

**PostgreSQL - Synergizing an agile,  
customer-focused database with the needs  
of a Web 0.2beta-savvy community**

- or -

## **Some stuff about PostgreSQL**

17<sup>th</sup> October 2008

Dave Page

***Enterprise*DB™**

Magnus Hagander

Redpill  
Linpro 



# POSTGRES

- Originated at UC Berkeley under Dr Michael Stonebraker in 1986
- Successor to Ingres – hence POST-GRES, meaning after Ingres – but no code shared
- A research project to explore new ideas, such as ‘object relational’ technologies
- Spun off commercial database Illustra, now part of Informix



## Postgres95

- In 1994, Andrew Yu and Jolly Chen renamed the project, and replaced the POSTQUEL query language with SQL
- Converted to ANSI C, compiled with unpatched GCC
- Released as BSD licensed open source code, but with no real outside development



# PostgreSQL

- Postgres95 was discontinued at Berkeley in 1996
- The PostgreSQL project was born on July 8, but retained the Postgres95 name initially
- Project started by Marc Fournier, Bruce Momjian and a gang of interested hackers



# PostgreSQL

- CVS & Mailing list set up
- Existing patches applied to the new CVS
- Developers start to attack bugs and learn the code
- PostgreSQL 6.0 was released on 29<sup>th</sup> Jan 1997



## Early development

- No one really understands all the code yet
- Bugs are tackled one at a time
- Initial confusion over process - what do I do with my patch? Has anything happened with it?
- Bruce starts gathering bug reports, and managing the process



## Core team is created

- Project organization modeled on FreeBSD
- Separate developer, committer and core roles



## PostgreSQL 6.5

- Multi Version Concurrency Control (MVCC)
- Hot Backup
- Numeric datatype
- Temporary tables





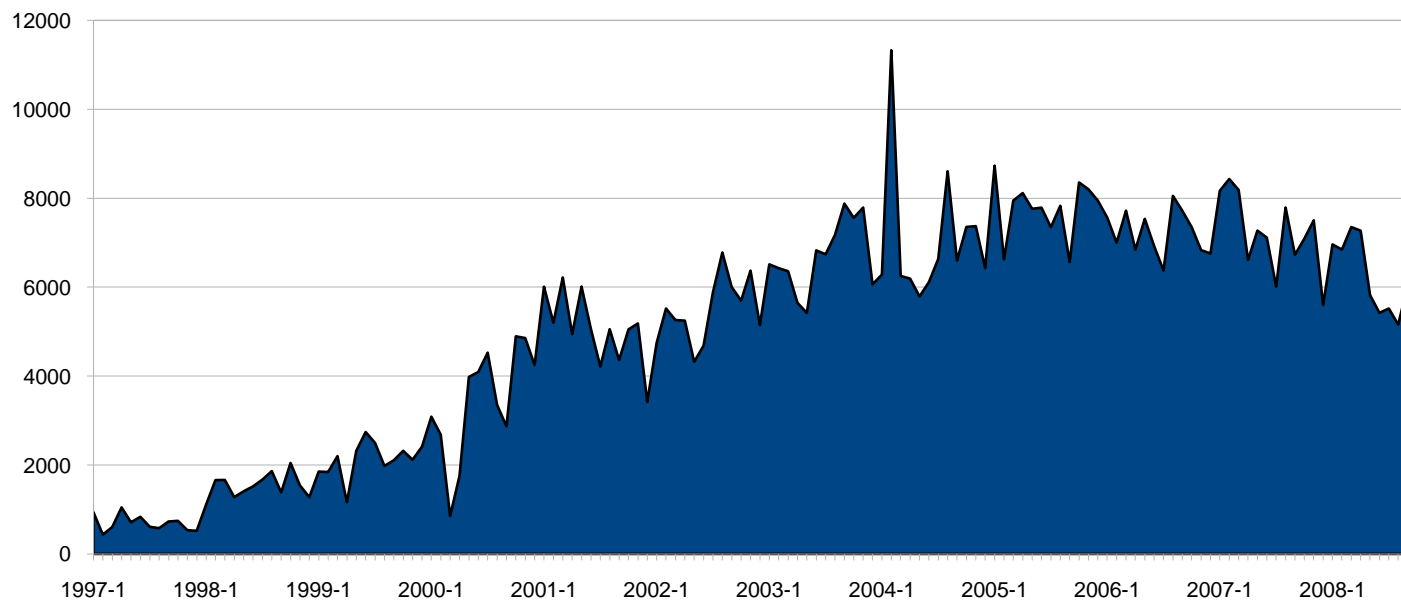
## Japan PostgreSQL User Group (JPUG)

- Formed July 23, 1999 with 942 initial members(!)
- Much work on translation and enterprise adoption
- Incorporated as a non-profit organisation in 2005
- Present: 2300 members, 51 staff and 24 directors
- Regular meetings and conferences, often with guest speakers from around the world.



# The community expands

- Additional mailing lists are added for specialist areas
  - pgsql-advocacy
  - pgsql-www
  - pgsql-performance



## Working groups

- Working groups form to undertake specific tasks
  - Web and sysadmin teams
  - Translators
  - Press contacts
  - Packagers



## PostgreSQL 7.1

- PostgreSQL 7.1 released on 13<sup>th</sup> April
- Many cool new features
  - WAL Logging
  - TOAST (The Oversized Attribute Storage Technique)
  - Outer joins
  - Full 64 bit support



# San Francisco PostgreSQL User Group

- Started in 2003 by Josh Berkus and David Fetter, now with 150 members
- 10 meetings per year in San Francisco, Oakland and South Bay
- Meetings usually centre around a presentation on hacking, apps, forks etc
- Booth and party at LinuxWorldExpo every year, this August hosting a full one-day mini-conference with over 100 attendees



## Getting serious about fundraising

- The PostgreSQL Foundation is formed with a board of directors and initial membership
- Intended to provide a vehicle for managing fund raising, expenditure and legal issues
- Failed after over a years worth of work due to problems attaining 501(c)(3) (non-profit) status from the US IRS
- The Foundation becomes the 'funds group', and looks at other alternatives to incorporation



## PostgreSQL 8.0

- Milestone release with many “enterprise features”
  - Windows support
  - Point In Time Recovery (PITR)
  - Savepoints
- Dedicated mailing lists needed for really large changes:
  - [pgsql-hackers-win32](#)
  - [pgsql-hackers-pitr](#)



## Next step for the Funds Group

- Decide to join Software in the Public Interest (SPI):
- Funds group members take an active role in SPI, with two members joining the board
- Other projects include Debian, Drupal, freedesktop.org, madwifi.org and GNUstep





## PostgreSQL's 10<sup>th</sup> Birthday

- Celebrated in style with a 2 day conference in Toronto on July 8/9 2006
- The first time most developers ever met in person
- Beer and food a-plenty!
- Culminated in a day long 'code sprint', with 30+ hackers working together at the Novotel Hotel



## Growing pains

- The development process starts to suffer as the project grows
- Patch authors can wait for months without feedback whilst their patches 'bit rot'
- The Feature Freeze period prior to release becomes longer and longer, nearing half of the entire development cycle



# Conference explosion

- PGCon (Ottawa) – 2007 & 2008
- PGDay (Prato, Italy) – 2007 & 2008
- PG:West (Portland) – 2007 & 2008
- PG:East (Maryland) – 2008



## More user groups

- Both regional and language-specific user groups are being formed independently:
  - PostgreSQLFr (French speaking, NPO)
  - ITPUG (Italian, NPO)
  - PostgreSQL.gr (Greek)
  - PostgreSQL PE (Peru)
  - Pgug.de (Germany)
  - and many more!



## PGEU formed

- PGEU planned at PGDay.IT 2007, at the PLUG bar
- Directors voted in at FOSDEM 2008
- Legally recognised as a non-profit organisation on April 26
- <http://www.postgresql.eu>



## PGEU – what's it for?

“The association has as its objectives the promotion, the support for deployment and the development of the open-source software PostgreSQL and its derivatives (and related projects) in Europe.”




# PGUK

- A small, ad-hoc group
- One conference, held on 2<sup>nd</sup> April 2008 in Birmingham
- Risk-free conference – held as part of the much larger UK Unix Users Group (UKUUG) conference
- Anyone can start a user group!
- <http://www.postgresql.org.uk/>



# The source tree is growing

- By January 2008, we have over 800,000 lines of code!





## Development meltdown!

- The development process becomes unworkable
- Some developers wait nine months or more for feedback
- Patches bit-rot so much they require significant reworking
- Developers become disillusioned and fed up, moving on to other projects or jobs
- The Core team and others work on a solution...



## Commit fests, new reviewers and collaborative management!

- The review process is documented and made easier for new, or inexperienced reviewers
- Every two months we hold a commit fest:
  - Patches are listed on the current commit fest wiki page.
  - Every two months, a commit fest starts:
    - Development stops
    - The commit fest manager recruits reviewers
    - All patches are reviewed and either committed, or returned with feedback



## Why does it work?

- Developers receive prompt feedback
- Patches don't bit-rot
- The simplified reviewing process has led to additional reviewers, easing the workload of the committers
- The process doesn't 'bottleneck' on a single person managing the patch queue



# PGUS

- A 'sister' organisation to PGEU
- The US PostgreSQL Association is a non profit corporation with the following purpose:
  - Educate, promote and support the creation, development and use of the PostgreSQL Open Source Database software, a software system which is available to the general public without charge;
  - Provide information and education regarding the use of PostgreSQL; and
  - Organize, hold and conduct meetings, discussion, and forums on the contemporary issues concerning the use of PostgreSQL.



## PGDay.EU

- The first example of PGEU working with a regional user group
- ‘On the ground’ and Italian track organisation by ITPUG
- English track organised by PGEU
- A success? We’ll soon find out!





## What's next?

- Further refinement of the development process
- More conferences – East/West series, PGDay, PGCon as well as smaller regionals
- Bigger presence at other conferences – LinuxLive, FOSDEM, OSCON
- PostgreSQL 8.4 and beyond





## PostgreSQL 8.4 – committed patches

- Free Space Map rewrite
- Case insensitive text type
- Variadic function support
- CASE in PL/PGSQL
- Per database collation
- libpq events
- Common Table Expressions (recursive queries)

plus much, much more...





## PostgreSQL 8.4 – work in progress

- Parallel restore
- Hot standby slave servers
- Synchronous replication
- Security Enhanced PostgreSQL
- Window functions
- IO improvements on RAID (posix\_fadvise)
- In place upgrade
- Column level privileges

plus much, much more....





Thank you!

