



# PostgreSQL

**What makes this database so powerful?**

GoOpen2008

Rafael Martinez, USIT, UiO

`r.m.guerrero@usit.uio.no`



## **Center for Information Technology Services, University of Oslo**

200 employees

A turnover of over 210 million kroner.

2/3 - Infrastructure and IT-services for the University of Oslo

1/3 - IT-services for other norwegian universities and the education sector


### **DBA department:**

186 Oracle databases, several TB of data

126 PostgreSQL databases / 235 GB of data

PostgreSQL -> Average of almost 18,000,000 transactions /day


**<http://www.usit.uio.no/>**



# PostgreSQL

## History





PostgreSQL

## **Ingres 1977-1985 - *The beginning***

Proof of concept for relational databases.

Michael Stonebraker, professor at Berkeley, California.

Ingres -> NonStop SQL, Sybase -> Microsoft SQL server

## **Postgres 1986-1994 - *As in "after Ingres"***

A project meant to break new ground in database concepts.

"Objects relational" technologies.

Commercialized to become Illustra.



## **Postgres95 1994-1995 - *New life in the OpenSource***


Two Ph.D. students from Stonebraker's lab, Andrew Yu and Jolly Chen started Postgres95.

Departed from academia to a new life in the open source world with a group of dedicated developers outside of Berkeley.


Establishment of the *PostgreSQL Global Development Team*.

Released as PostgreSQL 6.0 in 1996.

## **PostgreSQL 1996-today - *PostgreSQL project***



# PostgreSQL



**"Crash"**

**SQL -standards**

**Enterprise - 24/7**


Multiversion  
Concurrency Control  
(MVCC)  
Important SQL features  
Improved build-in types  
Speed

Improved performance  
Improved administration &  
maintenance  
24/7 ready



# PostgreSQL

## PostgreSQL Source code



**Total Physical Source Lines of Code PG-8.3.0 (SLOC) = 814,787**


**Development Effort Estimate, Person-Years (Person-Months) = 227.83 (2,734.05)**

**Schedule Estimate, Years (Months) = 4.21 (50.57)**

**Estimated Average Number of Developers (Effort/Schedule) = 54.06**

**Total Estimated Cost to Develop (Avg.salary: \$70,000/year, overhead: 2.40) = \$38,238,854**

**REF: Basic COCOMO (Constructive COst MOdel for software cost estimation model)**




PostgreSQL

## Features

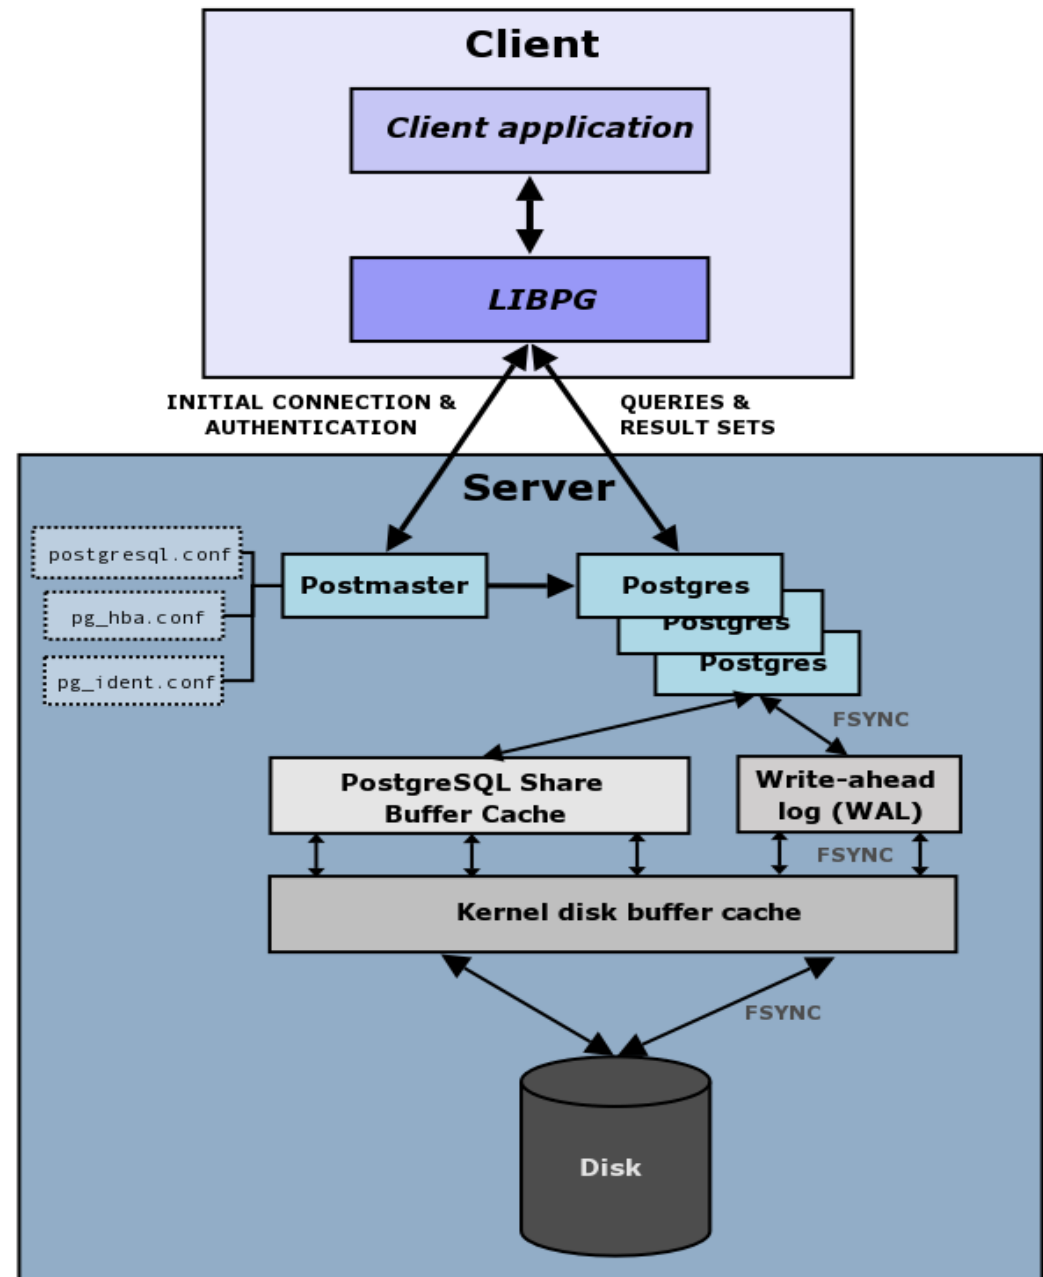
<http://www.postgresql.org/about/featurematrix>





# PostgreSQL

It uses a multi-process model  
It does not use multi-threading.





## **General features**

Fully ACID compliance (Atomic, Consistent, Isolated, Durable)

ANSI-SQL 92/99/2003 compliance

Referential integrity

Multi-version concurrency control (MVCC)

Write-Ahead logging (WAL), REDO recovery

Point-in-time recovery PITR / Online backups

Replication

Tablespaces

Savepoints, two-phase commits

Functional and partial indexes

B-tree, R-tree, Hash, GiST and GIN index types

Full text search

Native SSL, Kerberos, GSSAPI and SSPI support

**Linux, UNIX (AIX, BSD, HP-UX, SGI, IRIX, Mac OS X, Solaris, Tru64), Windows.**




## Development features

Stored procedures, PL/pgSQL, PL/Perl, PL/Python, PL/Tcl, PL/php, PL/java, PL/R, PL/Ruby, PL/sh, ...

Native interfaces for ODBC, JDBC, C, C++, PHP, Perl, TCL, ECPG, Python, Ruby, Lisp, Scheme, Qt, .Net, OpenOffice SDBC, ...

User defined data types, functions and operators, SPI  
Open and documented API.



PostgreSQL

## SQL features

Rules  
Views  
Triggers  
Cursors  
Sequences

Inheritance  
Outer joins  
Sub-selects  
Unicode  
SQL/XML standard



## **Some upcoming features**

- Auto-tuning / auto-configuration
- Easy upgrade-in-place - 'pgmigrator'
- More SQL99 and SQL2003 features
- More OLTP performance enhancements
- Auto partitioning / Dynamic partitioning
- External tables interfaces (SQL/MED compliant)
- More exotic datatypes
- More query optimizer improvements
- Faster vacuum with reduced impact
- Improved XML support



## Support / sponsors

**EnterpriseDB™**



**OmniTI**

**UNISYS**


**FUJITSU**



**2ndQuadrant +**

**SRA**





PostgreSQL

## **So why is this database so powerful?**

Open source project and quality source code

Immunity to over-deployment – BSD license

Professional support

Low maintenance and tuning requirements

Reliability and stability

Excellent performance

Designed for high volume environments

Extensible


Cross platform

Command line & GUI database design and administration tools



**<http://www.postgresql.org/>**





# PostgreSQL

## References

- [1] PostgreSQL documentation, 8.0.x online manual - <http://www.postgresql.org/docs/>
- [2] PostgreSQL (second edition), *Korry Douglas / Susan Douglas*, Developer's library.
- [3] **Beginning databases with PostgreSQL (second edition)**, *Richard Stones & Neil Matthew*
- [4] PostgreSQL mailing lists, <http://www.postgresql.org/community/lists/>
  
- [5] History of PostgreSQL – presentation, *Bruce Momjian* - <http://candle.pha.pa.us/>
- [6] PostgreSQL Performance tuning – presentation, *Bruce Momjian* - <http://candle.pha.pa.us/>
- [7] Mastering PostgreSQL administration – presentation, *Bruce Momjian* - <http://candle.pha.pa.us/>
- [8] Data processing inside postgresql– presentation, *Bruce Momjian* - <http://candle.pha.pa.us/>
- [9] PostgreSQL internals through pictures – presentation, *Bruce Momjian* - <http://candle.pha.pa.us/>
- [10] PostgreSQL replication solutions – presentation, *Bruce Momjian* - <http://candle.pha.pa.us/>
- [11] PostgreSQL: Past, Present, and Future – presentation, *Bruce Momjian* - <http://candle.pha.pa.us/>
- [12] PostgreSQL hardware performance tuning – article, *Bruce Momjian* - <http://candle.pha.pa.us/>
- [13] Get to know PostgreSQL – presentation, *Oddbjørn Steffensen* - <http://www.tricknology.org/foilware/>
- [14] Power PostgreSQL - <http://www.powerpostgresql.com/>

