

Here's what that looks like in SQL:

```
hideCopy 📄  
  
1 BEGIN;  
2  
3 CREATE EXTENSION IF NOT EXISTS hll;  
4  
5 ALTER TABLE posts ADD COLUMN seen_count_hll hll  
6 NOT NULL DEFAULT hll_empty();  
7  
8 COMMENT ON COLUMN posts.seen_count_hll  
9 IS 'HyperLogLog storing user IDs';  
10  
11 COMMIT;
```

Here we need the `citus/postgresql-hll` extension, which is generously made (truly) open source by [citusdata](#).

NOTE that we still have access to the association table -- and while we still insert rows into it, we can *drop* the primary key index, and simply update our HLL (and leave ourselves a note on when we last updated it).

## Caveats

There's not much to add to this solution, as the heavy lifting is mostly done by `postgresql-hll`, but there's one big caveat:

This approach *will* need a custom Postgres image for this, since `hll` is not an official `contrib` module

There are also a few optimizations that are easy to imagine:

Batching inserts to the association table (storing them in some other medium in the meantime -- local disk, redis, etc)

Writing our association table entries in a completely different storage medium altogether (like object storage) and use [Foreign Data Wrappers](#) and `pg_cron` and delay or put off processing all together

## Performance

The most complicated solution by far, let's see how it fares:

```
1
2
3
4
5
6
7
8
9
10
11
12
13
14
15
16
17
18
19
20
21
22
23
24
25
26
27
28
```

Stat	2.5%	50%	97.5%	99%	Avg
Latency	0 ms	2 ms	6 ms	6 ms	2.28

Stat	1%	2.5%	50%	9
Req/Sec	272	285	351	4
Bytes/Sec	49.5 kB	51.9 kB	63.9 kB	8

Req/Bytes counts sampled once per second.  
# of samples: 60

Percentile	Latency (ms)
0.001	0
0.01	0
0.1	0
1	0
2.5	0

28	2.5	0
29		
30	10	0
31		
32	25	1
33		
34	50	2
35		
36	75	4
37		
38	90	6
39		
40	97.5	6
41		
42	99	6
43		
44	99.9	9
45		
46	99.99	28
47		
48	99.999	59
49		
50		
51	21k requests in 60.03s, 3.86 MB read	

Another somewhat nuanced degradation in performance -- while the 99.99%ile latency was nearly 2x higher, the average latency was actually *lower* than the assoc-table approach @ `2.28ms`.

The average latency on the HLL approach is 11% worse than simple-counter, 6% worse than simple-hstore, and *faster* than assoc-table alone, which is an improvement.

## Oh, the other places we could go

One of the great things about Postgres is it's expansive ecosystem -- while Postgres may (and frankly *should not*) beat the perfect specialist tool for your use case, it often does an outstanding job in the general case.

Let's look into some more experiments that could be run -- maybe one day in the future we'll get some numbers behind these (community contributions are welcome!).

Incremental view maintenance powered by `pg_ivm`

If you haven't heard about `pg_ivm` it's an extension for handling Incremental View Maintenance -- updating `VI EW`s when underlying tables change.

IVM is a hotly requested feature whenever views (particularly materialized views) are mentioned, so there has been much fanfare to its release.

There are a couple advantages we could gain by using `pg_ivm` :

Ability to time constrain calculations (newer posts which are more likely to be seen can exist in instant-access views)

We could theoretically remove the complicated nature of the HLL all together by using `COUNT` with IVM

`pg_ivm` is quite new and cutting edge but looks to be a great solution -- it's worth giving a shot someday.

## Doing graph computations with AGE

As is usually the case in academia and practice, we can make our problem drastically easier by simply changing the data structures we use to model our problem!

One such reconfiguration would be storing the information as a graph:

As you might imagine, finding the number of "seen-by" relations would simply be counting the number of edges out of one of the nodes!

Well, the Postgres ecosystem has us covered here too! AGE is an extension that allows you to perform graph related queries in Postgres.

We won't pursue it in this post but it would be a great way to model this problem as well -- thanks to the extensibility of Postgres, this data could live right next to our normal relational data as well.

## So what's the best way to do it?

OK, so what's the answer at the end of the day? What's the best way to get to that useful v1? Here are the numbers:

In tabular form:

Approach	Avg (ms)	99%ile (ms)	99.999%ile (ms)
simple-counter	2.03	6	23
simple-hstore	2.15	6	16
assoc-table	2.5	8	30
hll	2.16	7	27

**If we go strictly with the data, the best way *looks* to be the `hstore`-powered solution, but I think the HLL is probably the right choice.**

The HLL results were quite variable -- some runs were faster than others, so I've taken the best of 3 runs.

Even though the data says `hstore`, knowing that posts will be seen by more and more people over time, I *might* choose the HLL solution for an actual implementation. It's far less likely to pose a bloated row problem, and it has the absolute correctness (and later



recall) of the assoc-table solution, while performing better over all (as you can imagine, no need to `COUNT` rows).

Another benefit of the HLL solution is that [PostgreSQL tablespaces](#) allow us to put the association table on a different, slower storage mechanism, and keep our `posts` table fast. Arguably in a real system we might have the HLL in something like `redis` but for a v1, it looks like Postgres does quite well!

## Wrap-up

I hope you enjoyed this look down the trunk hole, and you've got an idea of how to implement solutions to surprisingly complex problems like this one with Postgres.

As usual, Postgres has the tools to solve the problem *reasonably* well (if not completely) before you reach out for more complicated/standalone solutions.

**See any problems with the code, solutions that haven't been tried? -- reach out, or open an issue!**

## More Postgres resources

[Partial data dumps using Postgres Row Level Security](#)

[Postgres Views](#)

[Postgres Auditing in 150 lines of SQL](#)

[Cracking PostgreSQL Interview Questions](#)

[What are PostgreSQL Templates?](#)

[Realtime Postgres RLS on Supabase](#)

Share this article

Last post

**Supabase Flutter SDK 1.0 Developer**

**Preview**

2 August 2022

Next post

**Revamped Auth Helpers for Supabase (with  
SvelteKit support)**

13 July 2022

#### Related articles

[Supabase Beta May 2023](#)

[Supabase Vecs: a vector client for Postgres](#)

[Flutter Hackathon Winners](#)

[ChatGPT plugins now support Postgres & Supabase](#)

[Building ChatGPT Plugins with Supabase Edge Runtime](#)

[View all posts](#)

# Build in a weekend, scale to millions

Start your project



## Product

Database

Auth

Functions

Realtime

Storage

Vector

Pricing

Launch Week 7

## Developers

Documentation

Changelog

Contributing

## Resources

Support

System Status

Integrations

Experts

Brand Assets / Logos

DPA

SOC2

## Company

Blog

Customer Stories

Careers

[Open Source](#)

[SupaSquad](#)

[DevTo](#)

[RSS](#)

[Company](#)

[Terms of Service](#)

[Privacy Policy](#)

[Acceptable Use Policy](#)

[Support Policy](#)

[Service Level Agreement](#)

[Humans.txt](#)

[Lawyers.txt](#)

[Security.txt](#)

---

© Supabase Inc

