



PostgreSQL Web Projects: From Start to Finish


Simon Riggs
2nd Quadrant






PostgreSQL

- Open Source Relational Database
 - 11 year project, 21 years since inception
 - BSD Licence
- Database
 - Server-based SQL:2003 compliant
 - Highly concurrent
 - Fast, Robust, Secure




Stunning Features

- Visual EXPLAIN
- Procedural Language Debugger
- (Amazing) Index Support
- Cost Base Optimizer, Histograms
- SQL:2003




Principles

- “It just works”
- “Failure is not an Option”
- All features, Great and Small
- Respect Published Standards
- Build on the work of others
- Peer Review of Everything
- Good documentation is fundamental
- Solid, Simple Code (...works at 4am)
- Improvement is always possible




Sponsors (VC Startups)






Sponsors (Hosting Providers)






Sponsors (Services)






Sponsors (Major Vendors)






Sponsors (Major Users)






Sponsors






UK Users

- Web Infrastructure
- Independent Service Providers
- Major Telcos and Telecom Equipment Vendors
- Gaming, Gambling, Media
- Local Government
- University Projects
- The City

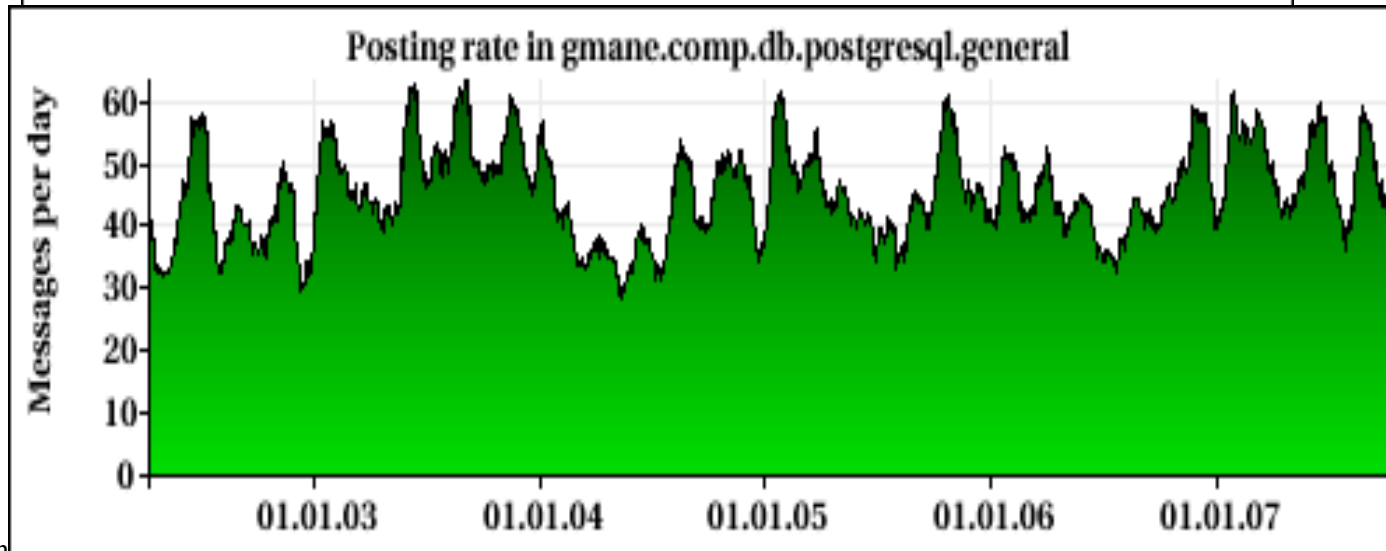
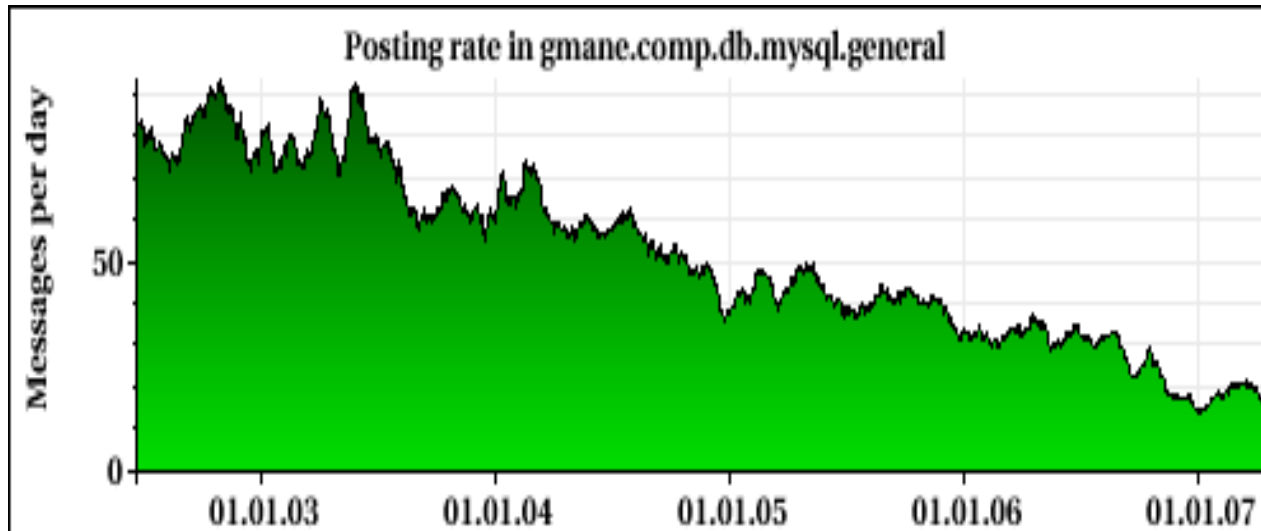



International Users

- Afilias
- Skype
- NTT
- ...
- 25% market share in Japan
- 4,000+ user group members




Mailing List Postings






Use Cases

- OLTP
- Web Applications
- Data Warehouse
- Audit & Review




Web Project Evolution

- Startup
- Success
- The Big Time




Get Started Quickly

- Zero Licence Cost
 - Greater % of money goes to Development
 - No licence planning, no vendor negotiations
 - Faster Time To Market
 - Faster Time To Profit
- Install and Go
 - Runs everywhere, with everything
 - Easy setup
 - SQL:2003 Database
 - Graphical Admin & Design tools
 - Stack Integration




Basic Web Features

- SQL:2003
- UTF-8
- XML Support
- National Language Support
- Great documentation
- Responsive, Busy Community




Works the way you do

- Multiple client interfaces
 - Java, perl, python, PHP, odbc, .net, native
- Multiple stack integrations
 - LAMP, Java, .Net, Rails
- Users, Path, environment variables
- Multiple Server side languages
 - PL/pgSQL, PL/java, PL/perl, PL/R, PL/PHP, PL/Ruby
- regex




Security

- Security by default
- Multi-layer security model
- Full encryption
- Enterprise security infrastructure integration
- Full disclosure of vulnerabilities
- Security bug fix SLA




Cool Web Features

- It just works!
- Arrays
- Text > 2000, Autocompression
- Datatypes
 - enum, uuid, cidr, macaddr
- SET parameters
- Asynchronous Commit




Enhance Incrementally

- Advanced Features
- Extensible Features
- Transactional DDL
- Production Tuning facilities
 - Logging, profiling, server-side debugging
- **CREATE INDEX CONCURRENTLY**




Major Extensions

- Full Text Search
 - Snowball
 - GIN – inverted search index
 - GIST – high concurrency
- Geographical Information System (GIS)
 - Full implementation
 - R-Tree indexes
 - Custom statistics functions




Data Availability

- No block locking
- COPY (load...)
- VACUUM
- CREATE INDEX CONCURRENTLY




High Availability

- Hot Backup
- Point in Time Recovery
- Warm Standby




Replication

- “Postgres doesn't have replication”
 - True: Not part of the core package
 - False: There's loads of replication alternatives
- Slony
- Warm Standby
- Bucardo
- Londiste



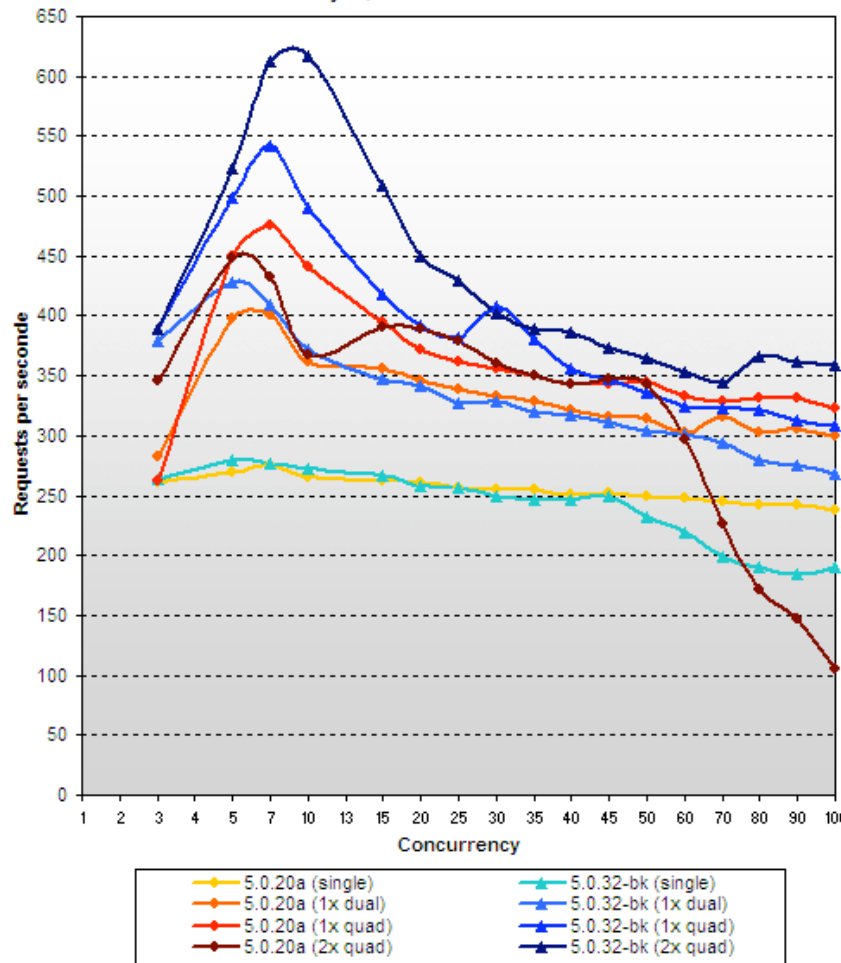
Scale Up, VLDB

- SMP scalability
 - 16 cores verified by Unisys Scalability Lab
- VLDB
 - Partitioning
 - Many Terabyte references worldwide
 - Frequently used in the 100GB+ range





MySQL v PostgreSQL Scalability

Tweakers.net Database-simulatie
MySQL 5.0.20a vs. 5.0.32-bk



Tweakers.net Database-simulatie
Clovertown vs. Woodcrest - PostgreSQL 8.2-dev







MySQL v PostgreSQL Scalability


<http://jamonation.com/node/734>

pgsql sysbench on 8-core opteron, FreeBSD 7.0 with N CPUs active




mysql sysbench on 8-core opteron






Scale Out

- Result Cache
 - Memcached, pgmemcache
- Multi-node scalability
 - PgBouncer, pgpool
 - PL/Proxy



8.3 Features

- HOT, Distributed Checkpoint, Autotuning BgWriter
- Further Scalability improvements
- Asynchronous Commit
- Synchronous Scanning
- XML Datatypes
- Warm Standby, Replication improvements



PostgreSQL