


Tame the growth of your database

Adam Wolk
PGConf.EU
2023 Prague



Multiversion Concurrency Control, MVCC

- Each SQL statement sees a snapshot
- Less lock contention
- Reads don't block writes
- Writes don't block reads



copy-on-write

- Hidden table columns
- xmin, xmax, ctid



Time travel

<https://dsf.berkeley.edu/papers/ERL-M90-34.pdf>

Lastly, POSTGRES supports the notion of **time travel**. This feature allows a user to run historical queries. For example to find the salary of Sam at time T one would query:

```
retrieve (EMP.salary)
using EMP [T]
where EMP.name = "Sam"
```

POSTGRES will automatically find the version of Sam's record valid at the correct time and get the appropriate salary.

What is bloat and why too much is bad?

- Redundant
- Sparse data
- Disk size
- Memory
- IOPS



Vacuum

- Recycles storage
- Freeze txids
- Removing clog files
- Rewrites (FULL)
- Update (Free Space Map) FSM
- Update Visibility Map (VM)
- Gather stats (ANALYZE)



Autovacuum

- Prevent txid wraparound
- Gather statistics



Doctor, it
hurts when I
do this –
**then do it
more often!**

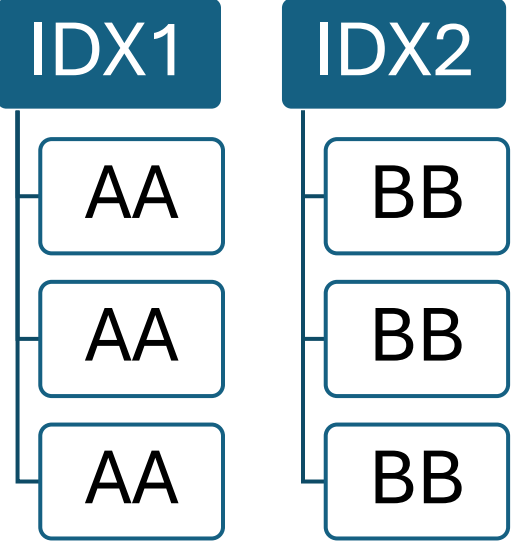
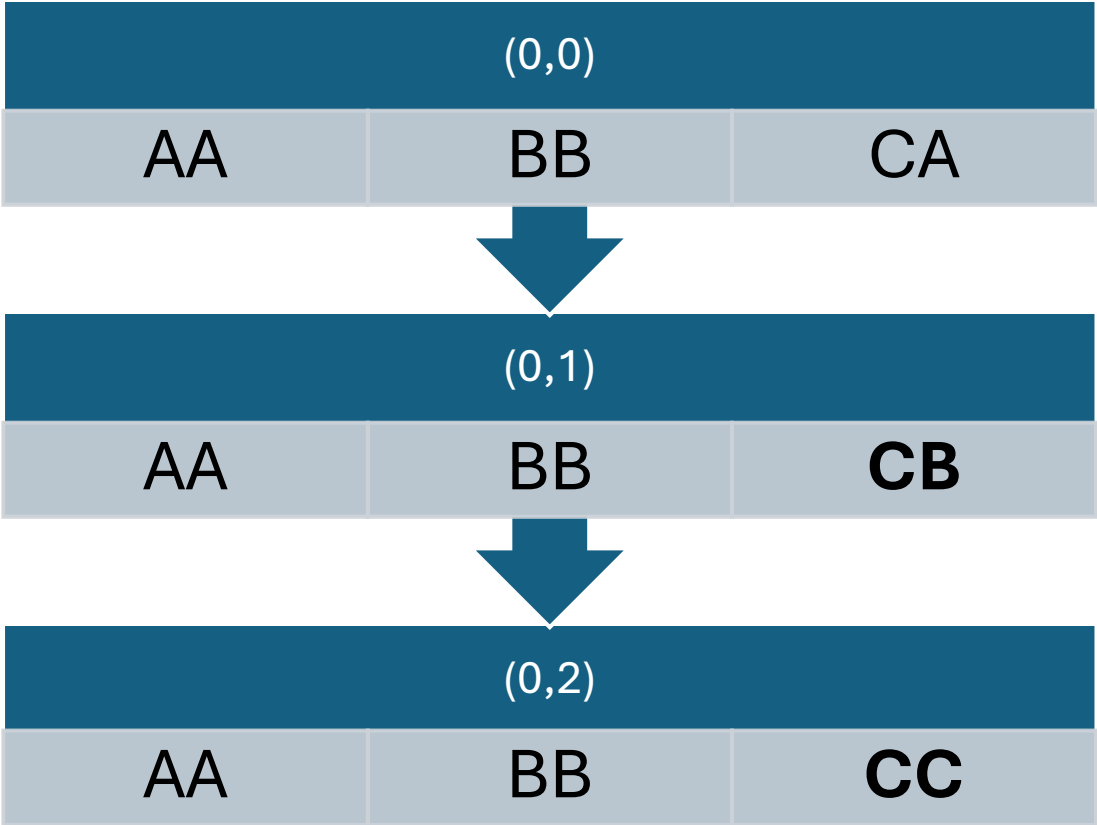
- VACUUM after loads
- Less rows frequently
- autovacuum starvation

Heap-Only Tuples (HOT) updates

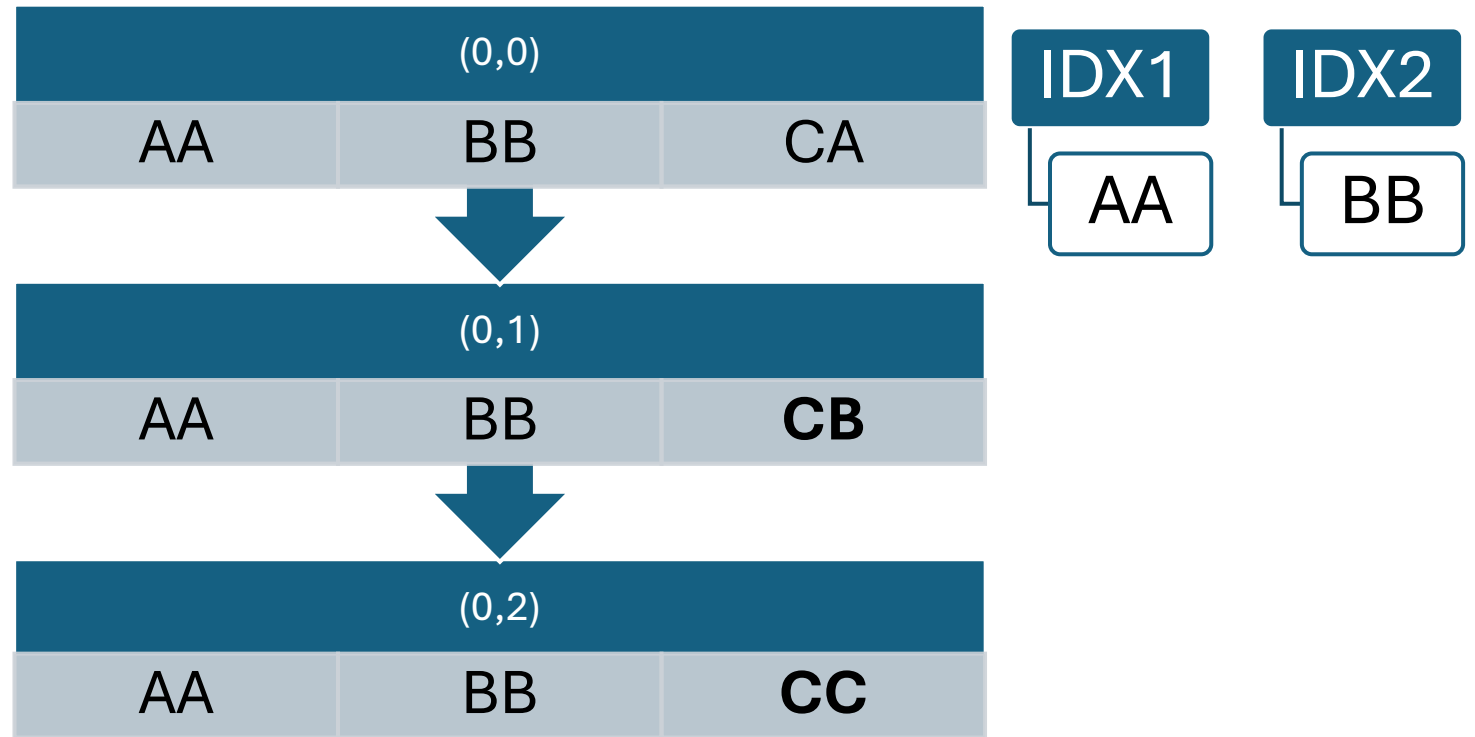
- Write amplification
- UPDATE can't touch indexed columns
- Space available on the page
- Some bloat can help (or decreased fillfactor)
- Inline vacuum




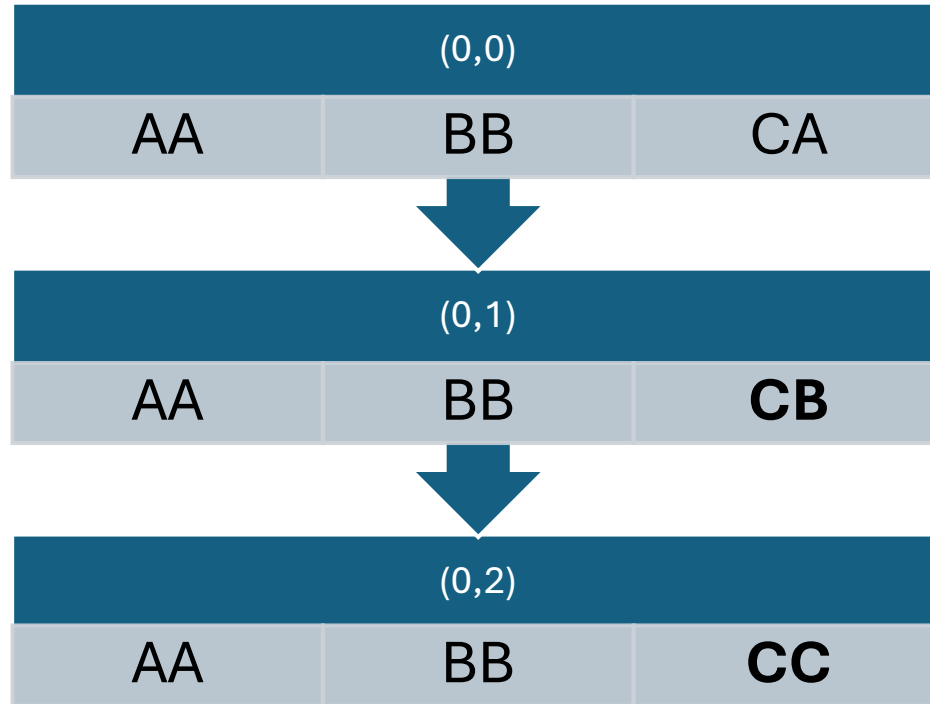
Without HOT updates



When HOT



HOT
prevented



Doctor, it
hurts when I
do this –
**then don't
do that!**


- Understand when data is copied
- Prevent generating bloat
- Don't blindly copy patterns
- Know your tools

Django ORM

- Same value is still a copy
- `col = col` is still a copy
- `old != new` prevents HOT updates when indexed
- Use `.save(update_fields=..)`
- Careful with auto fields and triggers

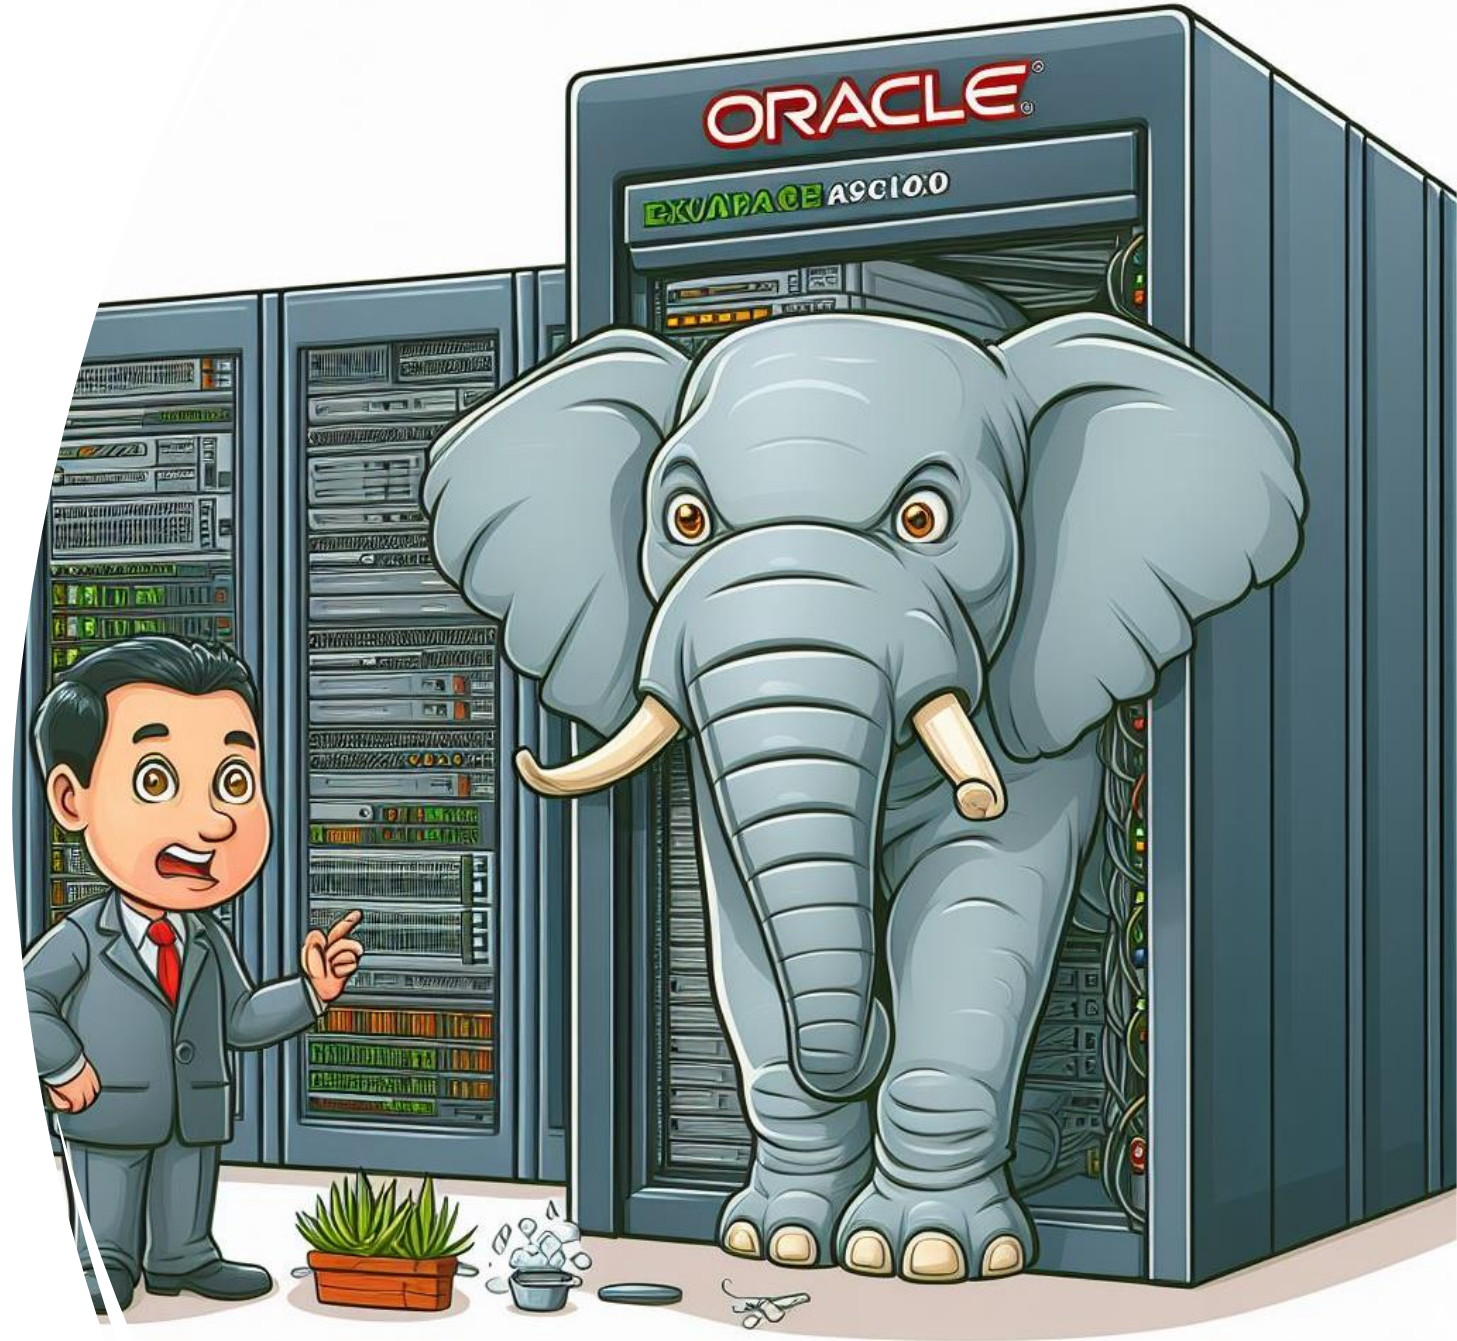


Django DEMO



Oracle ROWID

- ROWID != ctid
- UPDATE tombstone by ctid instead of DELETE
- Index on tombstone preventing HOT updates



DELETE DEMO





Thank you

-
- @mulander
 - adamwolk@microsoft.com