

PostgreSQL

&



SQLite




McAfee®



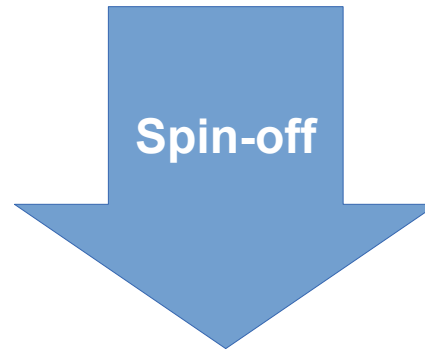


The most widely used SQL engine


- More than 2,000,000,000 running instances
- More than 500,000 applications



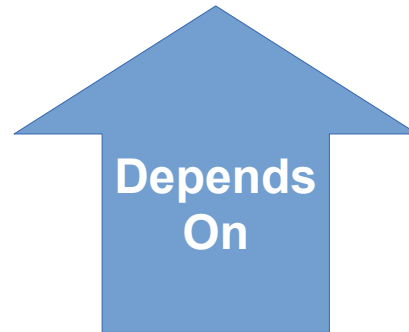
PostgreSQL



SQLite



PostgreSQL



SQLite

WWWPD

<http://www.sqlite.org/sqllogictest>

- SQLite
 - MySQL
 - Oracle
 - SQL Server
 - PostgreSQL
- Crashes and/or incorrect answers
- ← Never crashed and always correct



versus



Enterprise Data Depot

Client and Server

Scale Up

Hidden Files

Dump and Restore

Multiple binaries

Application File Format

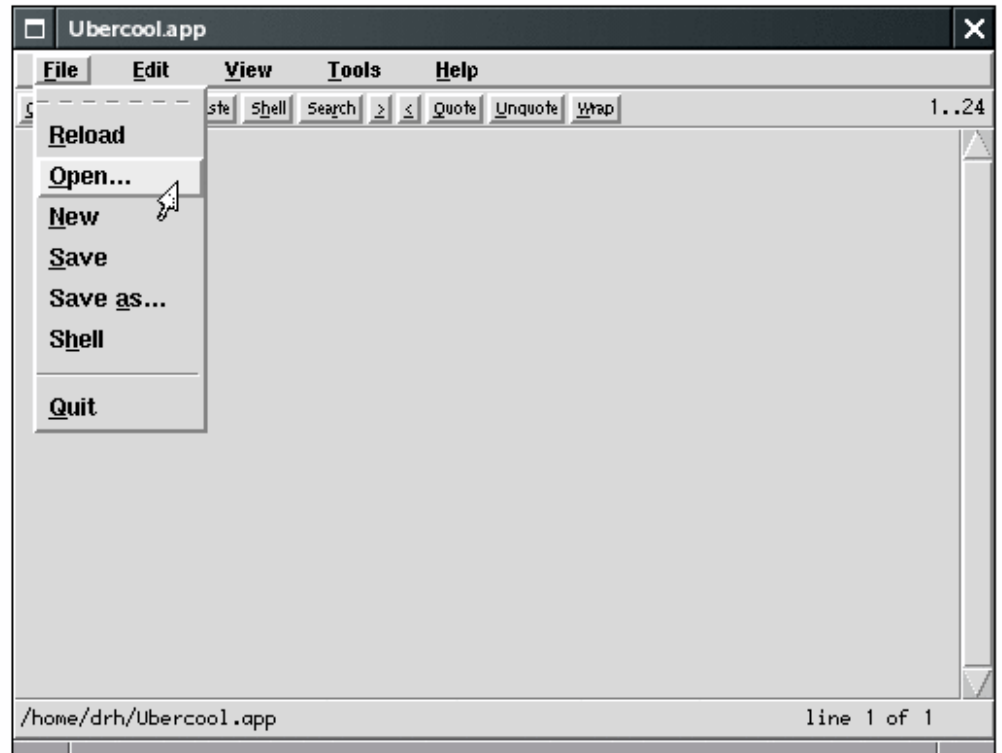

Serverless

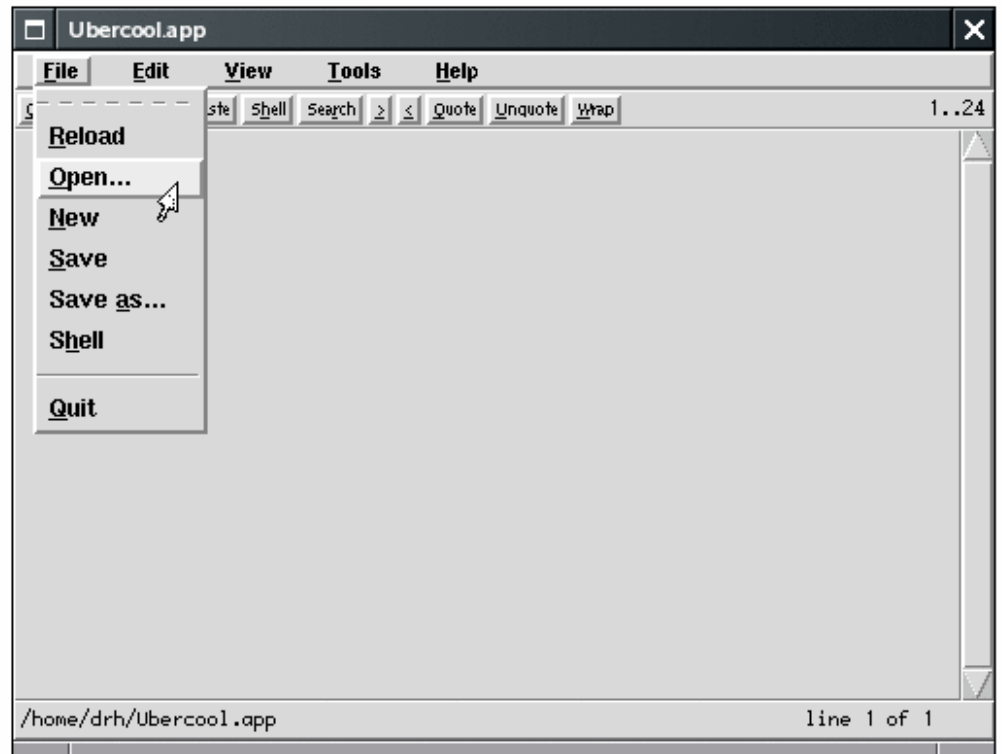

Scale Down

Single Disk File

Backwards Compatible

One file of ANSI-C

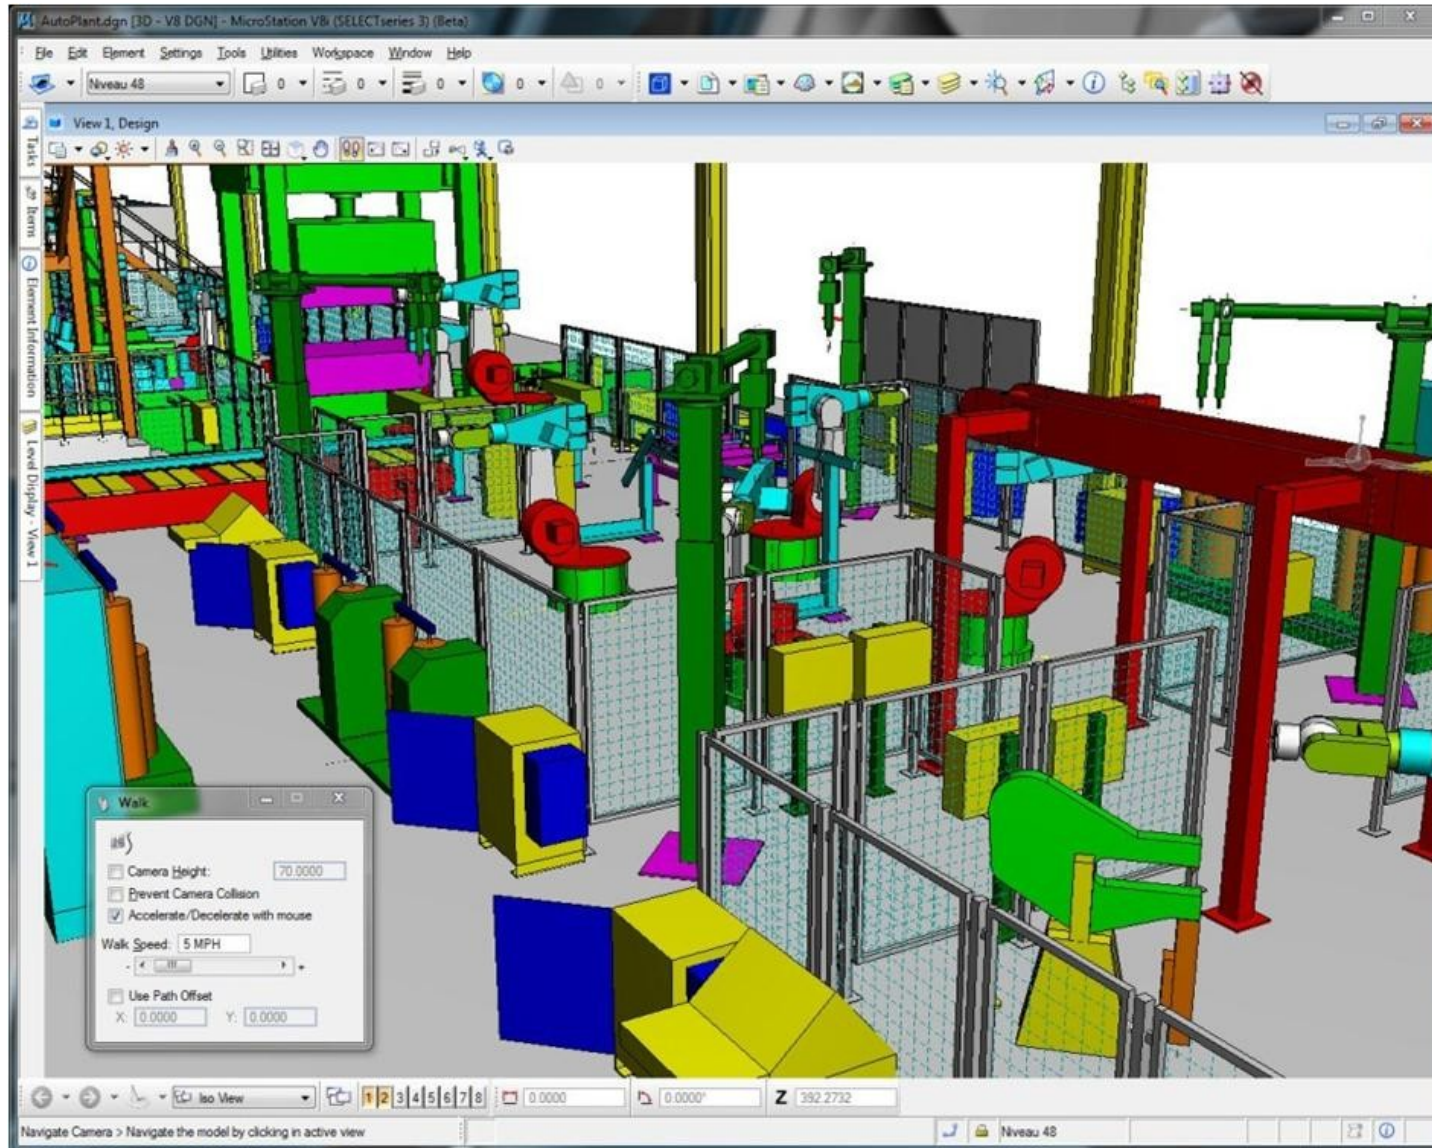




SQLite as Application File Format

- No parsing code to write and debug
- Single-file documents
- High-level query language
- Accessible content
- Cross-platform and cross-language
- Atomic transactions
- Incremental and continuous updating
- Easily extensible
- Multi-process and multi-thread safe
- Improved performance

Real Example: MicroStation



Real Example: Adobe Lightroom



What If....

OpenDocument was an SQLite database
Instead of a ZIP archive of XML files...

- Fast and low-I/O save of small changes
- Fast startup
- Reduced memory usage (no need to hold the entire presentation in memory at once)
- No need for “recovery” after a crash
- No need to “file save”
- Undo across sessions
- Large searchable database of slides

What If....

ePub was an SQLite database
Instead of a ZIP archive of XML files...

- Full text search
- Faster open of large documents
- Multiple image resolutions

SQLite Archiver

```
CREATE TABLE sqlar(  
  name TEXT PRIMARY KEY, -- name of the file  
  mode INT,              -- access permissions  
  mtime INT,             -- last modification time  
  sz INT,                -- original file size  
  data BLOB              -- compressed content  
);
```

- <http://www.sqlite.org/sqlar>
- Size similar to ZIP
- Transactional
- Concurrent & random access

SQLAR is 0.47% larger than ZIP

```
-rw-r--r-- 1 drh drh 9677708 May 17 12:44 pgcon2014.odp  
-rw-r--r-- 1 drh drh 9722880 May 17 12:46 pgcon2014.sqlar
```

What If....

Git stored content in an SQLite database instead of in a pile-of-files under the “.git” folder.

- Advanced queries for a richer user interface
- Proof against crashes
- Wiki and Tickets
- Concurrent access
- Coding errors less likely to corrupt repository
- On-the-fly repository compression
- Single-file repository



versus



Enterprise Data Depot

Client and Server

Scale Up

Hidden Files

Dump and Restore

Multiple binaries

Application File Format

Serverless

Scale Down

Single Disk File

Backwards Compatible

One file of ANSI-C



versus



Enterprise Data Depot

Application File Format

“SQLite is not a replacement for PostgreSQL.

SQLite is a replacement for fopen().”



versus



OK

```
CREATE TABLE abc(xyz INT);  
INSERT INTO abc VALUES('123');
```

OK

OK

```
CREATE TABLE abc(xyz TEXT);  
INSERT INTO abc VALUES(123);
```

OK

NO!

```
CREATE TABLE abc(xyz INT);  
INSERT INTO abc VALUES('Hello!');
```

OK



versus



| | | |
|------------|---|-----------|
| NO! | CREATE TABLE abc(xyz); | OK |
| NO! | CREATE TABLE abc(xyz ANY); | OK |
| OK | CREATE TABLE abc(xyz TEXT); | OK |
| NO! | CREATE TABLE abc(xyz BLOB); INSERT INTO abc VALUES(x'FEDC'); | OK |

United by SQL

- The SQL front-end is the secret sauce
- Most widely known programming language
- Let the engine figure out the algorithm
- Representation is the essence of programming

“Representation is the essence of computer programming.”

“Show me your flowcharts and conceal your tables, and I shall continue to be mystified. Show me your tables, and I won't usually need your flowcharts; they'll be obvious.”

- Fred Brooks, *The Mythical Man-Month*, pp102-103





"Data dominates. If you've chosen the right data structures and organized things well, the algorithms will almost always be self-evident. Data structures, not algorithms, are central to programming."

- Rob Pike: Rule of Programming #5

"Bad programmers worry about the code. Good programmers worry about data structures and their relationships."

- Linus Torvalds, on the Git mailing list, 2006-06-27




United by SQL


- The SQL front-end is the secret sauce
- Most widely known programming language
- Let the engine figure out the algorithm
- Representation is the essence of programming

*Key/value and “eventually consistent”
data stores are a passing fad.*

SQL is here for the long term.

Postmodern Databases

- 
- Absence of objective truth
 - Queries return opinions rather than facts



PostgreSQL
SQLite
Oracle
SQL Server



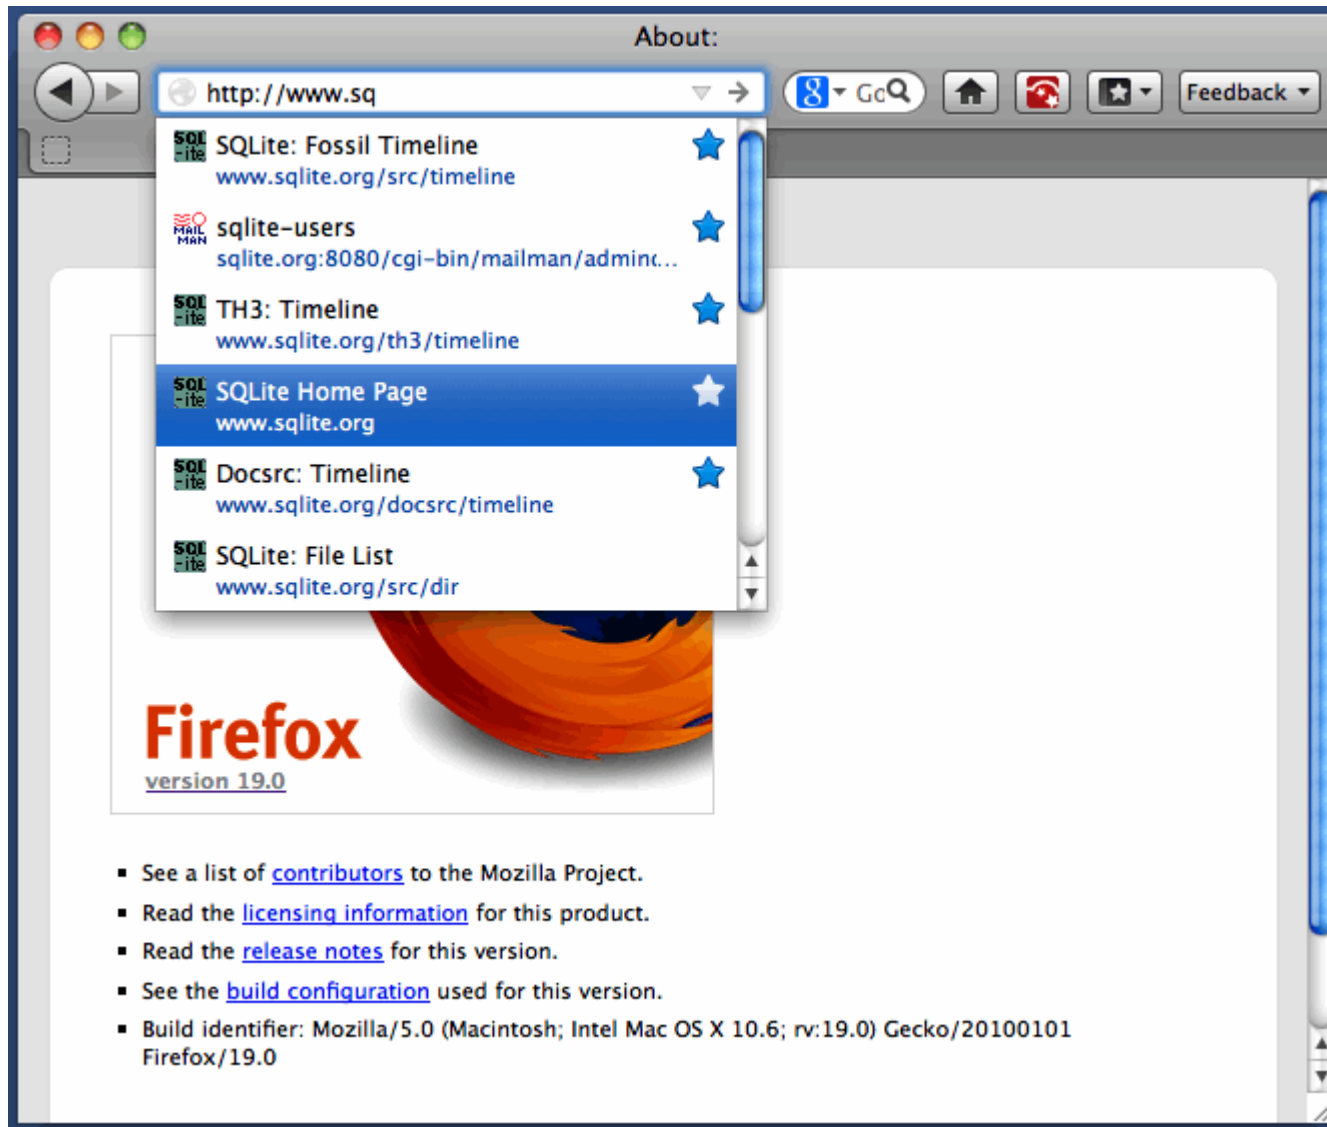
Cassandra
MongoDB
CouchDB
BigTable



Alexander Lloyd at Berlin Buzzwords, 2012

*There was another big cultural shift at Google...
The SQL-based analytics system, Dremel... made a lot
of SQL converts at Google. People realized it is incredibly
powerful to just push the semantics of your query down into
the storage system and let it figure out what to do.*


*NoSQL databases that only have weak consistency
are enforcing a broadly applied premature optimization
on the entire system.*



PostgreSQL: Timeline

www.sqlite.org/postgresql/timeline?c=2000-05-29&nd&n=8&y=ci

News Projects Local Docs Banks Biz Admin Gmail Imported From Fi...



PostgreSQL

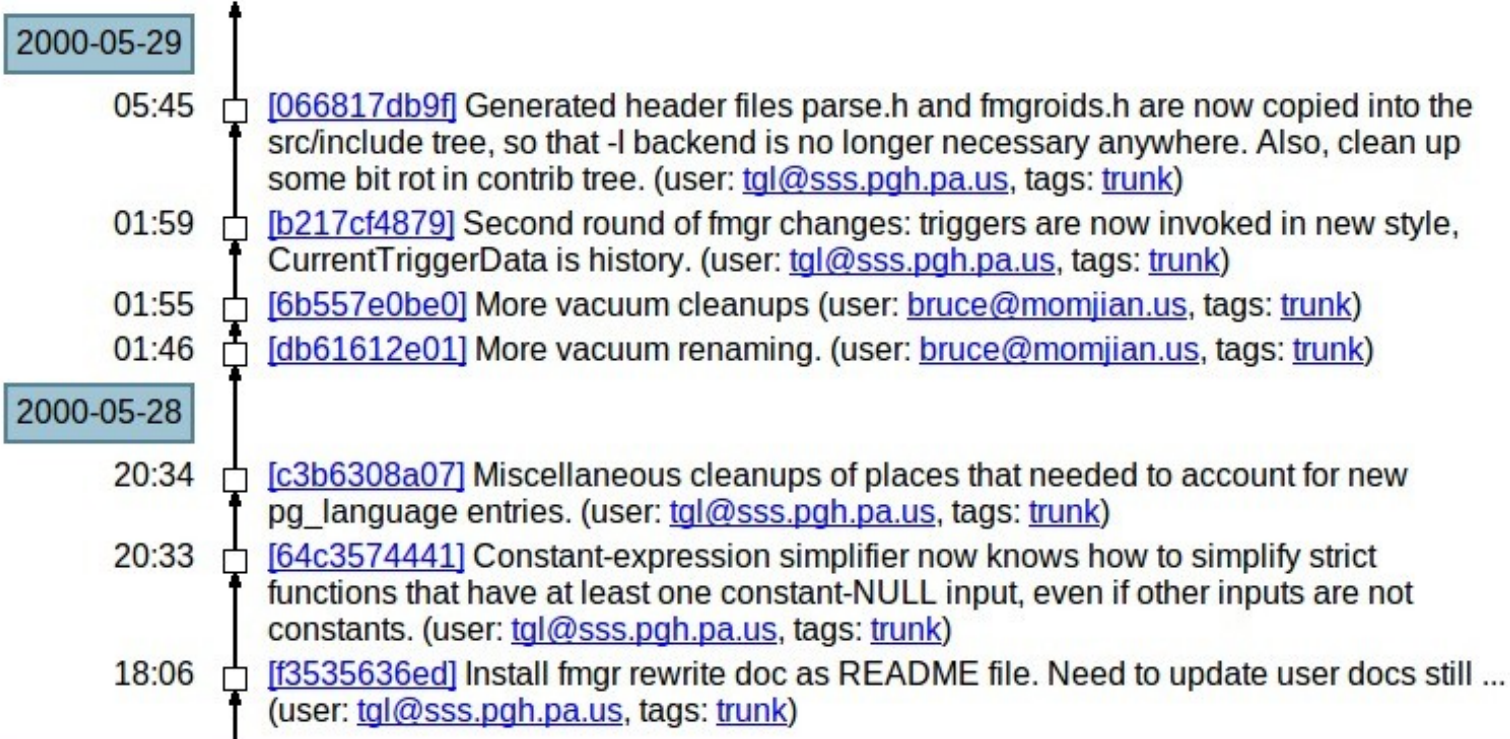
Timeline

Not logged in

Home Timeline Files Branches Tags Tickets Wiki Login

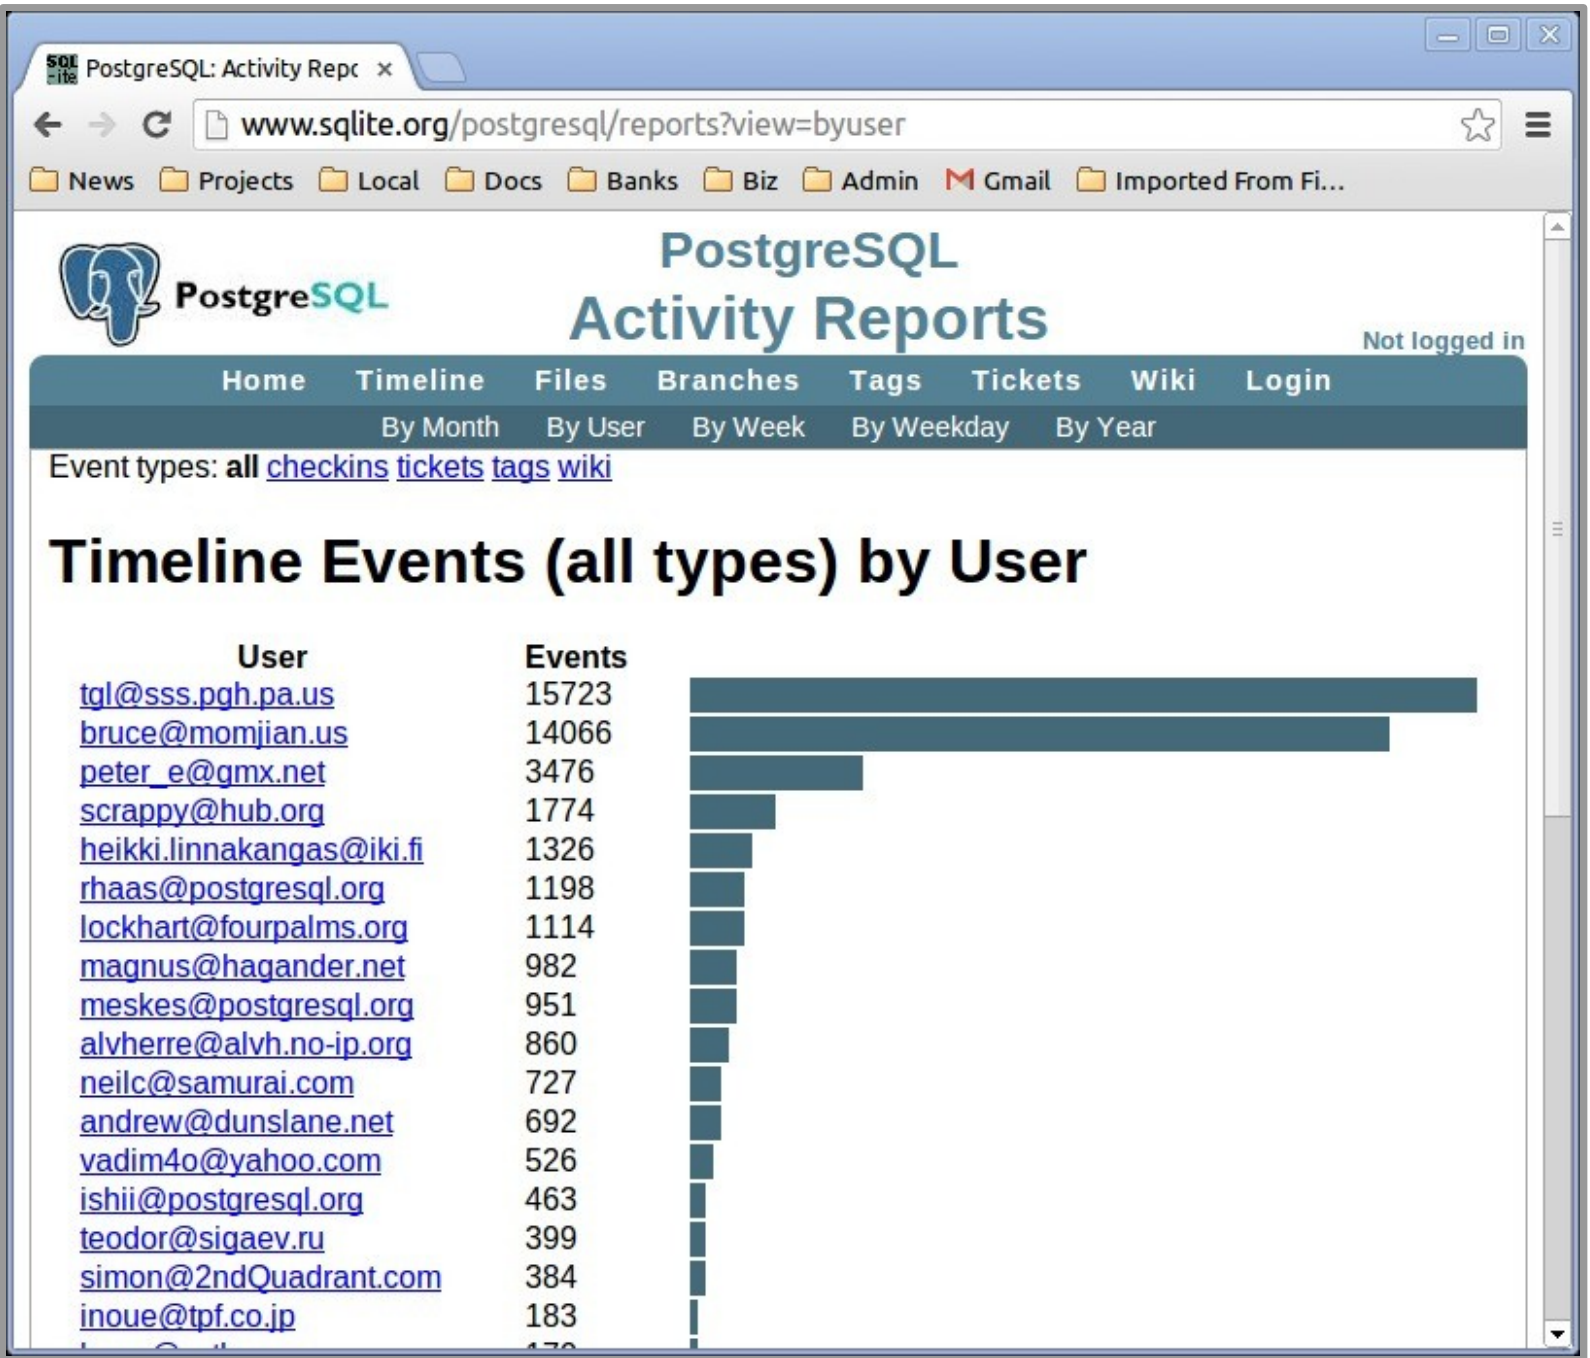
200 Entries All Types Events Only Show Files Tags Only Tickets Only Unhide Wiki Only

8 checkins occurring around 2000-05-29.



| Date | Time | Commit ID | Description | User | Tags |
|------------|-------|--------------|--|-------------------|-------|
| 2000-05-29 | 05:45 | [066817db9f] | Generated header files parse.h and fmgroids.h are now copied into the src/include tree, so that -l backend is no longer necessary anywhere. Also, clean up some bit rot in contrib tree. | tgl@sss.pgh.pa.us | trunk |
| 2000-05-29 | 01:59 | [b217cf4879] | Second round of fmgr changes: triggers are now invoked in new style, CurrentTriggerData is history. | tgl@sss.pgh.pa.us | trunk |
| 2000-05-29 | 01:55 | [6b557e0be0] | More vacuum cleanups | bruce@momjian.us | trunk |
| 2000-05-29 | 01:46 | [db61612e01] | More vacuum renaming. | bruce@momjian.us | trunk |
| 2000-05-28 | 20:34 | [c3b6308a07] | Miscellaneous cleanups of places that needed to account for new pg_language entries. | tgl@sss.pgh.pa.us | trunk |
| 2000-05-28 | 20:33 | [64c3574441] | Constant-expression simplifier now knows how to simplify strict functions that have at least one constant-NULL input, even if other inputs are not constants. | tgl@sss.pgh.pa.us | trunk |
| 2000-05-28 | 18:06 | [f3535636ed] | Install fmgr rewrite doc as README file. Need to update user docs still ... | tgl@sss.pgh.pa.us | trunk |

```
SELECT blob.rid AS blobRid,  
       uuid AS uuid,  
       datetime(event.mtime) AS timestamp,  
       coalesce(ecomment, comment) AS comment,  
       coalesce(euser, user) AS user,  
       blob.rid IN leaf AS leaf,  
       bgcolor AS bgColor,  
       event.type AS eventType,  
       (SELECT group_concat(substr(tagname,5), ', ' )  
        FROM tag, tagxref  
        WHERE tagname GLOB 'sym-*'   
              AND tag.tagid=tagxref.tagid  
              AND tagxref.rid=blob.rid  
              AND tagxref.tagtype>0) AS tags,  
       tagid AS tagid, brief AS brief,  
       event.mtime AS mtime  
FROM event CROSS JOIN blob  
WHERE blob.rid=event.objid  
      AND NOT EXISTS(SELECT 1 FROM tagxref  
                    WHERE tagid=5 AND tagtype>0  
                    AND rid=blob.rid)  
  
      AND event.type='ci'  
      AND event.mtime>=2451693.500000  
ORDER BY event.mtime ASC  
LIMIT 8
```



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

Concluding Thoughts

- Think of PostgreSQL as a programming language
- PostgreSQL is the best available reference platform for SQL
- Use PostgreSQL as an enterprise data depot, and use SQLite for application files.