

# The next ten years of Postgres

Álvaro Herrera – Postgres developer, EDB



Postgres Conference Germany 2026  
Essen 20-21 April 2026

## The Next Ten Years

Thank you from EDB for being here!

I'm feeling lucky

Ten years is a long time

I'm feeling lucky

Ten years is a long time ...

... but ten years ago feel like yesterday

## Postgres in 2016

2016: Postgres 9.6 was released

## Postgres in 2016

2016: Postgres 9.6 was released

**It was great!**

# What was Postgres 9.6 lacking

- no declarative partitioning
- no logical replication
- no MERGE
- no query paralellism
- no generated columns
- no extended statistics

no temporal SQL features, no advanced SQL/JSON features, no table access methods, no JIT  
( ... and so on)

# Postgres 19 is coming

- SQL/PGQ
- `pg_plan_advise`
- additional temporal query features
- concurrent vacuum full

But let's talk about Postgres 29

Let's talk about Postgres 29

# Table Access Methods (table AMs)

- today we have an API ...  
but we don't offer an in-core choice
- In 10 years we could have several

## Table Access Methods (table AMs)

- today we have an API ...  
but we don't offer an in-core choice
- In 10 years we could have several
- zheap (no need for vacuum)

## Table Access Methods (table AMs)

- today we have an API ...  
but we don't offer an in-core choice
- In 10 years we could have several
- zheap (no need for vacuum)
- columnar store

## Table Access Methods (table AMs)

- today we have an API ...  
but we don't offer an in-core choice
- In 10 years we could have several
- zheap (no need for vacuum)
- columnar store
- timeseries

## Table Access Methods (table AMs)

- today we have an API ...  
but we don't offer an in-core choice
- In 10 years we could have several
- zheap (no need for vacuum)
- columnar store
- timeseries
- temporal data store?
- append-only store?
- index-organized tables?

# Logical replication

- Multi-primary replication
  - bi-directional replication

# Logical replication

- Multi-primary replication
  - bi-directional replication
- Better conflict handling

# Logical replication

- Multi-primary replication
  - bi-directional replication
- Better conflict handling
- Conflict-resistant datatypes

# Logical replication

- Multi-primary replication
  - bi-directional replication
- Better conflict handling
- Conflict-resistant datatypes
- Replication of DDL

# Concurrent DDL

- We just got concurrent vacuum full
- Infrastructure reuse for ALTER TABLE

## Concurrent DDL

- We just got concurrent vacuum full
- Infrastructure reuse for ALTER TABLE
- Unintrusive CREATE INDEX CONCURRENTLY

# Analytics

- Batched executor
- Vectorized execution
- Together with tailored table AMs and async I/O, these features will make Postgres a top-notch analytics platform

# Sharding / MPP – Massive Parallel Processing

- EDB has WarehousePG
  - inherited 21 years of Greenplum technology
  - Amazing for analytics
  - Come to the booth!

# Sharding / MPP – Massive Parallel Processing

- EDB has WarehousePG
  - inherited 21 years of Greenplum technology (and engineers)
  - Amazing for analytics
  - Come to the booth!
- ... but in Postgres, we need a plan!

# Multithread Server

- better performance
- hybrid model at first
- process-based model removed in a few years

# Asynchronous Input/Output (AIO)

- All reads and writes will use async I/O

# Asynchronous Input/Output (AIO)

- All reads and writes will use async I/O
- Direct I/O will be reasonable

# Asynchronous Input/Output (AIO)

- All reads and writes will use async I/O
- Direct I/O will be reasonable
- Better performance?
- Better RAM utilization

## Development community

- Postgres has gotten more and more complex
- It will not get any simpler!
- Contributor mentoring program is crucial
- Looking forward to twice as many committers as today, four times as many other core code contributors as today

# All for what?

Upwards movement of:

- usability
- scalability
- performance
- observability
- extra features

... so that Postgres continues to be  
the Most Loved Database System