

The History of PostgreSQL Open Source Development

BRUCE MOMJIAN



POSTGRES SQL is successful only because of its vibrant open source development community. This presentation chronicles the building of that community.

Creative Commons Attribution License

<http://momjian.us/presentations>

Last updated: January, 2017

Programming Before Relational Databases

- ▶ OS Files
- ▶ ISAM, indexed on a key
- ▶ No joins, WHERE clauses, or complex sorts

The University of California at Berkeley





Michael Stonebraker



Jolly Chen and Andrew Yu

PostgreSQL Code Base History

- ▶ Ingres — research prototype, spawned Relational Technologies, purchased by Computer Associates
- ▶ Postgres — research prototype, spawned Illustra, purchased by Informix
- ▶ Postgres95 — added SQL, spawned PostgreSQL

PostgreSQL Through the Years

1977-1985 Ingres

1986-1994 Postgres

1994-1995 Postgres95

1996- PostgreSQL

Transition from Single Maintainer To Internet Development

- ▶ Getting approval
- ▶ Scouring the mailing list
- ▶ Todo list
- ▶ Faq
- ▶ How does this thing work?

PostgreSQL Core Team




Top row: Thomas Lockhart, Jan Wieck, Tom Lane, Marc Fournier
Bottom row: Vadim Mikheev, Bruce Momjian

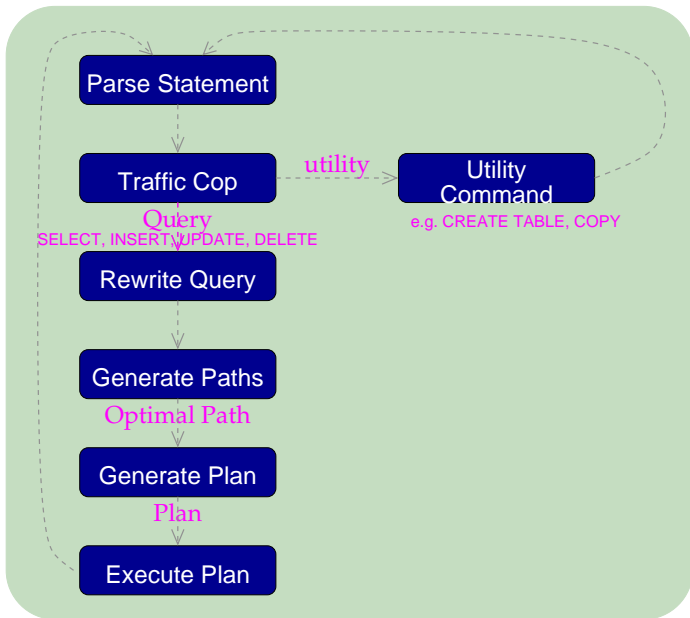
Assisting Our Developers

- ▶ Developer Education
 - ▶ Developer's Faq
 - ▶ Source code flowchart
 - ▶ Emphasize code clarity and people will respond
- ▶ Developer Recruitment
- ▶ Cvsup, remote Cvs
- ▶ Managing an increasing installed base

Internals Flowchart




Internals Flowchart (magnified)



Keeping the Code Clean

- ▶ automatic formatting
- ▶ remove unused functions
- ▶ mark single-file functions as static
- ▶ rename variables and functions
- ▶ require patches to be of acceptable quality

Developer's Globe



Daily Chores

- ▶ Mailing list
 - ▶ Answer simple questions
 - ▶ Identify valid bug reports
 - ▶ Work with users to test bug fixes
 - ▶ Collect feature requests
 - ▶ Encourage technical discussions
- ▶ Patch application
- ▶ Encourage
- ▶ Marketing

Packaging a Release

- ▶ Collect unapplied patches
- ▶ Resolve outstanding issues
- ▶ Regression testing
- ▶ Enter beta
- ▶ Resolve all bug reports
- ▶ Final release
- ▶ Collect and apply safe patches
- ▶ Issue subreleases

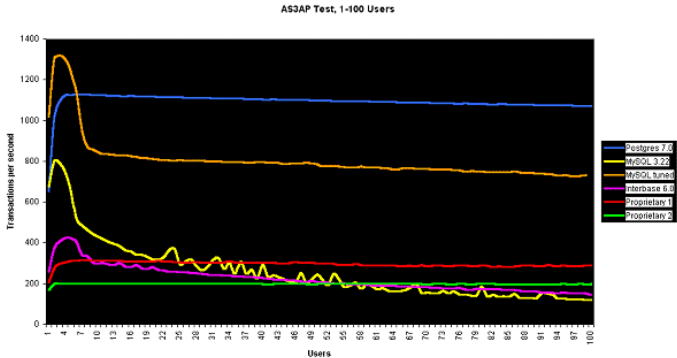
Release Dates and Sizes

Date	Release	Lines of code
1994		244,581
1996-08-01	1.02.1	
1996-10-27	1.09	178,976
1997-01-29	6.0	
1997-06-08	6.1	200,709
1997-10-02	6.2	225,848
1998-03-01	6.3	260,809
1998-10-30	6.4	297,918
1999-06-09	6.5	331,278
2000-05-08	7.0	383,270
2001-04-13	7.1	410,500
2002-02-04	7.2	394,274
2002-11-27	7.3	453,282
2003-11-17	7.4	508,523
2005-01-19	8.0	654,437
2005-11-08	8.1	630,422
2006-12-05	8.2	672,570


Development Directions

- ▶ Features
- ▶ Performance
- ▶ Reliability
- ▶ Code simplicity
- ▶ Standards conformance
- ▶ Documentation
- ▶ Pace of releases

Performance - AS3AP



Performance - TPC-C



Conclusion



Top row: Thomas Lockhart, Jolly Chen, Vadim Mikheev
Jan Wieck, Andrew Yu, Tom Lane
Bottom row: Bruce Momjian, Marc Fournier