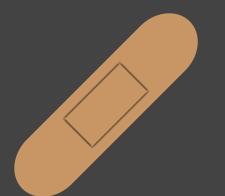




# Practical PgBouncer Pain Prevention



Nick Meyer @ Academia.edu PGConf NYC 2025



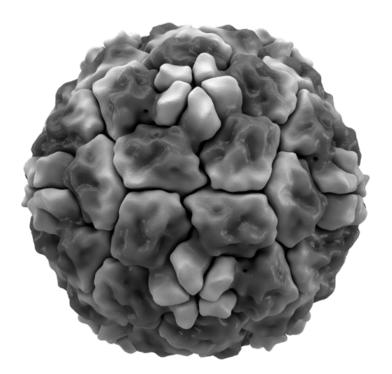


# Practical PgBouncer Pain Prevention



Nick Meyer @ Academia.edu PGConf NYC 2025

#### postgres problems:



"common cold"

#### pgbouncer problems:



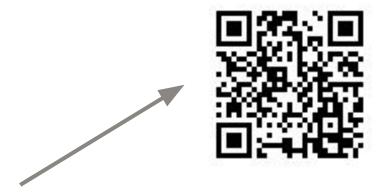


#### A bit about me (Nick Meyer)



- @ Academia.edu
- https://github.com/aristocrates





Slides:

https://github.com/aristocrates/pgconf\_nyc\_2025\_talks



#### **Objectives**

#### **PgBouncer:**

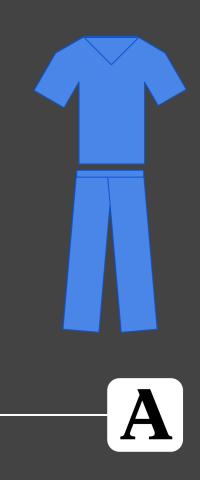
- 1. Configuration
- 2. Monitoring
- 3. Scaling ("multi-bouncer")



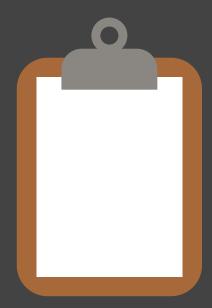


- PgBouncer alternatives
- TLS/authentication
- Medical advice





# Part 1 Configuration

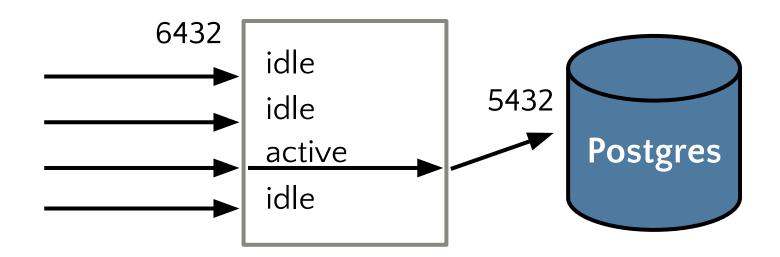


#### A Problem: too many connections

- Tons of clients
- Tons of connections
- Connections mostly idle
- Postgres connection = process = \$\$\$\$

## — A Solution: PgBouncer

"Connection pooler"



- Session
- Transaction
- Statement

#### BEGIN;

UPDATE users SET name = 'Nick' WHERE id = 1;

UPDATE user\_countries SET country = 'US' WHERE user\_id = 1;

COMMIT;

CREATE INDEX CONCURRENTLY
[...]

SELECT name FROM users WHERE id = 1;

- Session
- Transaction
- Statement

```
BEGIN;
```

UPDATE users SET name = 'Nick' WHERE id = 1;

UPDATE user\_countries SET country = 'US' WHERE user\_id = 1;

COMMIT;

CREATE INDEX CONCURRENTLY

SELECT name FROM users WHERE id = 1;

- Session
- Transaction
- Statement

```
BEGIN;
UPDATE users
SET name = 'Nick'
WHERE id = 1;
UPDATE user_countries
SET country = 'US'
WHERE user_id = 1;
COMMIT;
```

CREATE INDEX CONCURRENTLY
[...]

SELECT name FROM users WHERE id = 1;

- Session
- Transaction
- Statement



FROM users WHERE id = 1;

## pgbouncer.ini - databases section

#### /etc/pgbouncer/pgbouncer.ini

```
[databases]
db = host=A port=B dbname=C pool size=N
```

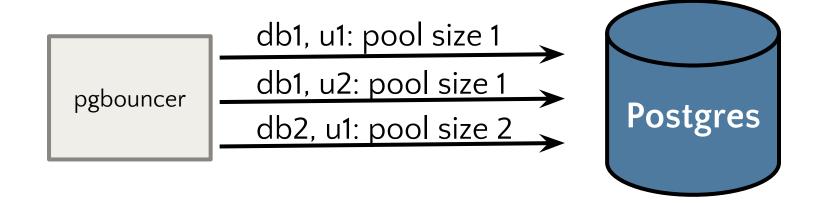
#### Many other params including:

- pool\_mode
- user



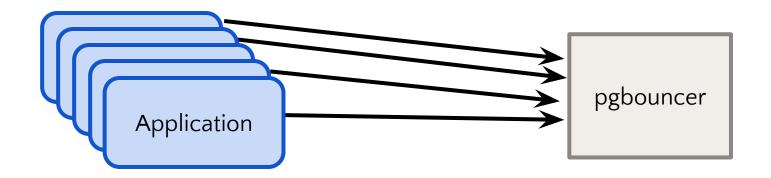
#### **Server connections**

- pool\_size
  - o per (database, user)
- sum(pool\_sizes) < max\_connections



## Client connections

- max\_client\_conn
- limitNOFILE
- /etc/security/limits.conf



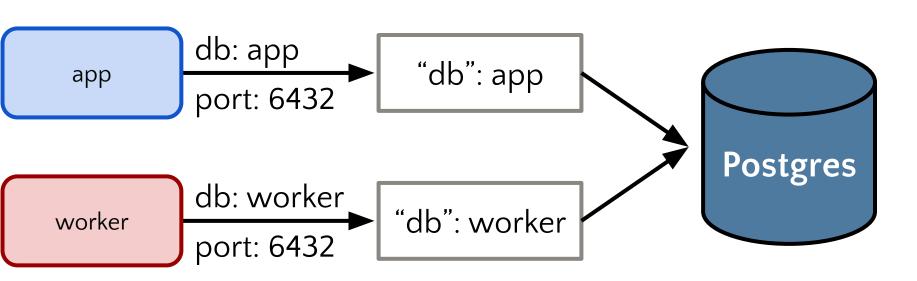
## A

#### Academia.edu uses multiple "databases"

```
[databases]
           = host= port= dbname=db pool size=A
   app
   worker = host= port= dbname=db pool size=B
                                                  enqueue
HTTP request
                               enqueue
                                         worker
                       app
response
```

## $\mathbf{A}$ Academia.edu uses multiple "databases"

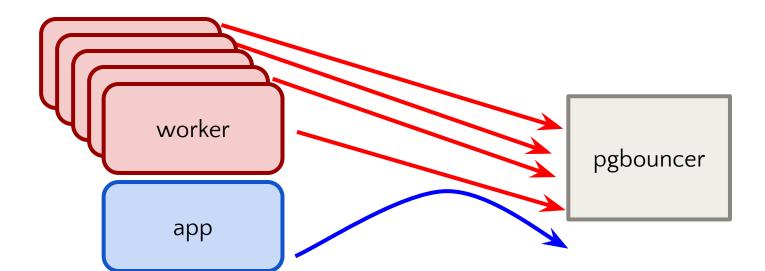
Separate "app" and "worker" with different pool sizes



## $\mathbf{A}$

#### Client connections (multiple workloads)

- max\_db\_client\_connections (1.24.0+)
- max\_user\_client\_connections (1.24.0+)



#### A Transaction mode: unsupported features

- Some statements are not supported
- Others can cause problems if they leak
- <a href="https://www.pgbouncer.org/features.html">https://www.pgbouncer.org/features.html</a>
- pgbouncer will not detect this



#### Transaction mode: unsupported features

SET/RESET	Never
LISTEN	Never
WITH HOLD CURSOR	Never
PREPARE / DEALLOCATE	Never
PRESERVE/DELETE ROWS temp tables	Never
LOAD statement	Never
Session-level advisory locks	Never

## A Prepared statements

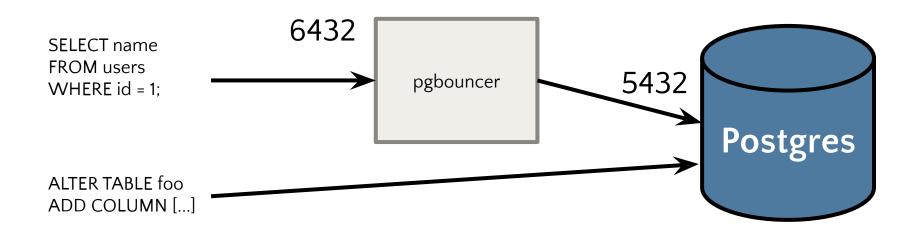
- (Before 1.21.0) need to avoid
- (1.21.0+) Protocol-level prepared statements
- **SQL** (PREPARE) not supported
- Deallocation
  - PQclosePrepared
  - Check language driver for support

## Schema changes

- lock\_timeout
- Make ORM set this
  - o Ruby on Rails: <u>Strong Migrations gem</u>
- SET lock\_timeout = '10s'
  - SET is not supported in transaction mode

## 

Most reliable: Connect directly to postgres for DDL



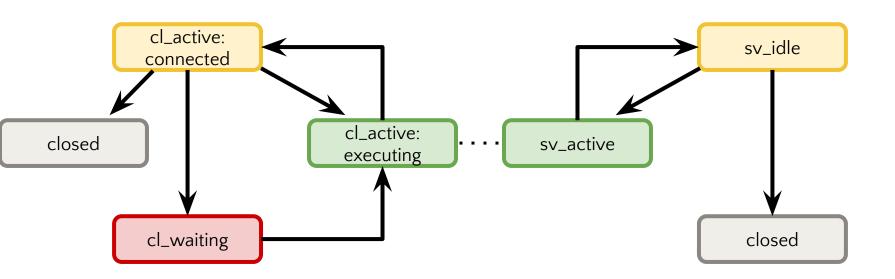
# Part 2 Monitoring



#### PgBouncer state machine (conceptual, not literal)

Client connections

Server connections





#### Connecting to pgbouncer

- Admin console
- psql -p 6432 -U pgbouncer -d pgbouncer
  - Commands: management, or viewing statistics
  - Does not actually allow general SQL
- You can avoid password if:
  - User pgbouncer + DB pgbouncer
  - Connection via socket (not -h localhost or IP)
  - Client user matches user running pgbouncer

#### $oxed{f A}$ SHOW pools;

- Pool = (database, user)
- cl\_active
- cl\_waiting
- sv\_active

```
pgbouncer=# show pools;
-[ RECORD 1 ]-----
database
                      pgbouncer
user
                      pgbouncer
cl active
cl waiting
cl active cancel req
cl waiting cancel_req |
sv active
sv active cancel
sv being canceled
sv idle
sv used
sv tested
sv login
maxwait
                      0
maxwait us
pool mode
                      statement
load balance hosts
```

## A cl\_active and cl\_waiting

- cl\_active includes both:
  - Actively running a query
  - Connected-but-idle, no queries waiting
- cl\_waiting
  - Sent a query
  - And: waiting for a server connection
- Want cl\_waiting to be 0

#### $|\mathbf{A}|$ sv\_active

- Server connections linked to a client
- If it "flatlines" at pool\_size, possible issue

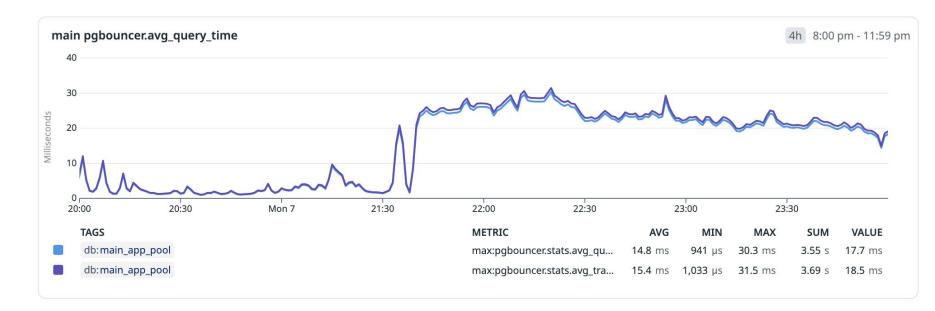
### $ig( \mathbf{A} ig)$ SHOW stats;

- SHOW stats\_averages;
- avg\_query\_time + avg\_xact\_time
- avg\_wait\_time
- Units: microseconds

## A

#### avg\_query\_time + avg\_xact\_time

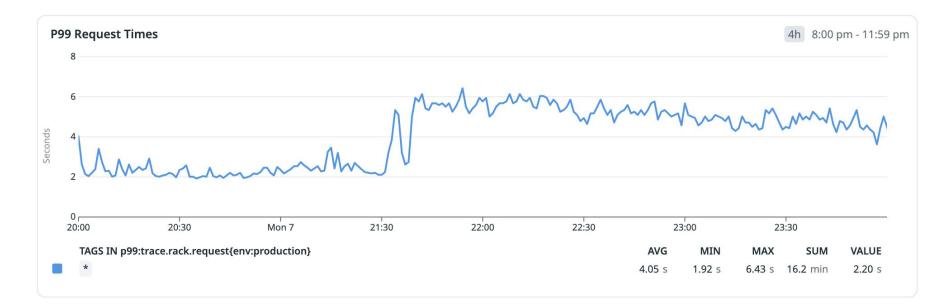
Useful for postgres query performance



## $\mathbf{A}$

#### avg\_query\_time + avg\_xact\_time

Useful for postgres query performance



#### $\{{f A}\}$ So-called "avg\_wait\_time"

- < 1.23.0: confusing units
  - Number of clients is **not** in the denominator
- >= 1.23.0: it's actually the average wait time
- Either way: watch for spikes

## $oldsymbol{A}$ SHOW clients;

- Interactive debugging
- psql [...] --csv

link not null => has server connection

```
-[ RECORD 1
type
user
                       app
database
                       app
replication
                       none
                       active
state
addr
                       unix
port
                       6432
local addr
                       บทเร
local port
                       6432
connect time
                       2025-09-30 13:30:00 EDT
request time
                       2025-09-30 13:30:05 EDT
wait.
                       31
wait us
                       912710
close needed
                       0x5db77e850be0
ptr
link
                       0x1c38c40
remote pid
                       403917
tls
application name
                       psql
prepared statements
id
```

pgbouncer=# show clients;

### $\mathbf{A}$

#### PgBouncer process CPU



- PgBouncer is single-threaded
- SHOW stats; does not show CPU util
- top
- pidstat -p \$PID 1 1
- If it hits 100%, other metrics start to look confusing



#### **Postgres monitoring - Queries**

#### Live queries

- pg\_stat\_activity
- idle in transaction count

#### Queries over time

pg\_stat\_statements

#### Wait events

- pg\_wait\_sampling
- Datadog database monitoring
- AWS RDS performance insights

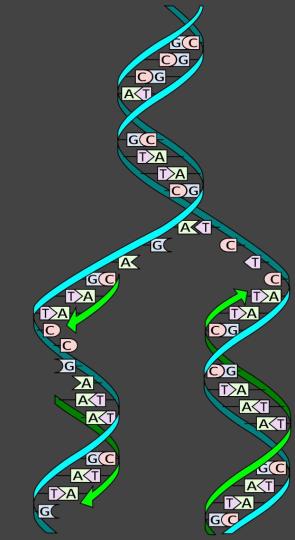


#### Postgres monitoring - PgHero

- Connections
- Live Queries

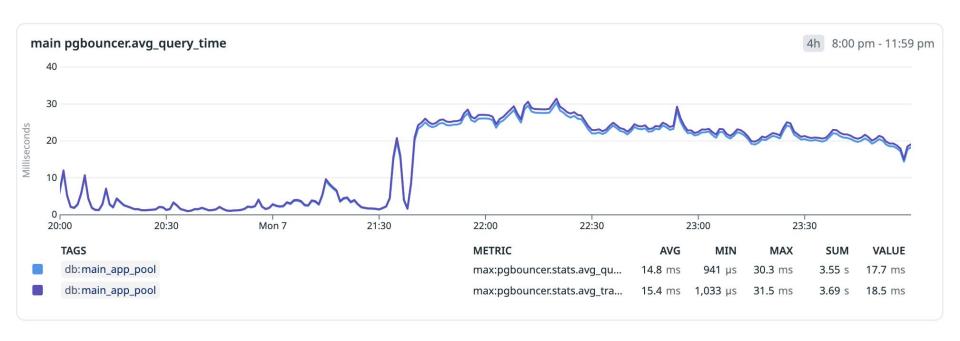


# Part 3 Scaling



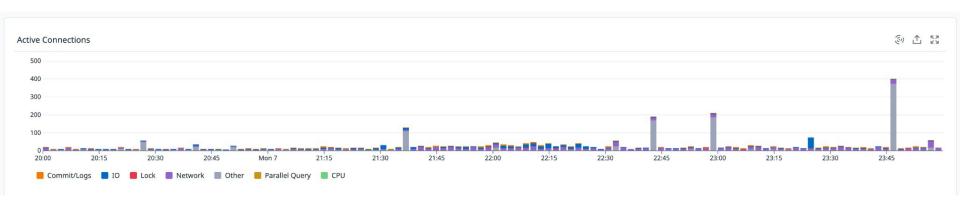






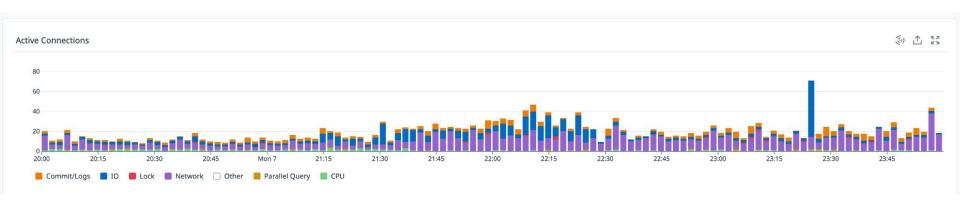


Wait events - not the same graph



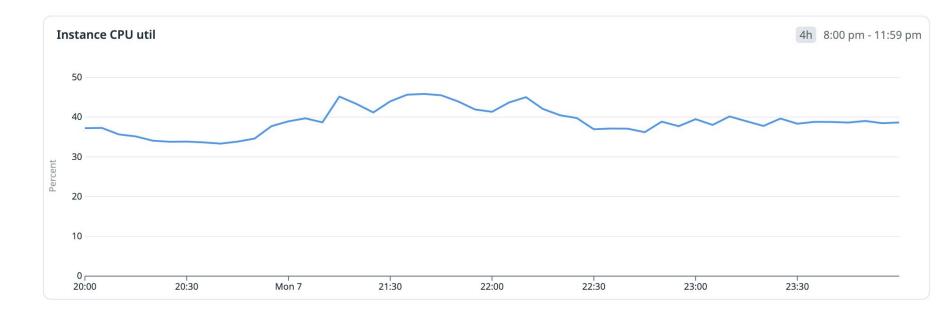


Wait events - not the same graph



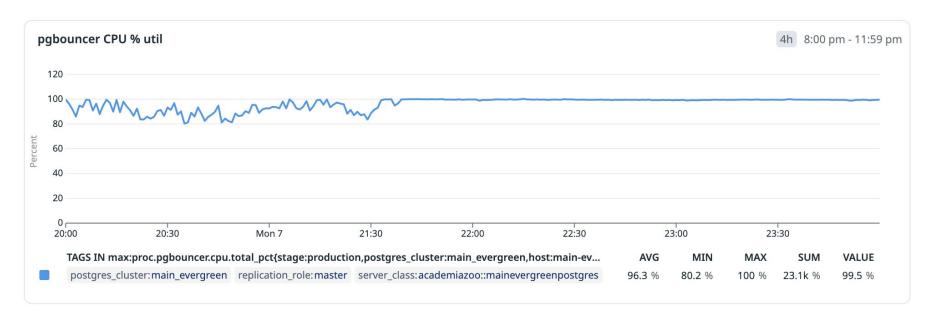


#### Instance CPU util not elevated





#### pgbouncer CPU at 100%





#### Signs that you need multiple PgBouncer processes

#### You will see:

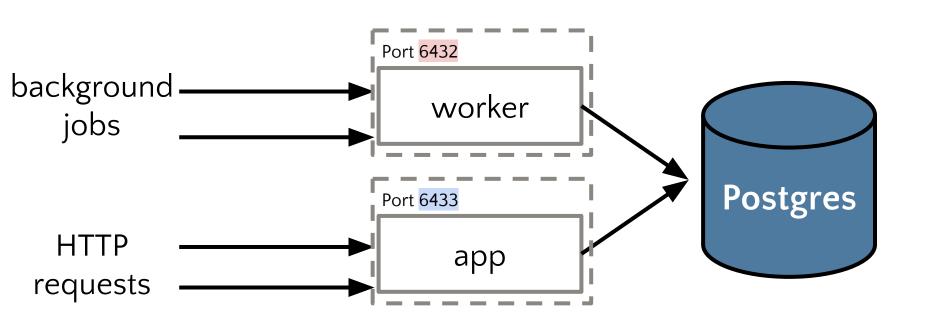
- PgBouncer 100% CPU
- Application sees high end-to-end query times

#### You might see:

- High avg\_query\_time and/or avg\_wait\_time
- High cl\_waiting
- Low/stable active sessions on postgres

# $-\mathbf{A}$

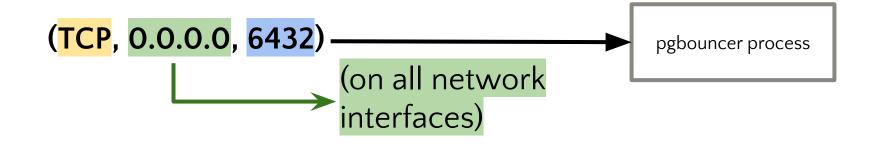
#### Multiple PgBouncers: by user



# A

#### Multiple PgBouncers: sharing the same port

- socket = abstraction over network communication, file
- Connection socket
  - remote ip, remote port, local ip, local port, protocol (TCP|UDP)
- Listening socket
  - Bind: (protocol, ip, port) -> process



### Multiple PgBouncers: sharing the same port

#### (normally) If you try to listen on 6432 on another process:

\$ pgbouncer /etc/pgbouncer/pgbouncer.ini
FATAL unix socket is in use, cannot continue

# A

#### Multiple PgBouncers: sharing the same port

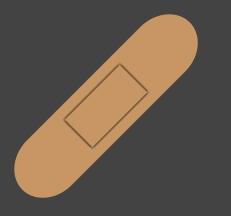
- SO\_REUSEPORT
- Load balancing from the networking stack
   Hash of (remote ip, remote port, local ip, local port)
- In the pgbouncer config:

```
so reuseport = 1
```

And then run multiple processes with that config

# Multiple PgBouncer tips

- Confirm client query times good
- Check connection counts to each
- Check all things that depend on pgbouncer
  - Monitoring
  - High availability solution
  - Upgrade process
  - Sysadmin scripts and playbooks





# Summary





#### Monitor PgBouncer process CPU util

- PgBouncer is single threaded
- CPU at 100% => bottleneck
- Might not be "out of the box" with monitoring solution



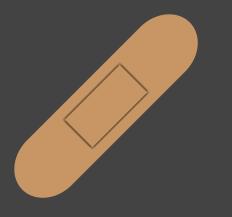
#### Do we need PgBouncer alternatives?

- Managing multiple processes can be a pain
- Transaction mode has limitations
- A lot of pain points have been fixed in recent versions
  - avg\_wait\_time
  - (protocol) prepared statements
  - max\_db\_client\_connections

### -A Summary

- Monitor pgbouncer process CPU util
  - If pgbouncer hits 100%, run more
- Upgrade
- Be aware of transaction mode limitations

# Questions?





# Appendix



# - A References

- https://www.heap.io/blog/decrypting-pgbouncers-diagnostic-information
- https://jpcamara.com/2023/04/12/pgbouncer-is-useful.html
- <a href="https://www.crunchydata.com/blog/postgres-at-scale-running-multiple-pgb">https://www.crunchydata.com/blog/postgres-at-scale-running-multiple-pgb</a>
  ouncers
- <a href="https://www.crunchydata.com/blog/prepared-statements-in-transaction-mode-for-pgbouncer">https://www.crunchydata.com/blog/prepared-statements-in-transaction-mode-for-pgbouncer</a>

# -A Image credits

- <a href="https://en.wikipedia.org/wiki/File:Rhinovirus\_isosurface.png">https://en.wikipedia.org/wiki/File:Rhinovirus\_isosurface.png</a> CC BY-SA
  - Author: <u>Thomas Splettstoesser</u>
- https://en.wikipedia.org/wiki/File:DNA\_replication\_split.svg
  - Author: <u>Madprime</u>



#### Schema changes - other approach

In theory: Transactions + SET LOCAL

- Still doesn't help with CREATE INDEX CONCURRENTLY
- Your ORM probably does not support this

