

The driving force behind PostgreSQL

Bruce Momjian



<http://momjian.us/main/writings/pgsql/Alibaba-2017.pdf>

Who Am I?

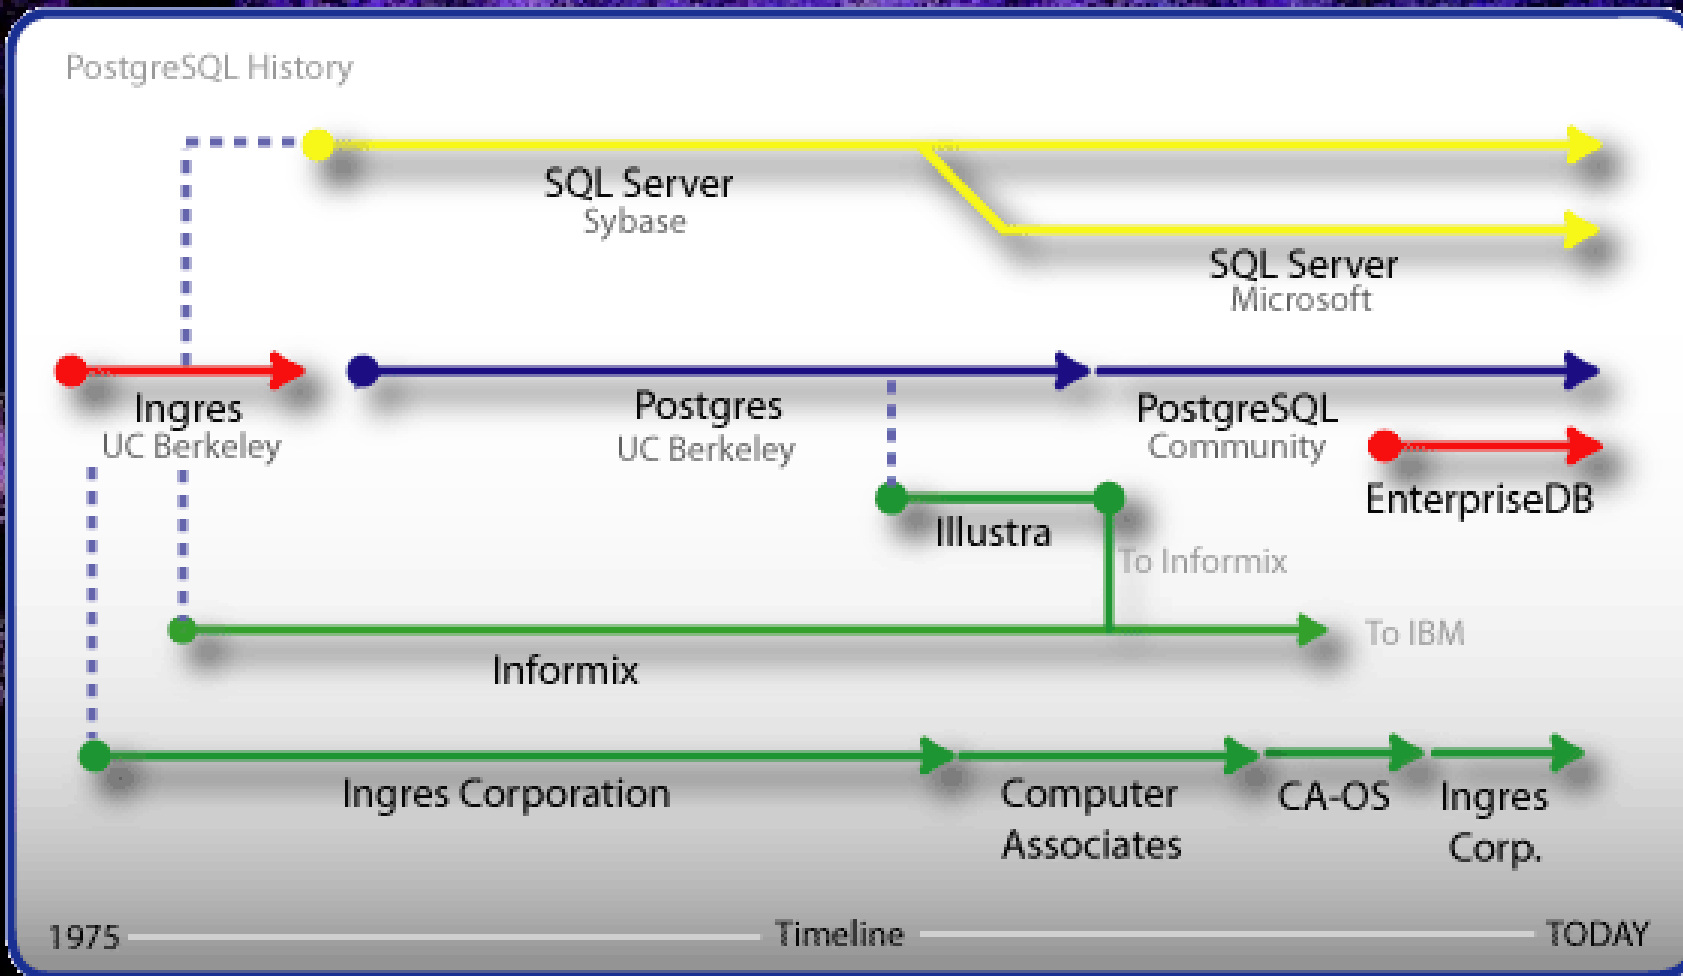
- Co-founder of Community PostgreSQL
- PostgreSQL core team member
- 21 years working on PostgreSQL
- Senior Database Architect at EnterpriseDB
- 11 years as an EDB employee



Agenda

1. PostgreSQL History
2. How open source benefits people and companies
3. Partnership between EnterpriseDB & PostgreSQL Community with Alibaba Cloud
4. Expectation for Open Source Databases

1. PostgreSQL History: Lineage



- Fully ACID Compliant
- MVCC
- Point in Time Recovery (PITR)
- Data and Index Partitioning
- Bitmap Indexes
- ANSI Constraints
- Triggers & Stored Functions
- Views & Data Types
- Nested Transactions
- Online Backup
- Online Reorganization
- Foreign Keys
- Streaming Replication
- Multi-Core Support
- JSON Support

PostgreSQL Version History

UCB
Postgres
[1986]

- PostgreSQL [1996]
 - V 6.0
 - Open Source

- PostgreSQL 8 [2005]
 - Microsoft Windows Native Server
 - PITR Tablespace

- PostgreSQL 9.4 [2014]
 - ALTER SYSTEM
 - Dynamic Background Workers
 - Replication Slots

- PostgreSQL 9.6 [2016]
 - Parallel Sequential Scan
 - Multiple Standbys in Synchronous Mode

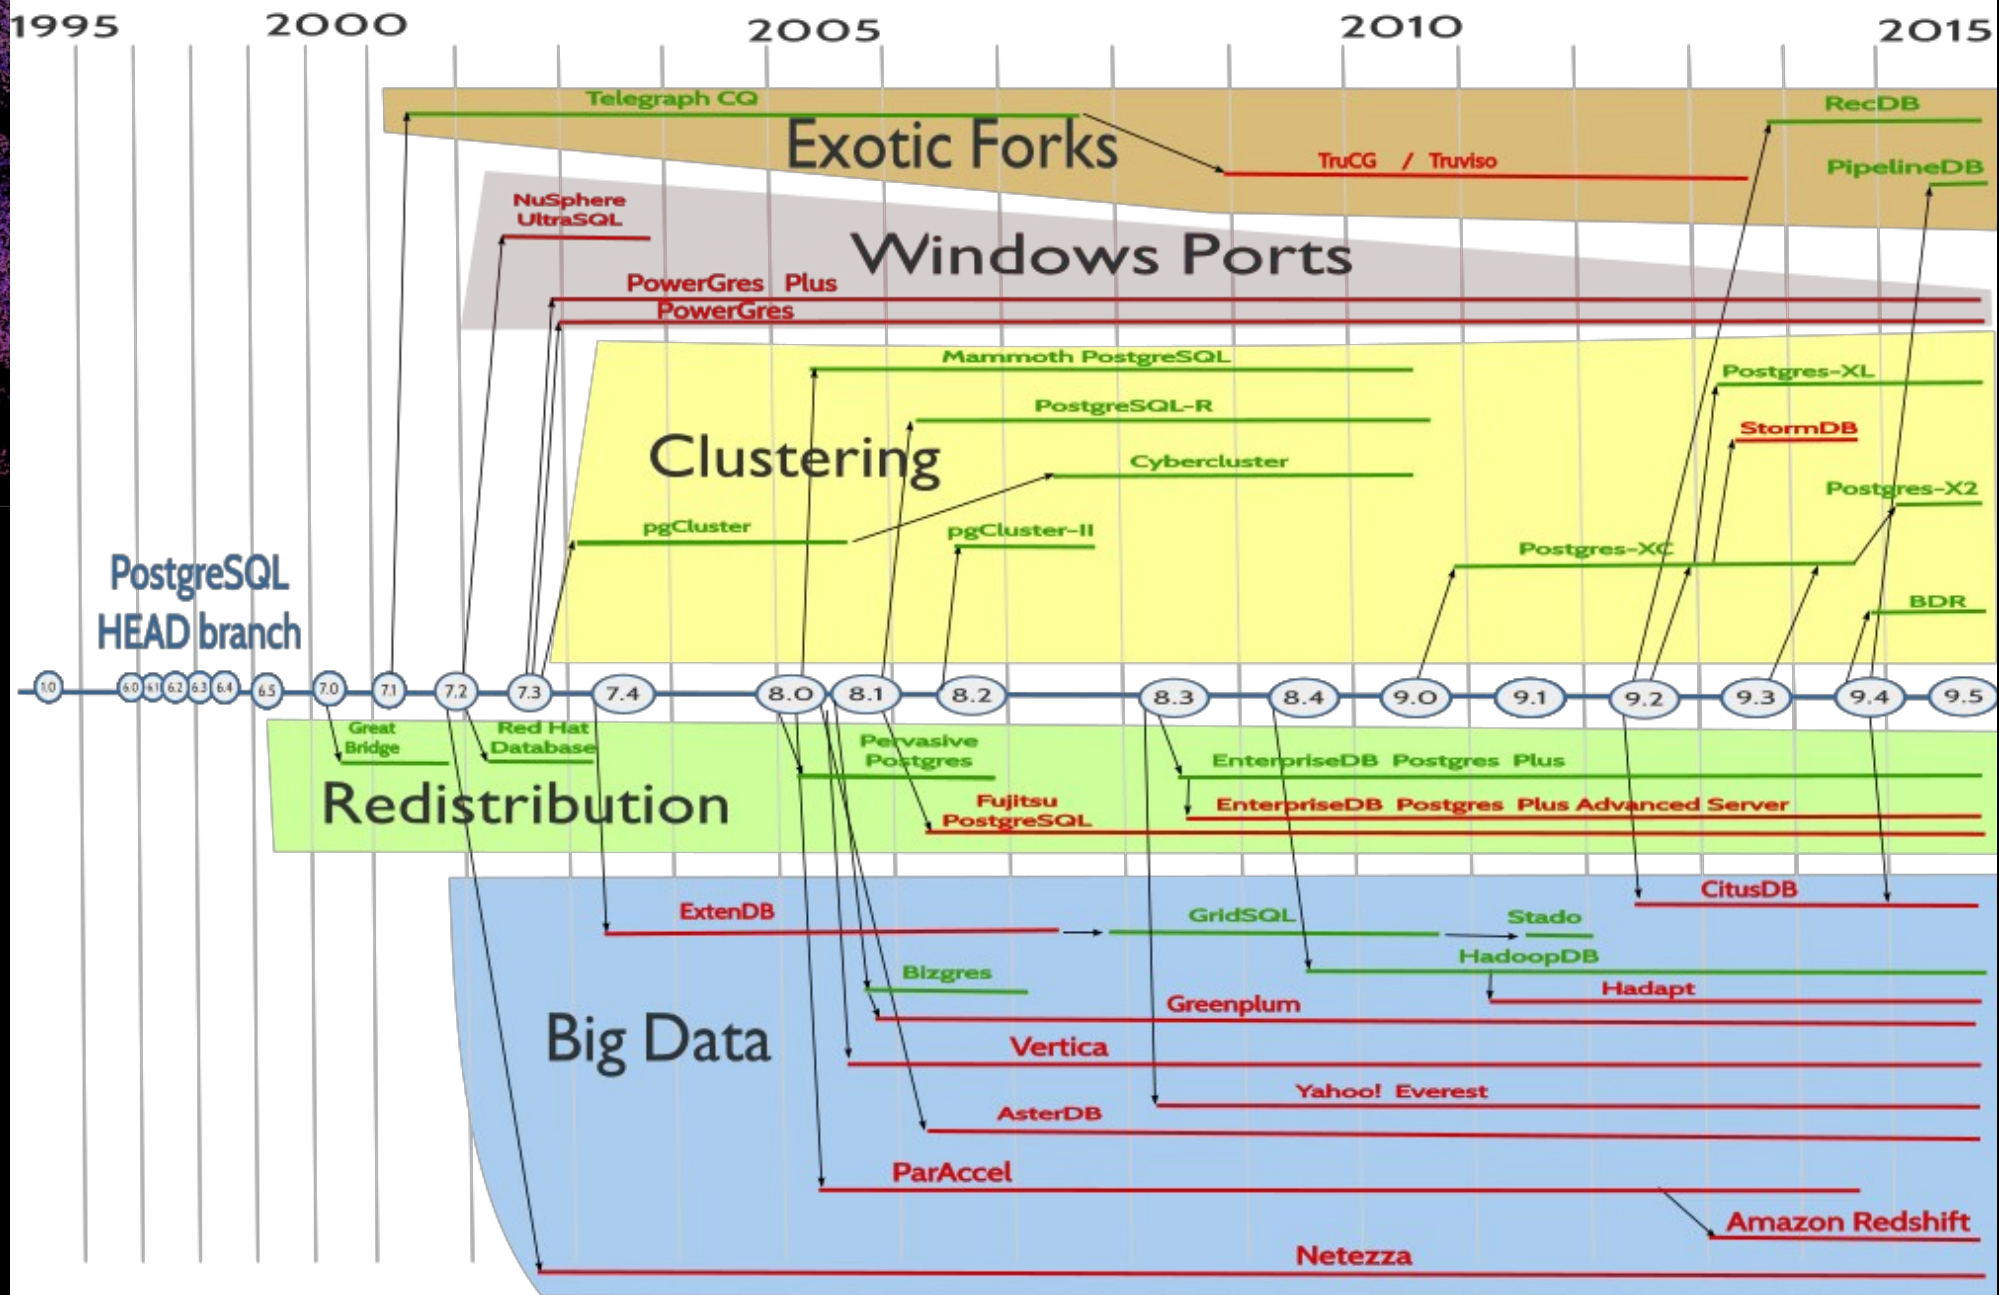
UCB
Postgres95
[1995]

PostgreSQL
7 [2000]

- PostgreSQL 9 [2010]
 - Streaming Replication
 - Major Performance Enhancements

- PostgreSQL 9.5 [2016]
 - Restrict Row level Access
 - BRIN Indexes
 - pg_rewind

the PostgreSQL Forks Timeline



2. How Open Source Benefits People and Companies


When the first survey launched 10 years ago, hardly anyone would have predicted that open source use would be ubiquitous worldwide just a decade later, but for many good reasons that's what happened. Its value in reducing development costs, in freeing internal developers to work on higher-order tasks, and in accelerating time to market is undeniable. Simply put, open source is the way applications are developed today.

Lou Shipley

President And CEO

Black Duck Software

<https://www.blackducksoftware.com/2016-future-of-open-source>



Advantages of Open Source

- Competitive features, innovation
- Freedom from vendor lock-in
- Quality of solutions
- Ability to customize and fix
- Cost
- Speed application development
- Reduce development costs
- Interoperability
- Breadth of solutions


<https://www.blackducksoftware.com/2016-future-of-open-source>

2016 Survey of Open Source Usage

- Operating Systems
- Database
- Development Tools

Database didn't appear in the top three the previous year's survey.

<https://www.blackducksoftware.com/2016-future-of-open-source>



Open Source Today


Open source today is unequivocally the engine of innovation; whether that's powering technology like operating systems, cloud, big data or IoT, or powering a new generation of open source companies delivering compelling solutions to the market.

Paul Santinelli

General Partner

North Bridge

<https://www.blackducksoftware.com/2016-future-of-open-source>



3. EDB Postgres Partnership with Alibaba


- Business Partners in Greater China since 2014
- EDB has been a solution in Alibaba's Public Cloud offerings for Greater China for three years
- License Partnership extended in 2017 with joint marketing model focused on:
 - Public Cloud Greater China
 - Public Cloud International Marketplace
- Exploring New Global Cloud DBaaS Partnership



Open Source Community Leader




Amit Kapila



Andres Freund



Bruce Momjian



Devrim Gunduz



Mohammad Usama



Thom Brown



Ashesh Vashi



Dave Page



Korry Douglas




Robert M Haas



Thomas Munro



Ashutosh Bapat



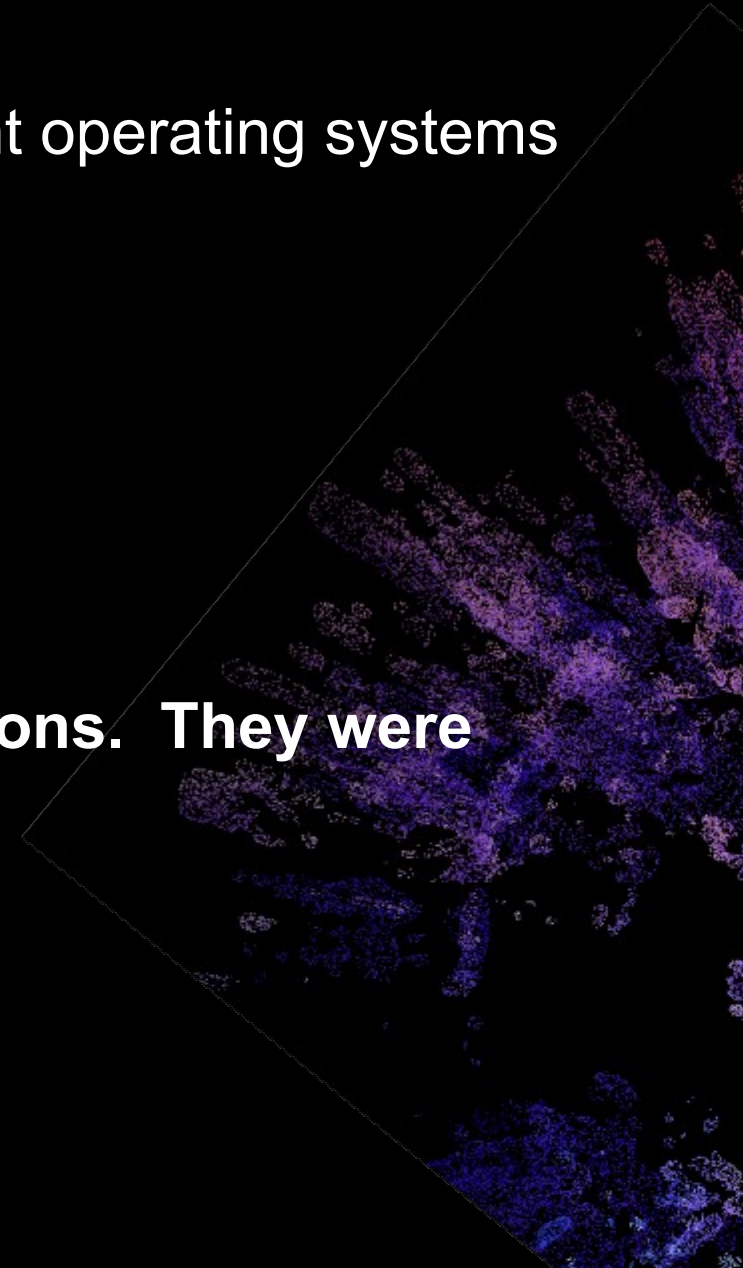
4. Expectations for Open Source Databases

When I started working on PostgreSQL in 1996, the dominant operating systems were:

- Solaris
- HP-UX
- AIX
- Windows NT

These are now gone or rarely deployed for new applications. They were produced by billion-dollar companies:


- Sun
- Hewlett Packard
- IBM
- Microsoft



Open Source Operating System Dominance

Open source operating systems now dominate:

- IT infrastructure, including cloud deployments
- Mobile
- Internet of Things (IOT)




Open Source Database Adoption Delay

Causes of the delay:

- New hardware often prompt new operating systems
- Few custom applications write to a custom operating system API
- Many custom applications write to a custom database API
- Database have varied workloads and require complex tuning for critical performance requirements
- Proprietary database systems quickly moved to open source operating systems





















Open Source Databases Today

The survey shows the open source database deployment has grown greatly in the last year.




DB Engines Ranking

131 systems in ranking, July 2017

Rank			DBMS	Database Model	Score		
Jul 2017	Jun 2017	Jul 2016			Jul 2017	Jun 2017	Jul 2016
1.	1.	1.	Oracle  	Relational DBMS	1374.88	+23.11	-66.65
2.	2.	2.	MySQL  	Relational DBMS	1349.11	+3.80	-14.18
3.	3.	3.	Microsoft SQL Server  	Relational DBMS	1226.00	+27.03	+33.11
4.	4.	4.	PostgreSQL  	Relational DBMS	369.44	+0.89	58.28
5.	5.	5.	DB2 	Relational DBMS	191.25	+3.74	+6.17
6.	6.	6.	Microsoft Access	Relational DBMS	126.13	-0.42	+1.23
7.	7.	7.	SQLite	Relational DBMS	113.86	-2.84	+5.33
8.	8.	8.	Teradata	Relational DBMS	78.37	+1.04	+4.43
9.	9.	9.	SAP Adaptive Server	Relational DBMS	66.91	-0.61	-3.82
10.	10.	10.	FileMaker	Relational DBMS	58.65	+1.57	+7.09
11.	11.	 13.	MariaDB 	Relational DBMS	54.36	+1.47	+18.56
12.	12.	12.	SAP HANA 	Relational DBMS	47.94	+0.45	+6.14
13.	13.	 11.	Hive 	Relational DBMS	46.20	+1.82	-1.34
14.	14.	14.	Informix	Relational DBMS	27.67	-0.16	-0.88
15.	15.	 16.	Microsoft Azure SQL Database	Relational DBMS	22.29	+0.96	+3.13
16.	16.	 15.	Vertica 	Relational DBMS	21.79	+0.88	+2.13
17.	17.	17.	Netezza	Relational DBMS	19.86	+0.20	+0.90
18.	18.	18.	Firebird	Relational DBMS	18.98	-0.16	+1.59
19.	19.	 23.	Impala	Relational DBMS	13.31	+0.51	+4.57
20.	20.	 19.	Amazon Redshift 	Relational DBMS	12.84	+0.31	+3.14

The Long Game



Conclusion

