


The Sugar-free Chocolate of **Databases**

Matt Butrovich
Carnegie Mellon University
#1 Ranked CMU-DB PhD Student

Eating Smarter

Food labels in the US are wild.




Staring at DBMS Traces

[illegible]

Zero Calorie Queries

- 1,462,909 queries from various workloads...
- CMDBAC data set shows that they are 27% of all queries!



- Look in your SQL logs and you'll see these queries over and over again!

```
File Edit View Bookmarks Plugins Settings Help
2022-10-09T05:02:55.618041Z 1906268 Query SELECT `core_system`.`id`, `core_system`.`slug`, `core_system`.`name`, `core_system`.`created`, `core_system`.`modified`, `core_system`.`secret_key`, `core_system`.`view_count`, `core_system`.`ver` FROM `core_system` WHERE `core_system`.`id` = 277 LIMIT 21
2022-10-09T05:02:55.618834Z 1906268 Query SELECT `core_systemvisit`.`id`, `core_systemvisit`.`system_id`, `core_systemvisit`.`ip_address`, `core_systemvisit`.`user_agent`, `core_systemvisit`.`created` FROM `core_systemvisit` WHERE (`core_systemvisit`.`ip_address` = '122.165.124.50' AND `core_systemvisit`.`user_agent` = 'Mozilla/5.0 (Windows NT 10.0; Win64; x64) AppleWebKit/537.36 (KHTML, like Gecko) Chrome/76.0.3809.132 Safari/537.36')
2022-10-09T05:02:55.679771Z 1907203 Connect peloton_user@localhost on using Socket
2022-10-09T05:02:55.679936Z 1907203 Query SET NAMES utf8mb4
2022-10-09T05:02:55.680196Z 1907203 Query SET NAMES 'utf8mb4' COLLATE 'utf8mb4_unicode_520_ci'
2022-10-09T05:02:55.680251Z 1907203 Query SELECT @@SESSION.sql_mode
2022-10-09T05:02:55.680347Z 1907203 Query SET SESSION sql_mode='NO_ZERO_IN_DATE,ERROR_FOR_DIVISION_BY_ZERO,NO_AUTO_CREATE_USER,NO_ENGINE_SUBSTITUTION'
2022-10-09T05:02:55.680405Z 1907203 Init DB peloton
2022-10-09T05:02:55.680955Z 1907203 Query SELECT option_name, option_value FROM wp_options WHERE autoload = 'yes'
2022-10-09T05:02:55.685831Z 1907203 Query SELECT option_value FROM wp_options WHERE option_name = 'gzipcompression' LIMIT 1
:
```

DBMS Proxies to the Rescue

- PgBouncer, RDS Proxy, ProxySQL
- Features:
 - Connection pooling
 - Query rewriting
 - Sharding
 - Query caching

What I Do

- **Tigger** is a proxy that pushes Application Layer (i.e., L7) DBMS protocol logic into kernel-space via eBPF.
- Perform things like transaction pooling and workload replication without ever going to user-space. **User-bypass**.



**I am graduating in early 2024.
I will be expensive to hire.**

<https://mattbutrovi.ch>