

Postgres Past Present, Future

Oleg Bartunov
Postgres Professional

March 2, 2017, Tel Aviv



When I started using Postgres

- No UTF-8, even no 8-bit
- No WAL
- No MVCC
- No replication
- No usable non-scalar data types
- No subselects, no window functions, no CTE
- It was Postgres95

What is PostgreSQL

PostgreSQL — The world's most advanced open source object-relational database.


Extendable — data types, operators, functions, indexes.

[ANSI SQL](#) (1992, 1999, 2003, 2008, 2011),
NoSQL (key-value, JSON, JSONB)

Developed and supported by independent international community (users, developers, companies)

Pronunciation: post-gress-Q-L, post-gres, pgsql

Web: <http://www.postgresql.org>, ([BSD](#), [MIT](#)) - like



Important features of PostgreSQL

Legendary reliability and stability

It is extremely common for companies to report that PostgreSQL has never crashed for them in several years of high activity operation. Not even once. It just works.

Cross platform

PostgreSQL is available for almost every brand of Unix. The last stable version runs on over 34 platforms. You can use PostgreSQL on Windows as well..

Appropriate for high volume environments

PostgreSQL uses a multiple row data storage strategy called MVCC (Multiversion Concurrency Control). This makes it especially suitable for high volume environments. Proprietary vendors use this strategy for that reason as well.

Scalability

PostgreSQL is well suited for the modern multicore CPU and his performance is linear scalable upto 512 cores. Cluster solutions like Postgres XL/XC provides horizontal scalability.

Extendability

Extendability is an ability to develop a new functionality (data types, queries, access methods) by application programmers (not PostgreSQL core developers) and without stopping server.

Availability

Liberal BSD, MIT-like license

- * More profitable business models with wide-scale deployment.
- * No possibility of being audited for license compliance at any stage.
- * Flexibility to do concept research and trial deployments without needing to include additional licensing costs.

Independency

PostgreSQL is developing by the international community (users and companies), no dependency on one vendor

Excellent Support

PostgreSQL has a vibrant community of committed professionals that are available to help you. In addition, many commercial companies offer support solutions.

PostgreSQL users

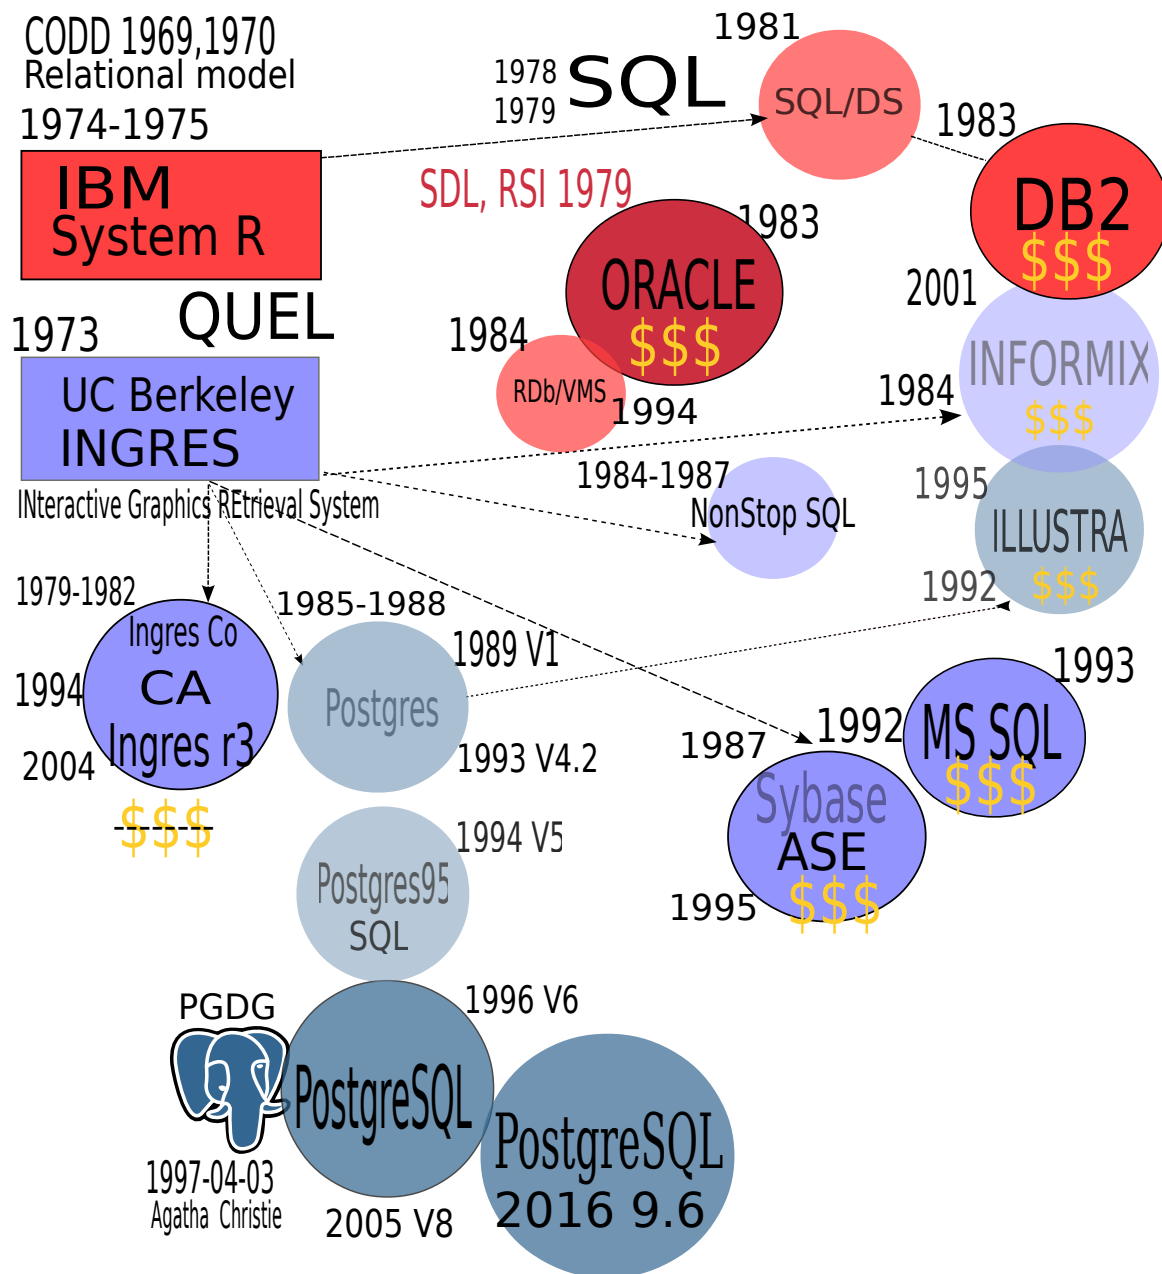


+BIG RUSSIAN Enterprise !



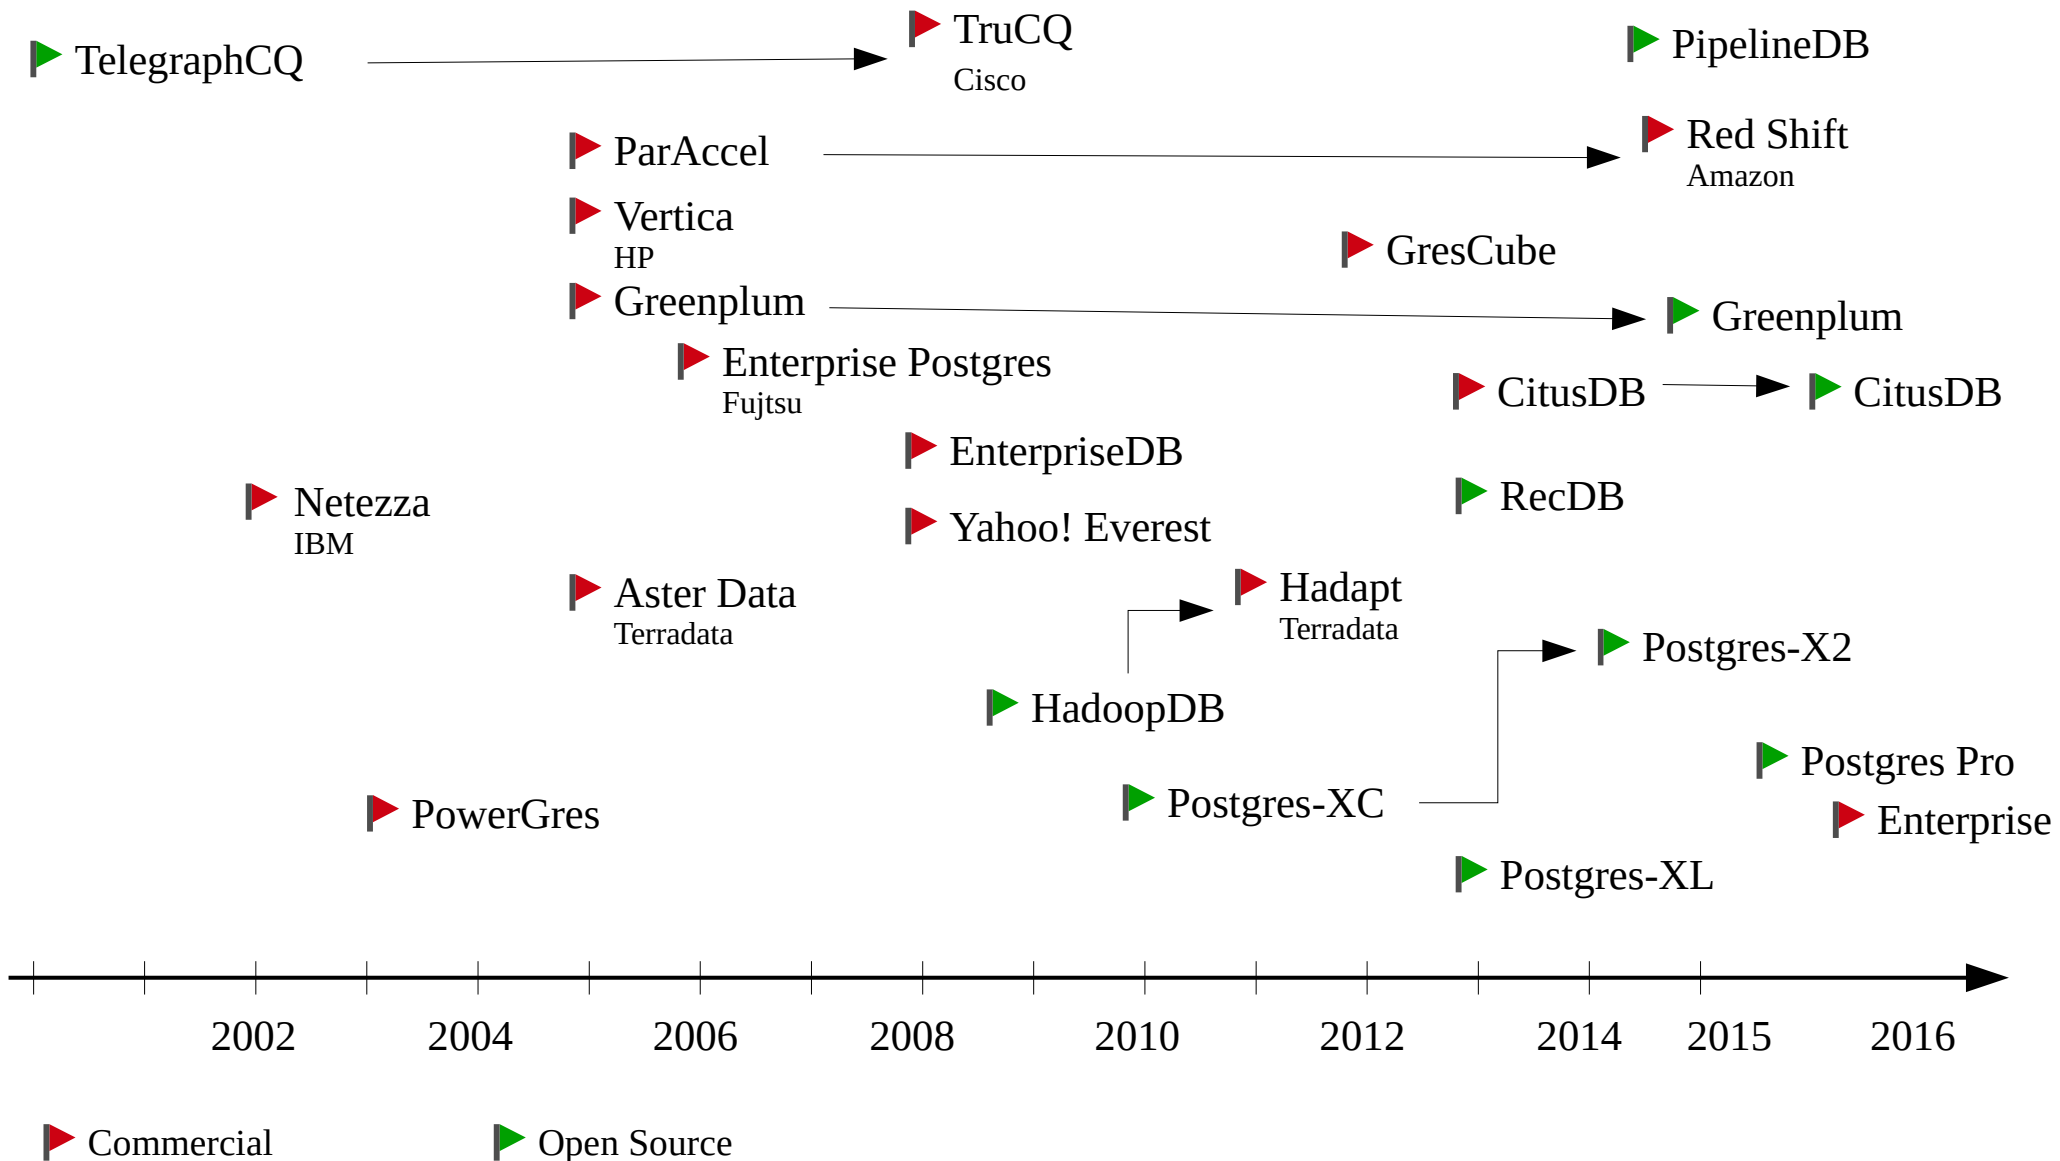
Michael Stonebreaker
Turing Award, 2015

PostgreSQL History




2017 10

PostgreSQL Forks (we love forks!)



PostgreSQL versions (~300)

PostgreSQL release timeline



Yesterday (9.5)

INSERT ... ON CONFLICT (aka Upsert)

- Row Level Security
- BRIN index
- Improved sorting (abbreviated keys)
- GROUP BY ... GROUPING SETS, CUBE, ROLLUP (multilevel aggregates)
- Improvement FDW (import schema)
- SELECT ... SKIP LOCKED
- SELECT ... TABLESAMPLE
- Vacuumdb -j (parallel vacuum)

UNFORK

- Citus DB – distributed database as an extension !
- GreenPlum – open sourced !

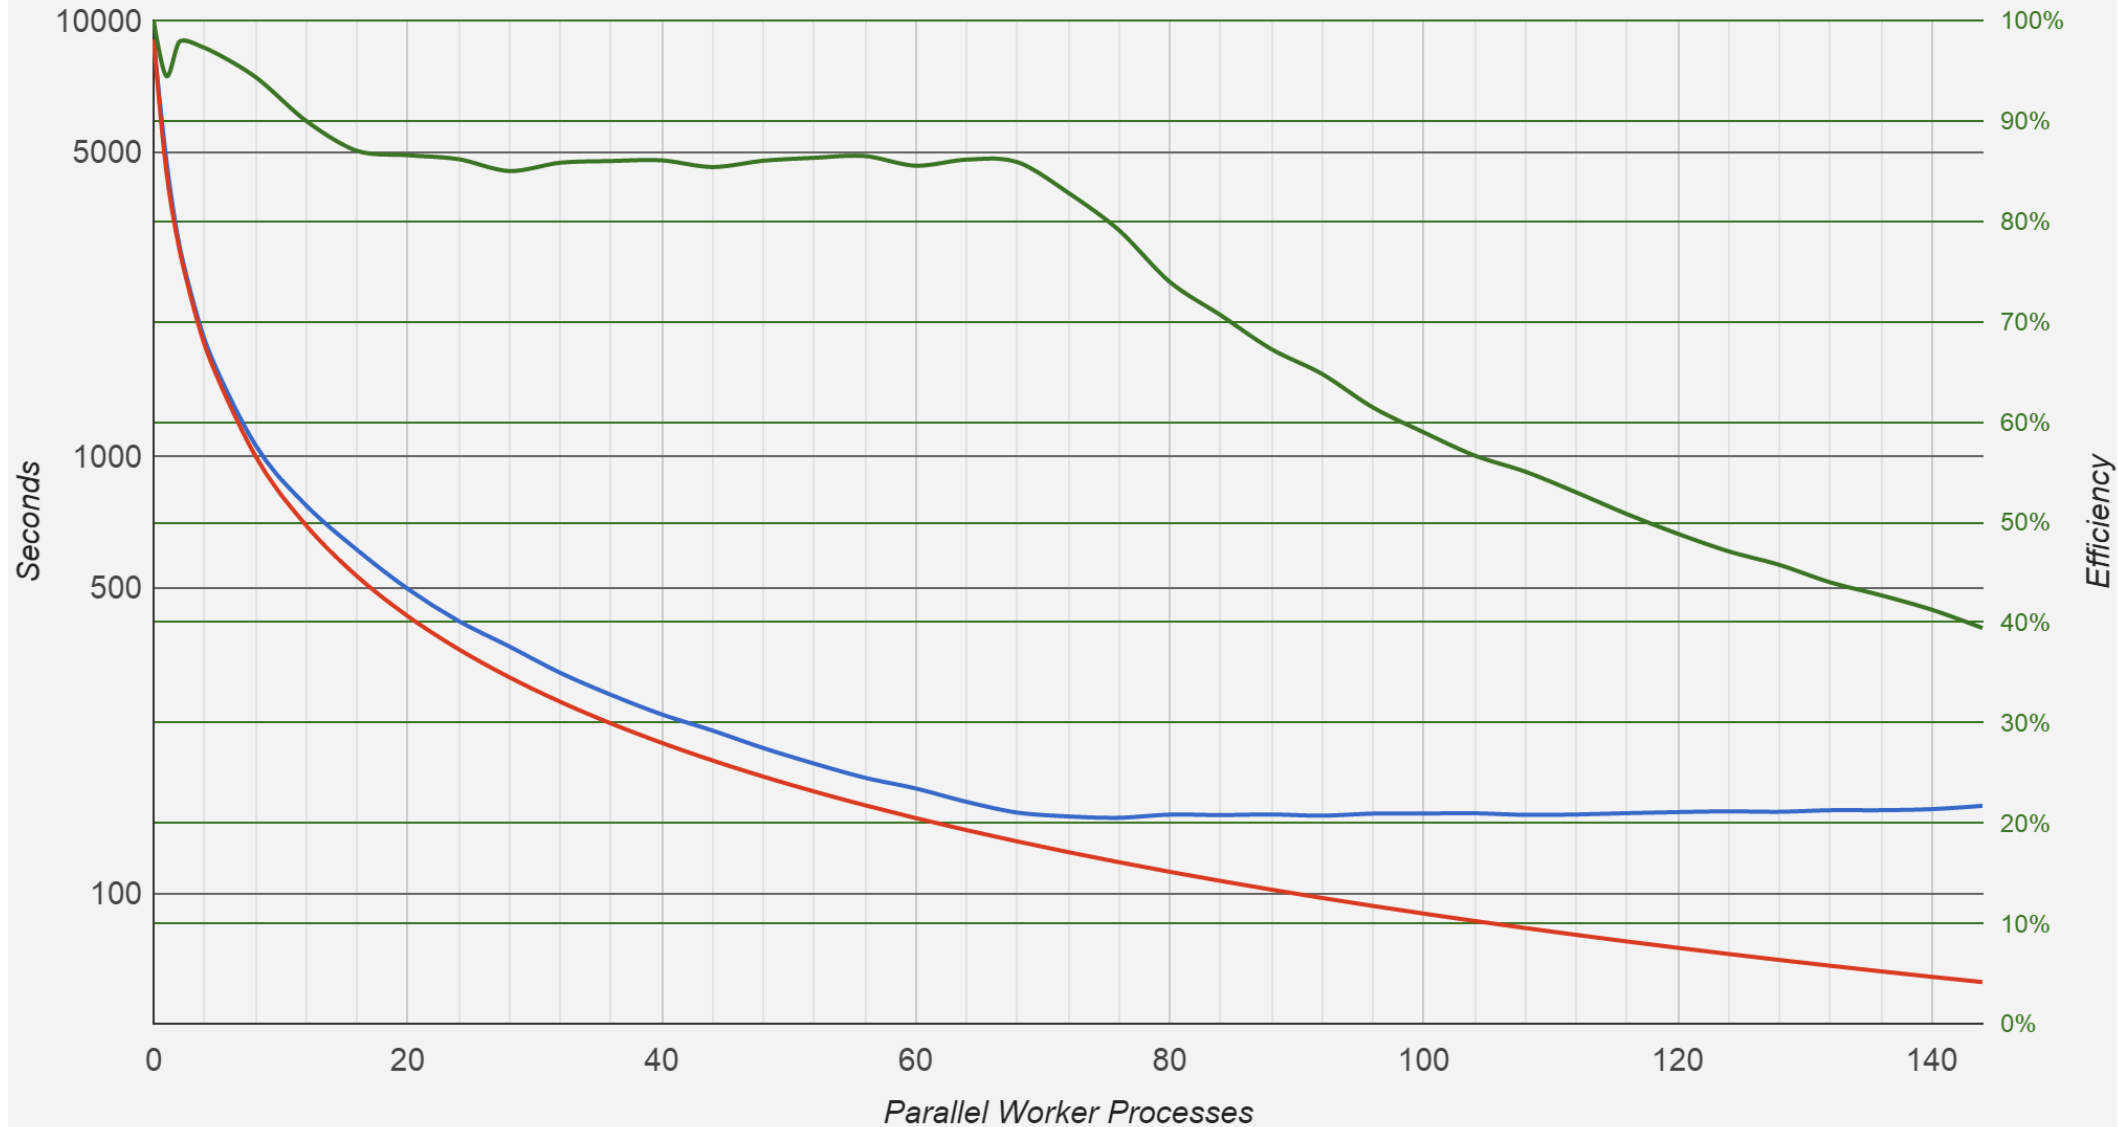
Today (9.6)

- Parallel query execution (sequence scan, join, aggregate)
- Better VACUUM with very big tables
- Extend relations multiple blocks at a time to improve scalability
- Improvement FDW (push down join, DML, sort)
- Indexing boxes and polygons with SP-GiST
- CREATE ACCESS METHOD, GENERIC WAL — (RUM access method)
- FTS improvement — phrase search
- Combining aggregates
- IOS for partial indexes
- Wait monitoring
- Multicore scalability improvement

Parallel Query (72 core)

TPC-H Q1 Parallel Query 1TB

— Actual Time (seconds) — Hypothetical time @ 100% efficiency (seconds) — Actual Efficiency (%)



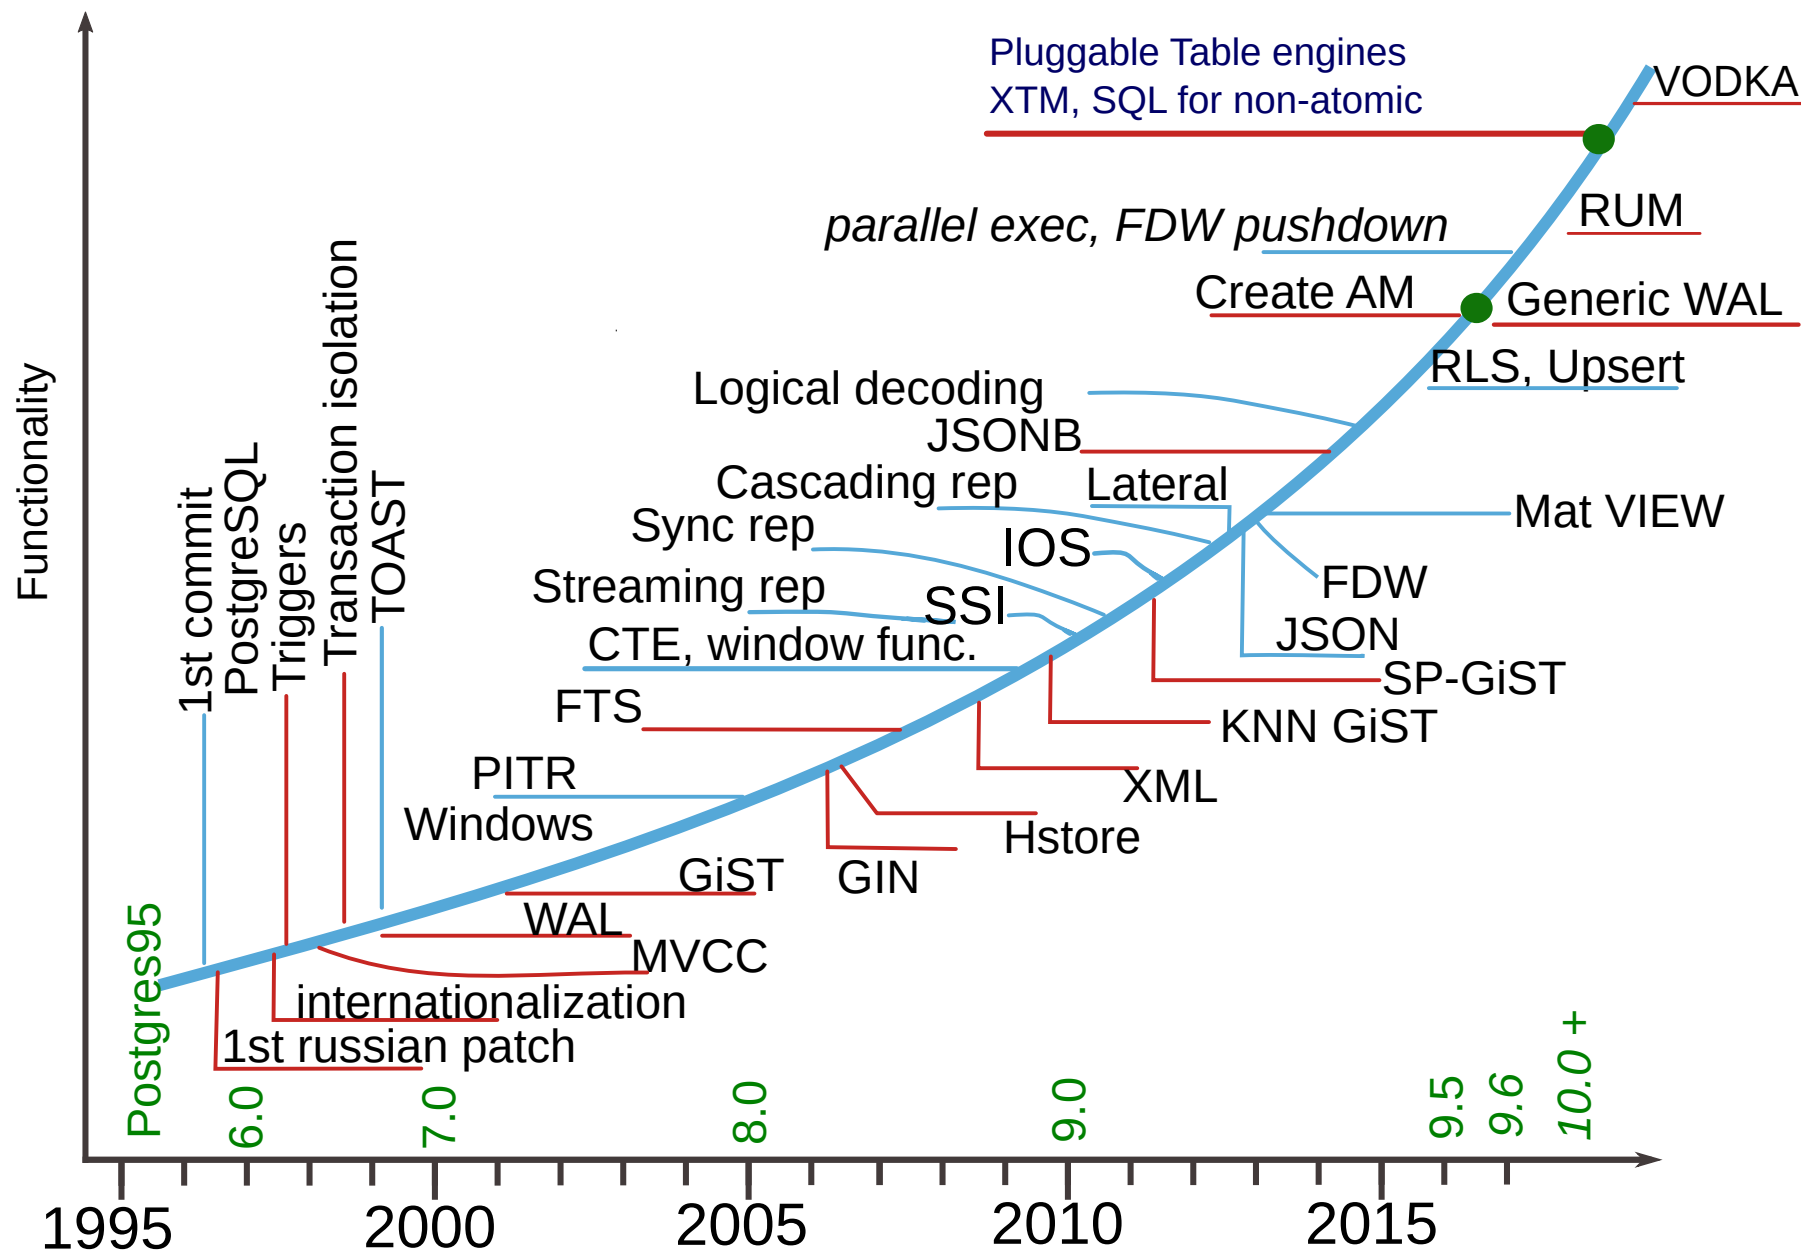
Tomorrow (10)

- Logical replication in core
- Declarative table partitioning
- FDW pushdown aggregates
- Parallel joins
- Parallel index scan
- Some features from the last commitfest

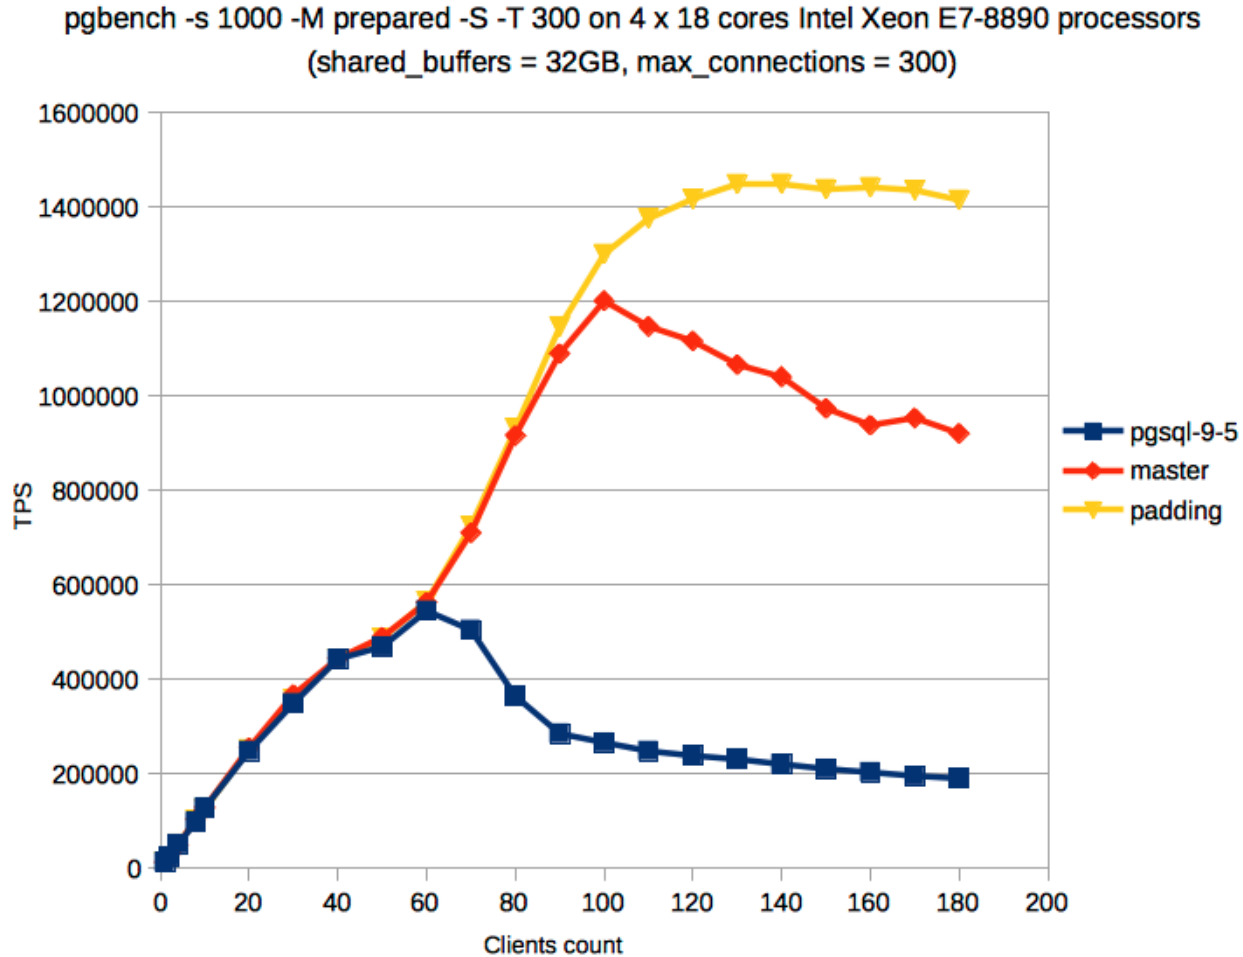
<https://commitfest.postgresql.org/13/>

- Covering indexes
- SQL/JSON ?
-many others....

20 years of Postgres Development




Multicore scalability



TODO:

- NUMA support

There are many ways to help Postgres



There are many ways to help Postgres

Core development

Development, review, testing, reporting bugs

Ecosystem

Extensions, drivers, ORM, monitoring tools... Postgres support in applications
Distributions, packages

Documentation

Improvement, translations, writing books, papers, ...blogging!

Meetings, Education

Creating of local communities,
Conference, meetups, seminars,
hackatons, educational and training
courses

Use PostgreSQL!

Use Postgres in your company !

Sponsorship

Help development, sponsor
community events.

~~When I started using Postgres~~ (Today)

- Excellent scalability on 2-socket machines
 - Waiting for improvement for 4-sockets
- Support a wide range of workloads
 - High security data, relational, unstructured, OLAP
- Postgres-centric companies
 - 2ndQuadrant, EDB, Postgres Professional ([roadmaps](#))
- Postgres groups
 - NTT, Fujitsu, Amazon, Alibaba, Tencent,...
- Postgres Development from
 - Community driven → Business driven
- Postgres is used in mission-critical enterprise systems

When I started using Postgres (Today)

- Postgres is #4 !

318 systems in ranking, February 2017

| Rank | | | DBMS | Database Model | Score | | |
|----------|----------|----------|------------------------|-------------------|----------|----------|----------|
| Feb 2017 | Jan 2017 | Feb 2016 | | | Feb 2017 | Jan 2017 | Feb 2016 |
| 1. | 1. | 1. | Oracle + | Relational DBMS | 1403.83 | -12.89 | -72.31 |
| 2. | 2. | 2. | MySQL + | Relational DBMS | 1380.30 | +14.02 | +59.18 |
| 3. | 3. | 3. | Microsoft SQL Server + | Relational DBMS | 1203.45 | -17.50 | +53.23 |
| 4. | ↑ 5. | ↑ 5. | PostgreSQL + | Relational DBMS | 353.68 | +23.31 | +65.02 |
| 5. | ↓ 4. | ↓ 4. | MongoDB + | Document store | 335.50 | +3.60 | +29.90 |
| 6. | 6. | 6. | DB2 + | Relational DBMS | 187.90 | +5.41 | -6.58 |
| 7. | 7. | ↑ 8. | Cassandra + | Wide column store | 134.38 | -2.06 | +2.62 |
| 8. | 8. | ↓ 7. | Microsoft Access | Relational DBMS | 133.39 | +5.94 | +0.31 |
| 9. | ↑ 10. | 9. | SQLite | Relational DBMS | 115.31 | +2.93 | +8.53 |
| 10. | ↓ 9. | 10. | Redis + | Key-value store | 114.03 | -4.66 | +11.96 |

Several Postgres groups are working on

- Postgres Distributed
- Postgres Vectorized
- Postgres Parallel
- Postgres Asynchronous
- Postgres Extendable+
- Postgres NoSQL
- Postgres Scalable (Vertical & Horizontal)

Cluster


FDW

GUI

CitusDB

GreenPlum


Postgres-X[CL]



PostgreSQL Core

Multicore parallelism


Backup




Pluggable storages




Backup: full, incremental, partial



Builtin HA cluster



Advanced GUI



Cloud version

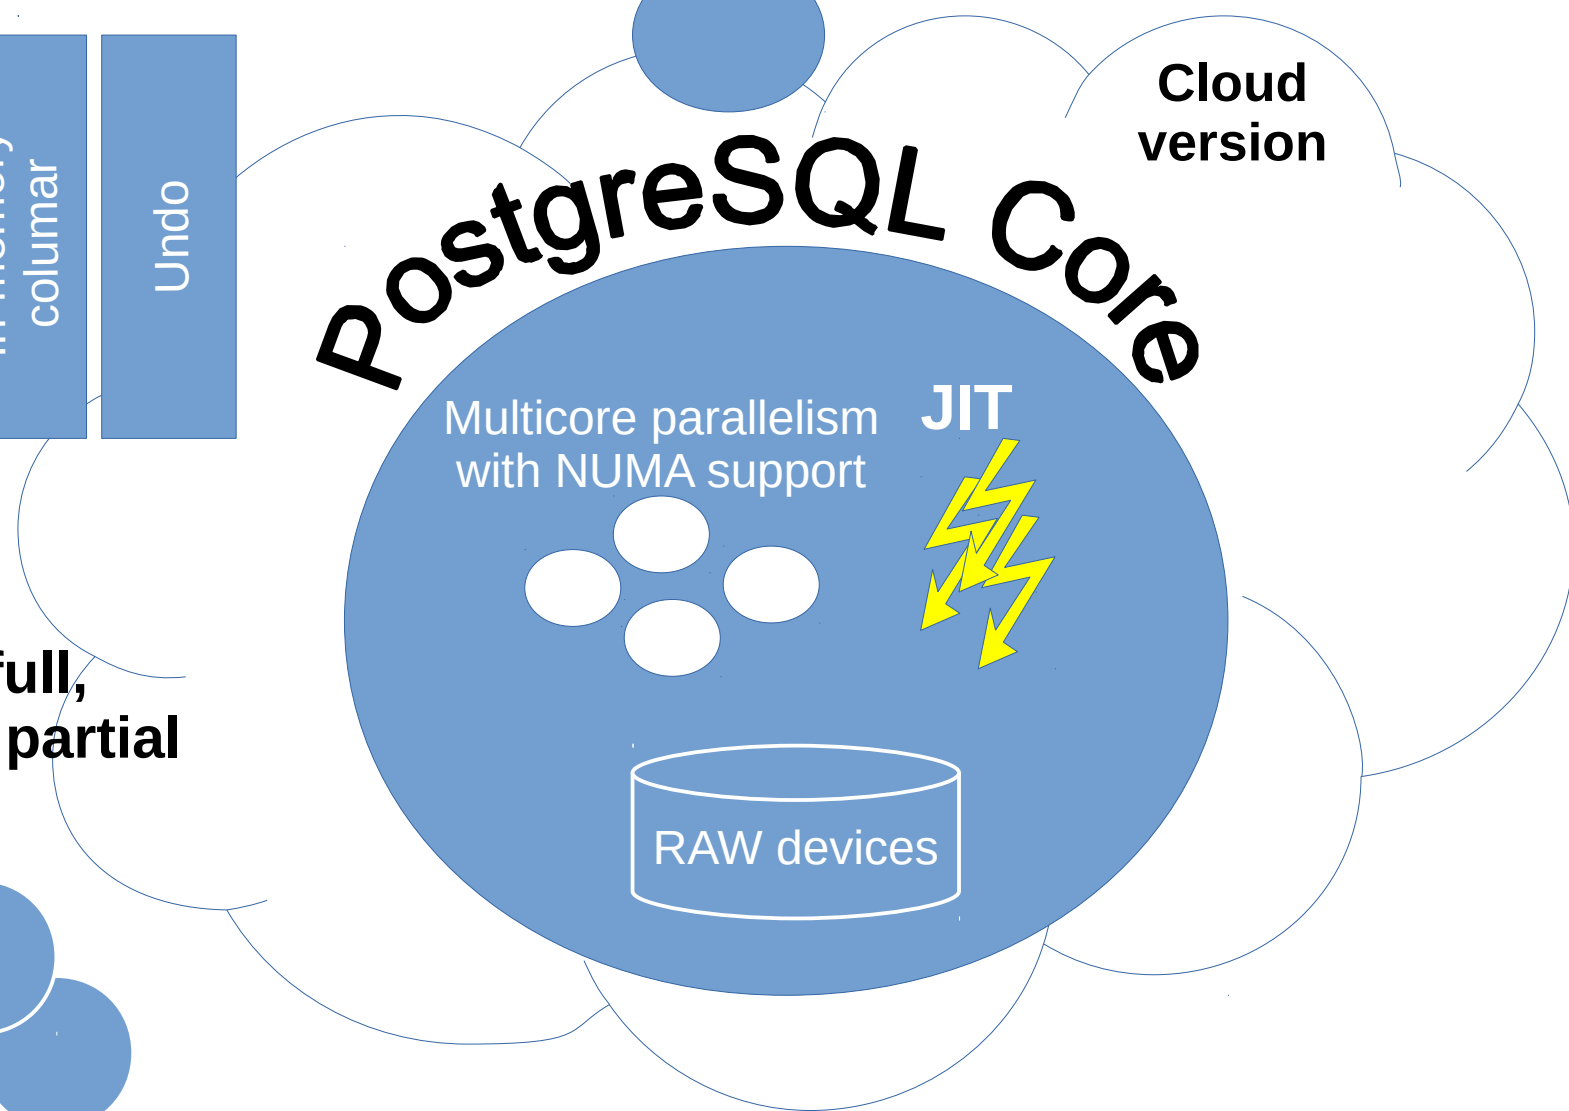
PostgreSQL Core

Multicore parallelism
with NUMA support

JIT



RAW devices





Thanks !

