

Selling Your Boss On Open Source *formerly PostgreSQL Adoption and Trends*

BRUCE MOMJIAN,
SOFTWARE RESEARCH ASSOCIATES

June, 2004




Abstract

This explores the adoption of Open Source, using PostgreSQL as an example.

Introduction

- How is open source development different from closed source?
- What is the pattern of PostgreSQL adoption?
- What does the future hold for PostgreSQL?

Open Source World



Revolutionary Change in Software Development: Challenging Totalitarianism

- Central Control
- Few Choices
- Limited feedback from population

Bringing Democracy to Software Development

- Distributed Control
- Many Choices
- Driven by feedback from population

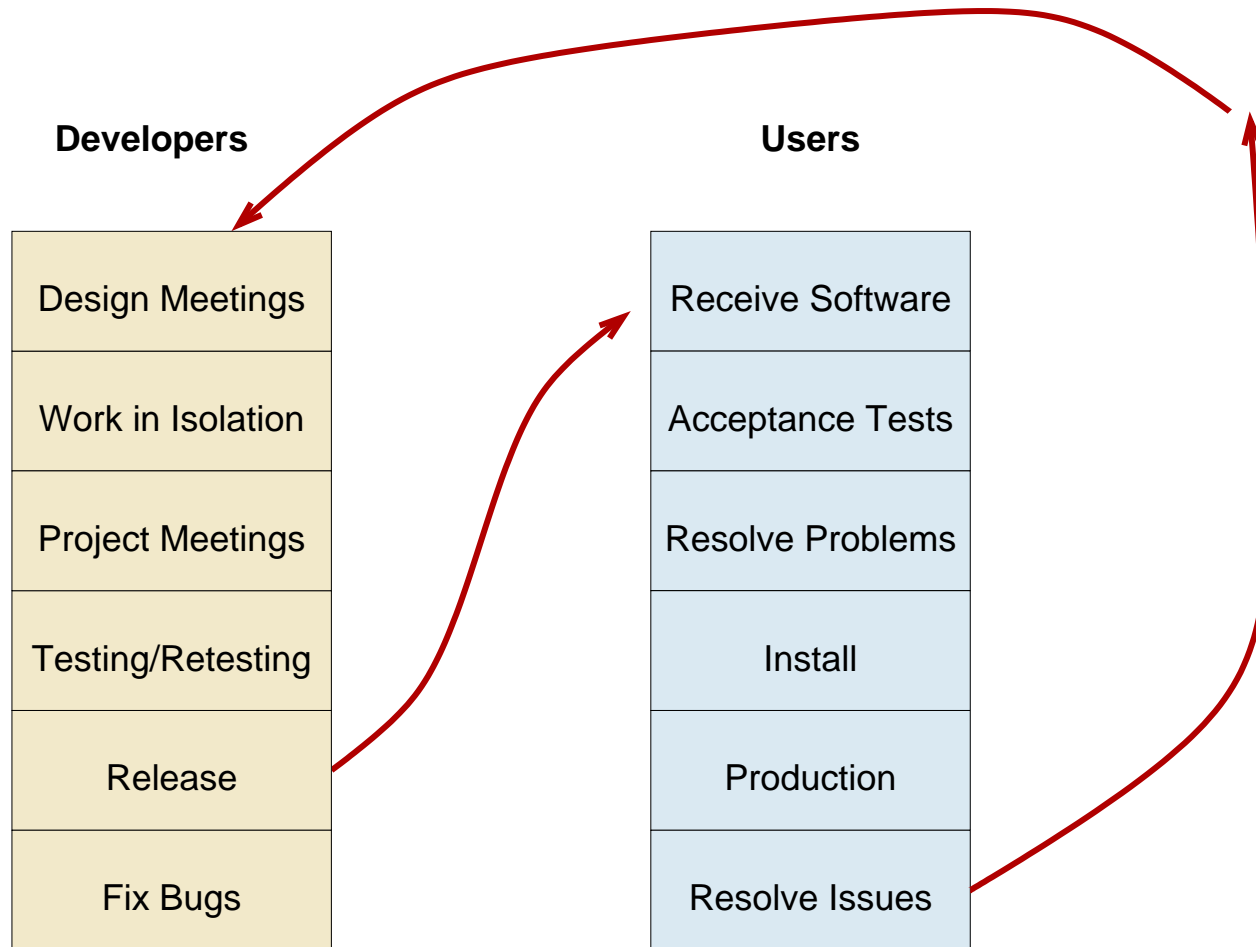
Proprietary Software Development

- Put everyone in a room
- Discuss new features and how to implement them
- Everyone goes back to their office
- More meetings
- Testing
- Release
- Gather user bug reports
- Repeat

Open Source Development

- Discuss via email, with access to a world-wide pool of talent
- New features discussed over several days by developers and users
- Coding
- Patch review
- Patch application and testing
- User testing
- Release

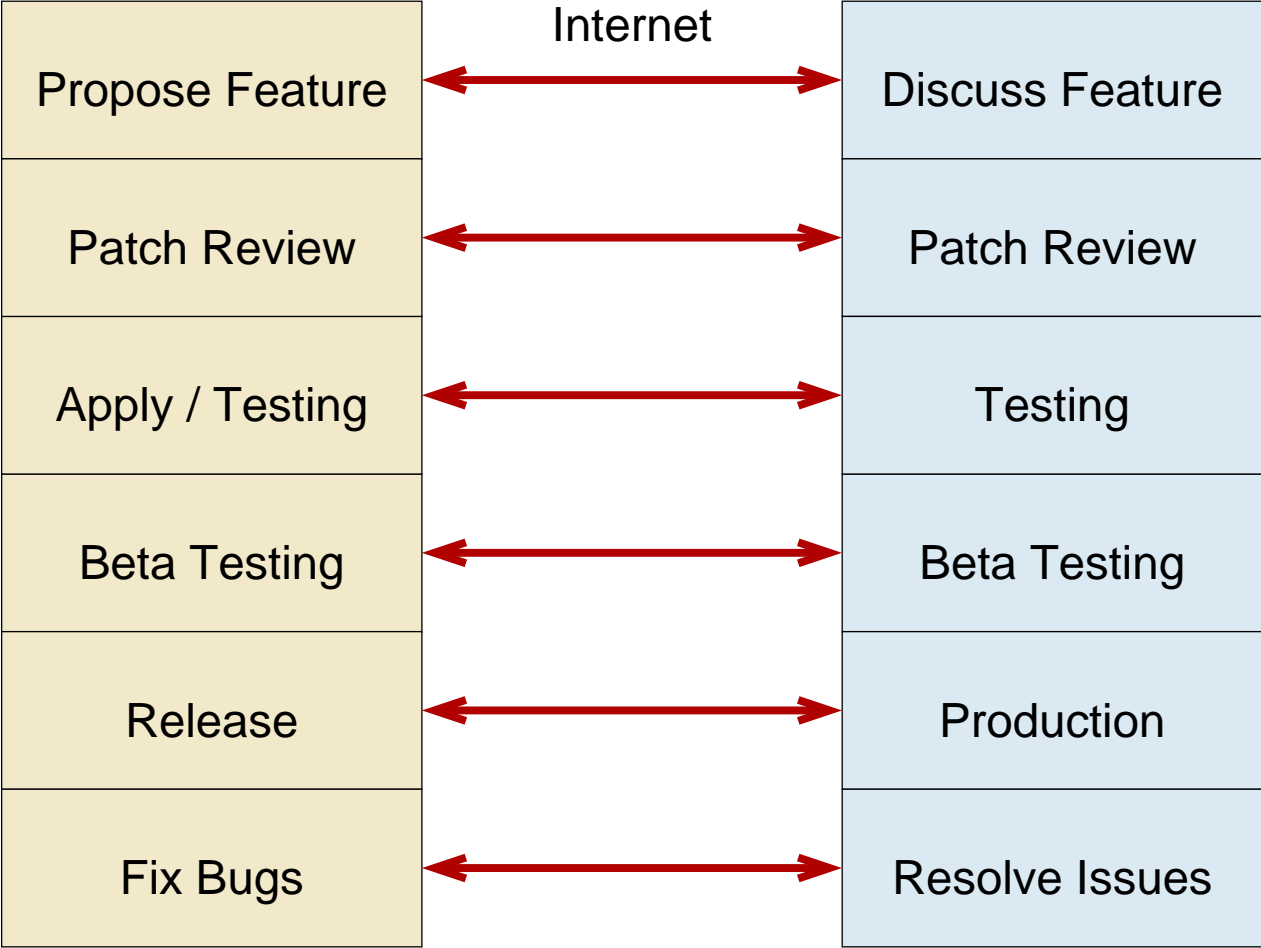
Proprietary Software Development



Open Source Development

Developers

Users



Typical Proprietary Database Product Life Cycle

- Create product
- Grow customer base
- Reduce development budget
- Decrease technical support costs
- Increase license fees
- Maximize revenue
- Customers leave
- End-of-life product

All possible because of high switching costs.

Typical Open Source Life Cycle

- Create project
- Developers create initial software
- Users and developers improve the software
- No development budget
- No technical support costs
- No license fees
- No maximizing of revenue
- No customers leaving!
- No end-of-life

Uniqueness of PostgreSQL as an Open Source Project

- Linux - no single gatekeeper, project is usable without enhancement
- Mozilla - no company history like AOL/Netscape
- Open Office - no controlling company like Sun
- Gnome - no controlling companies
- PHP - no Zend steering development
- Sendmail - no control by Sendmail, Inc.
- MySQL - no MySQL AB that does all server development

Similar Projects

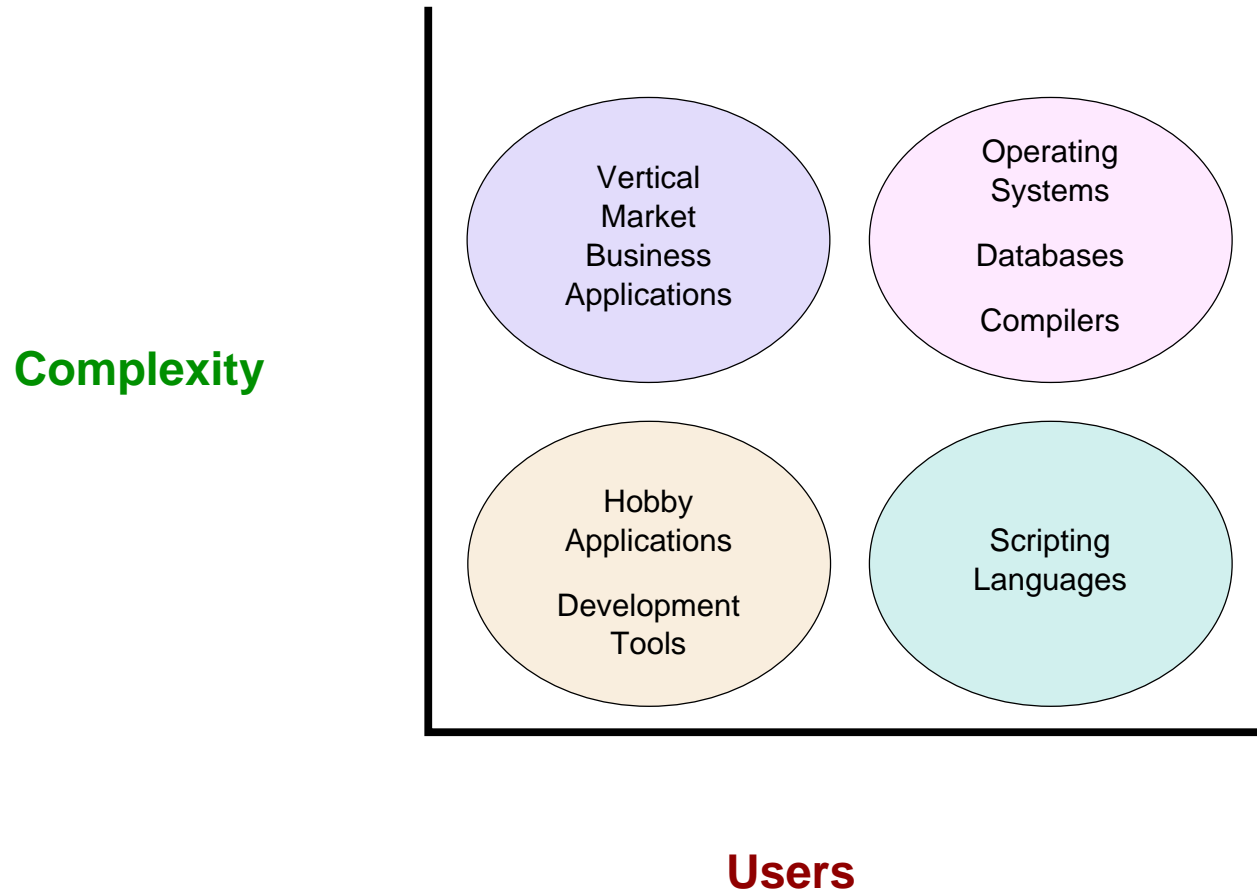
- FreeBSD - community project
- Samba - active developer community

Open Source Development Structures


- Distributed Leadership: PostgreSQL, FreeBSD, NetBSD, Samba
- Central Leadership: Linux, Perl, Python, OpenBSD
- Corporate Council: Gnome, XFree86
- Corporate Sponsored: Apache, PHP
- Corporate Controlled: MySQL, OpenOffice, Sendmail
- Single Developer: Many projects

Concepts taken from Josh Berkus, PostgreSQL core member.

Open Source Software Availability



Is It As Good As Proprietary Software?



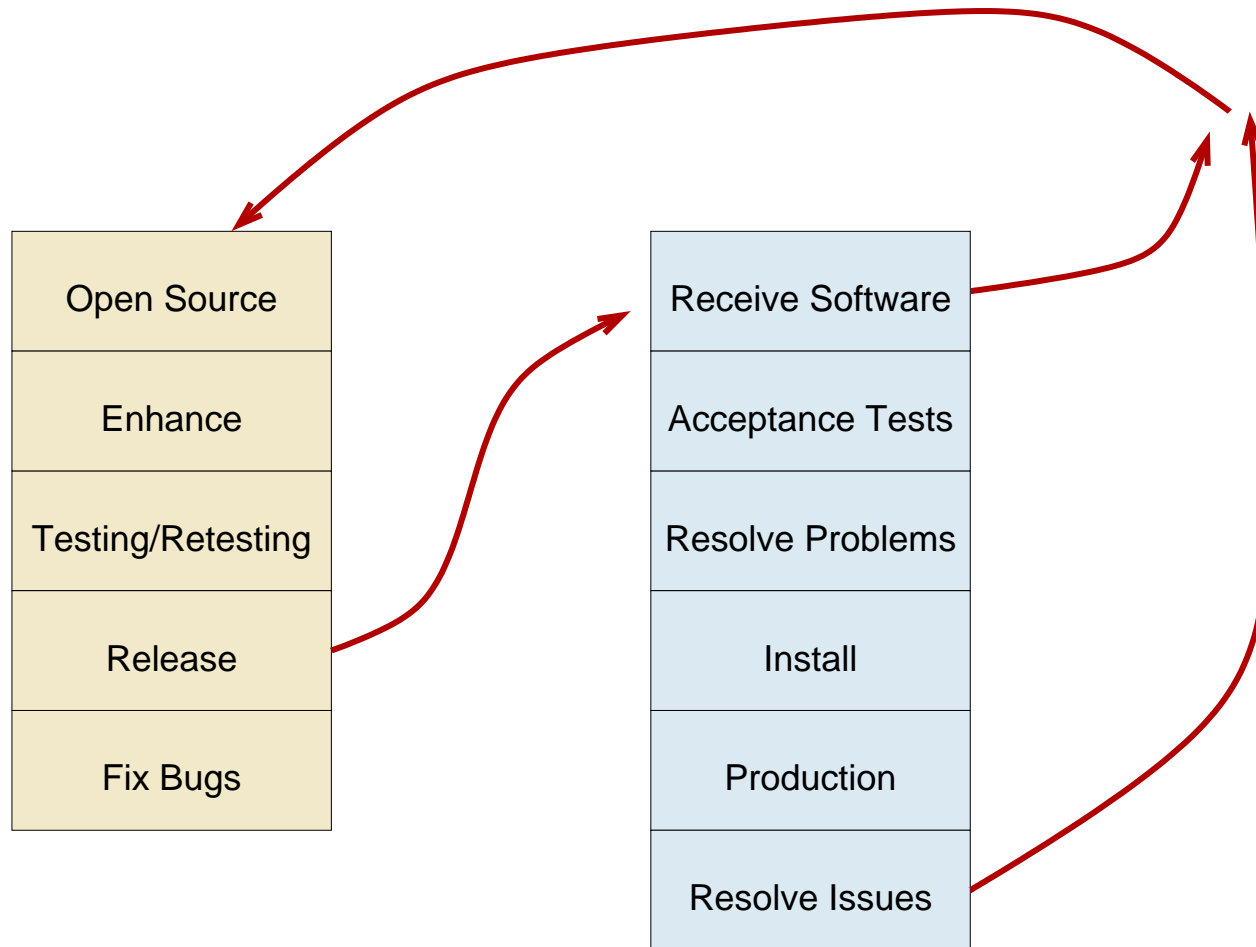
Problems With Open Source (and Democracy)

- Requires effort to stay informed
- Requires effort to receive assistance
- Documentation is not as thorough
- Database changes more rapidly
- More self-help

Businesses Supporting Open Source Users

- Customization
- Technical support
- Packaged releases
- Add-on software
- Training

Hybrid Development



Support Companies

- USA
- Canada
- Mexico
- Brazil
- France
- Spain
- Germany
- Austria
- Norway
- Israel
- Turkey
- Russia
- India
- Japan
- Australia

Value-Added Reseller (VAR) Advantages

- License
- No unsolvable problems
- Control database product life-cycle

End-user Advantages

- License costs
- Software control
- Customization

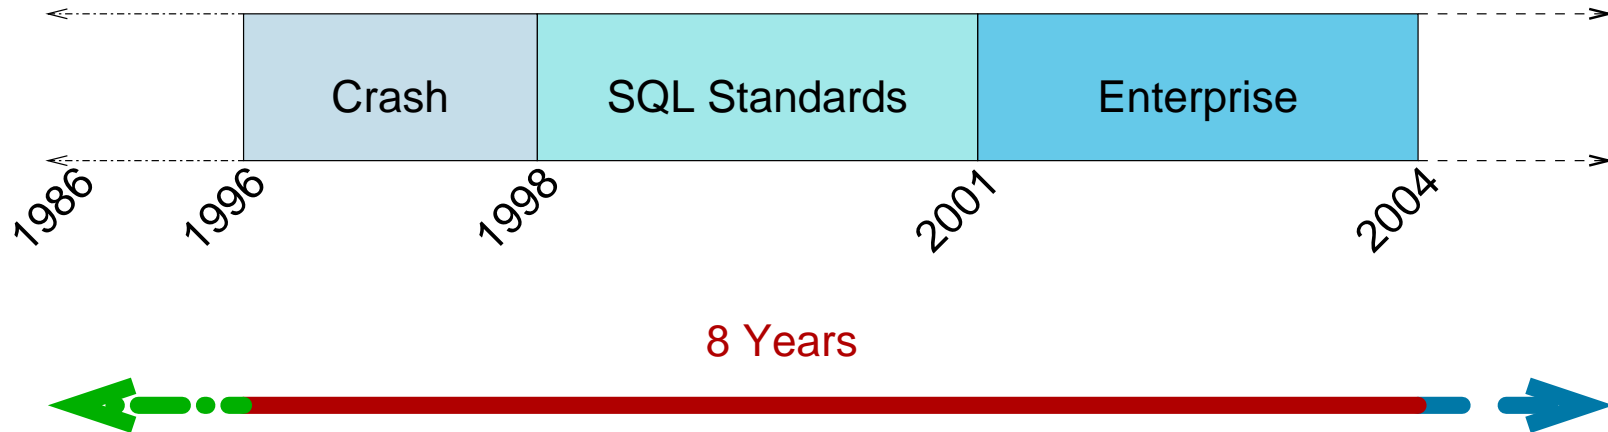
Adopters

- Financial
- Many servers
- Customization
- First adopters, future thinkers
- Company policy

Competition

- Oracle - huge code base, custom applications, \$\$\$
- DB2, MS SQL - like Oracle
- MySQL - GPL, closed development model, open source distribution
- Other database companies are declining because of open-source competition
- Other open-source database projects are having trouble attracting developers

PostgreSQL Evolution



Enterprise features include:

- Improved Performance
- Simplified Maintenance and Administration
- Efficient 24/7 Operation

Case Studies

- Retail
- Scientific
- Internet
- Product

Licensing

- BSD license, limitation of liability
- Not GPL, which requires unrestricted source distribution

Our Developers

Why is programming fun?

- Joy of making things
- Joy of making useful things
- Fun of a complex puzzle
- Joy of learning
- Tractable medium

Concepts derived from *The Mythical Man-Month* by Frederick P. Brooks, Jr.

Open Source Motivations

- Professional advancement, learn new skills
- Practical need for the software
- Mental Stimulation
- Belief in open source

Concepts derived from a survey conducted by OSDN and reported in “Open source a needed outlet for programming pros”, *The Register*, May 2, 2002, <http://www.theregister.co.uk/content/4/23935.html>.

PostgreSQL Motivations

- Professional advancement, learn new skills
- Practical need for the software
 - *Business needs the enhancement*
 - *Business wants PostgreSQL to thrive*
- Mental Stimulation
- Belief in open source

Extreme Case

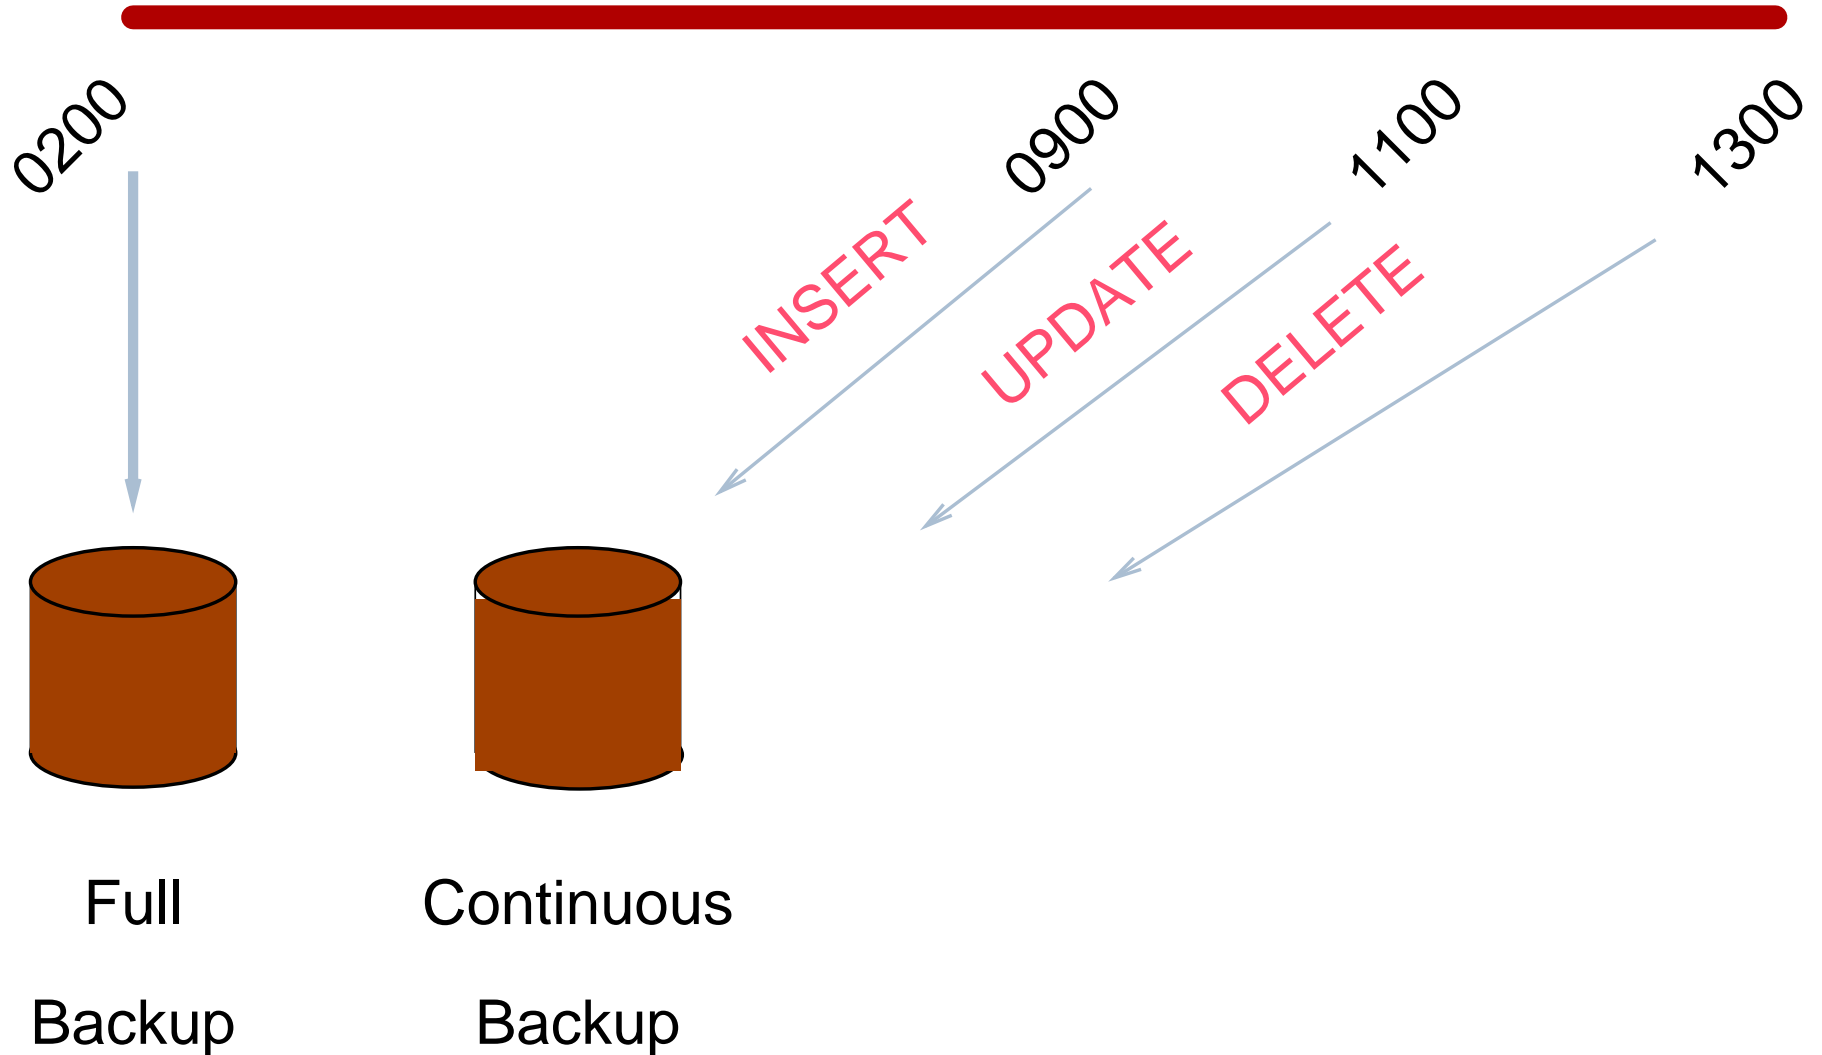
*It was cheaper
to pay to add missing features
to PostgreSQL
than to buy Oracle.*

New features, translations, and now ported to Unix.


Features in 7.5 and Later

7.5+

Point-in-time Recovery: Continuous Backup




MS Windows port: Windows 2000 & XP



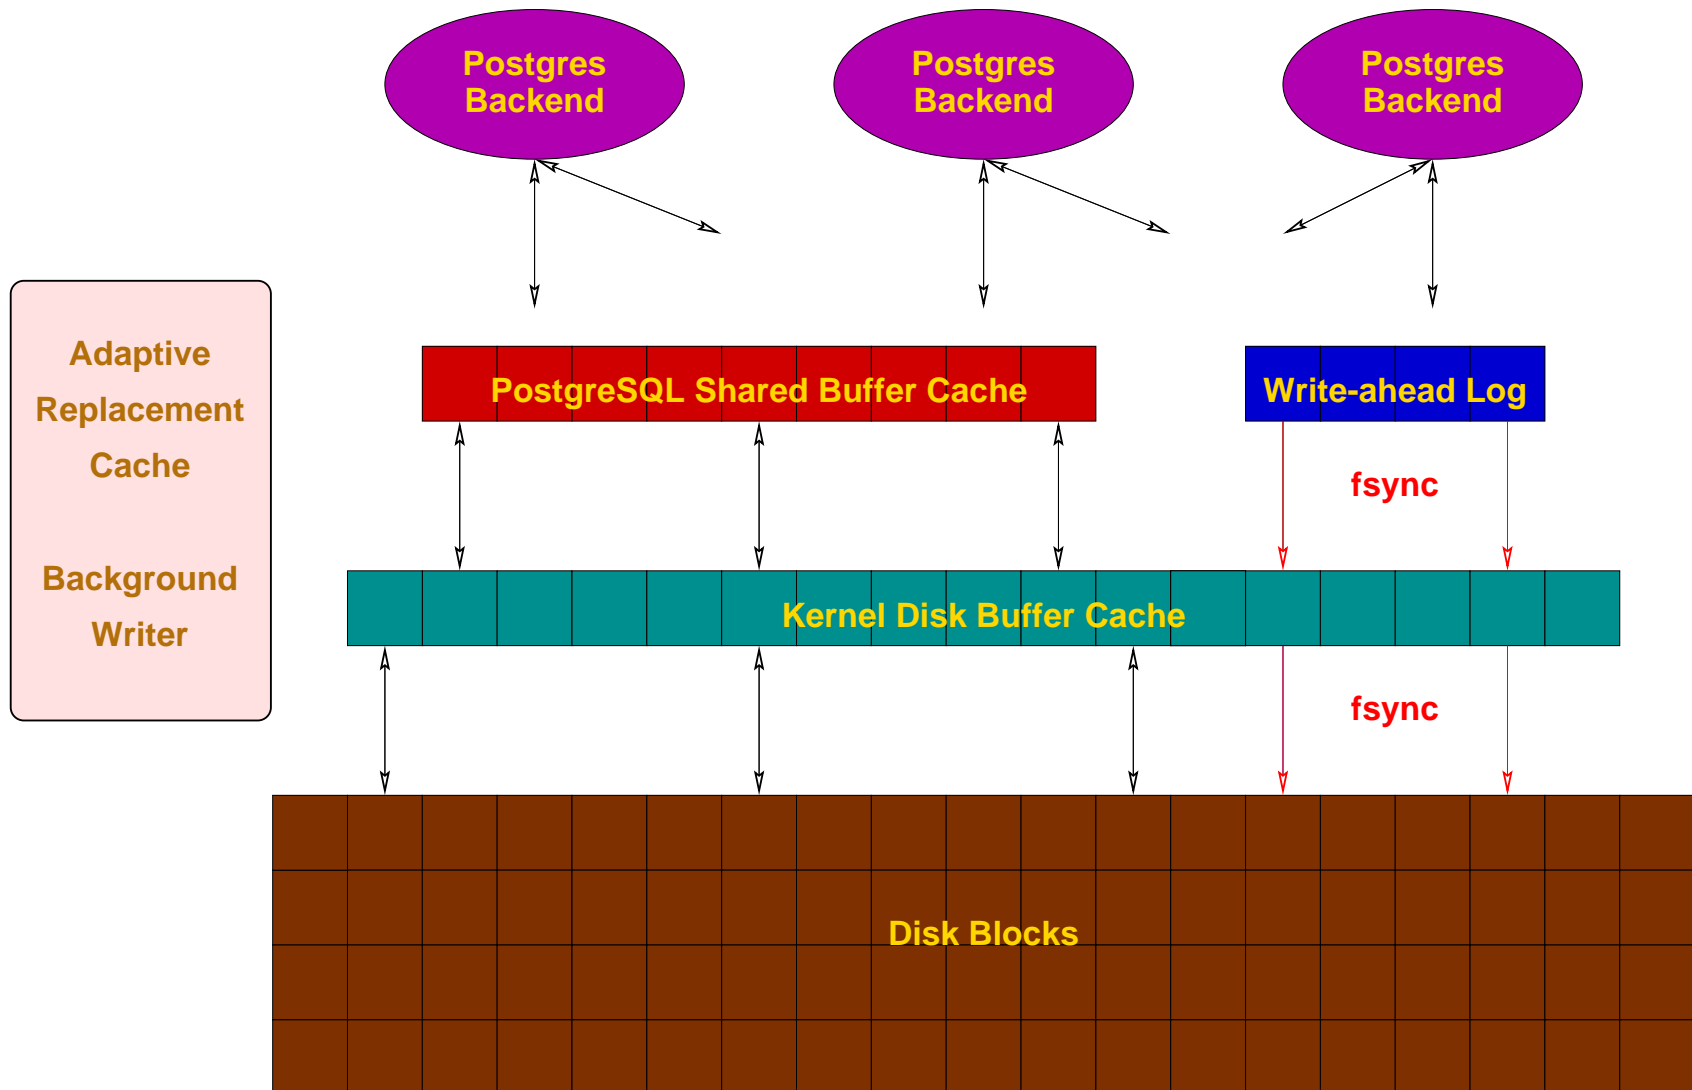
Nested Transactions

```
test=> BEGIN WORK;  
BEGIN  
test=> BEGIN WORK;  
BEGIN  
test=> INSERT INTO test VALUES (55);  
ERROR: duplicate key violates unique constraint "ii"  
test=> ABORT;  
ROLLBACK  
test=> UPDATE test SET total = total + 1;  
UPDATE 1  
test=> COMMIT WORK;  
COMMIT
```




Tablespaces



Improve Shared Buffer Management



More Master/Slave Replication Solutions



Asynchronous



Conclusion

