Which databases solve my problem?

a survey of open source databases

Selena Deckelmann
End Point Corporation
@selenamarie
PostgreSQL Global Development Group

LCA 2010

2005:

BerkeleyDB

MySQL

PostgreSQL

SQLite

2010:

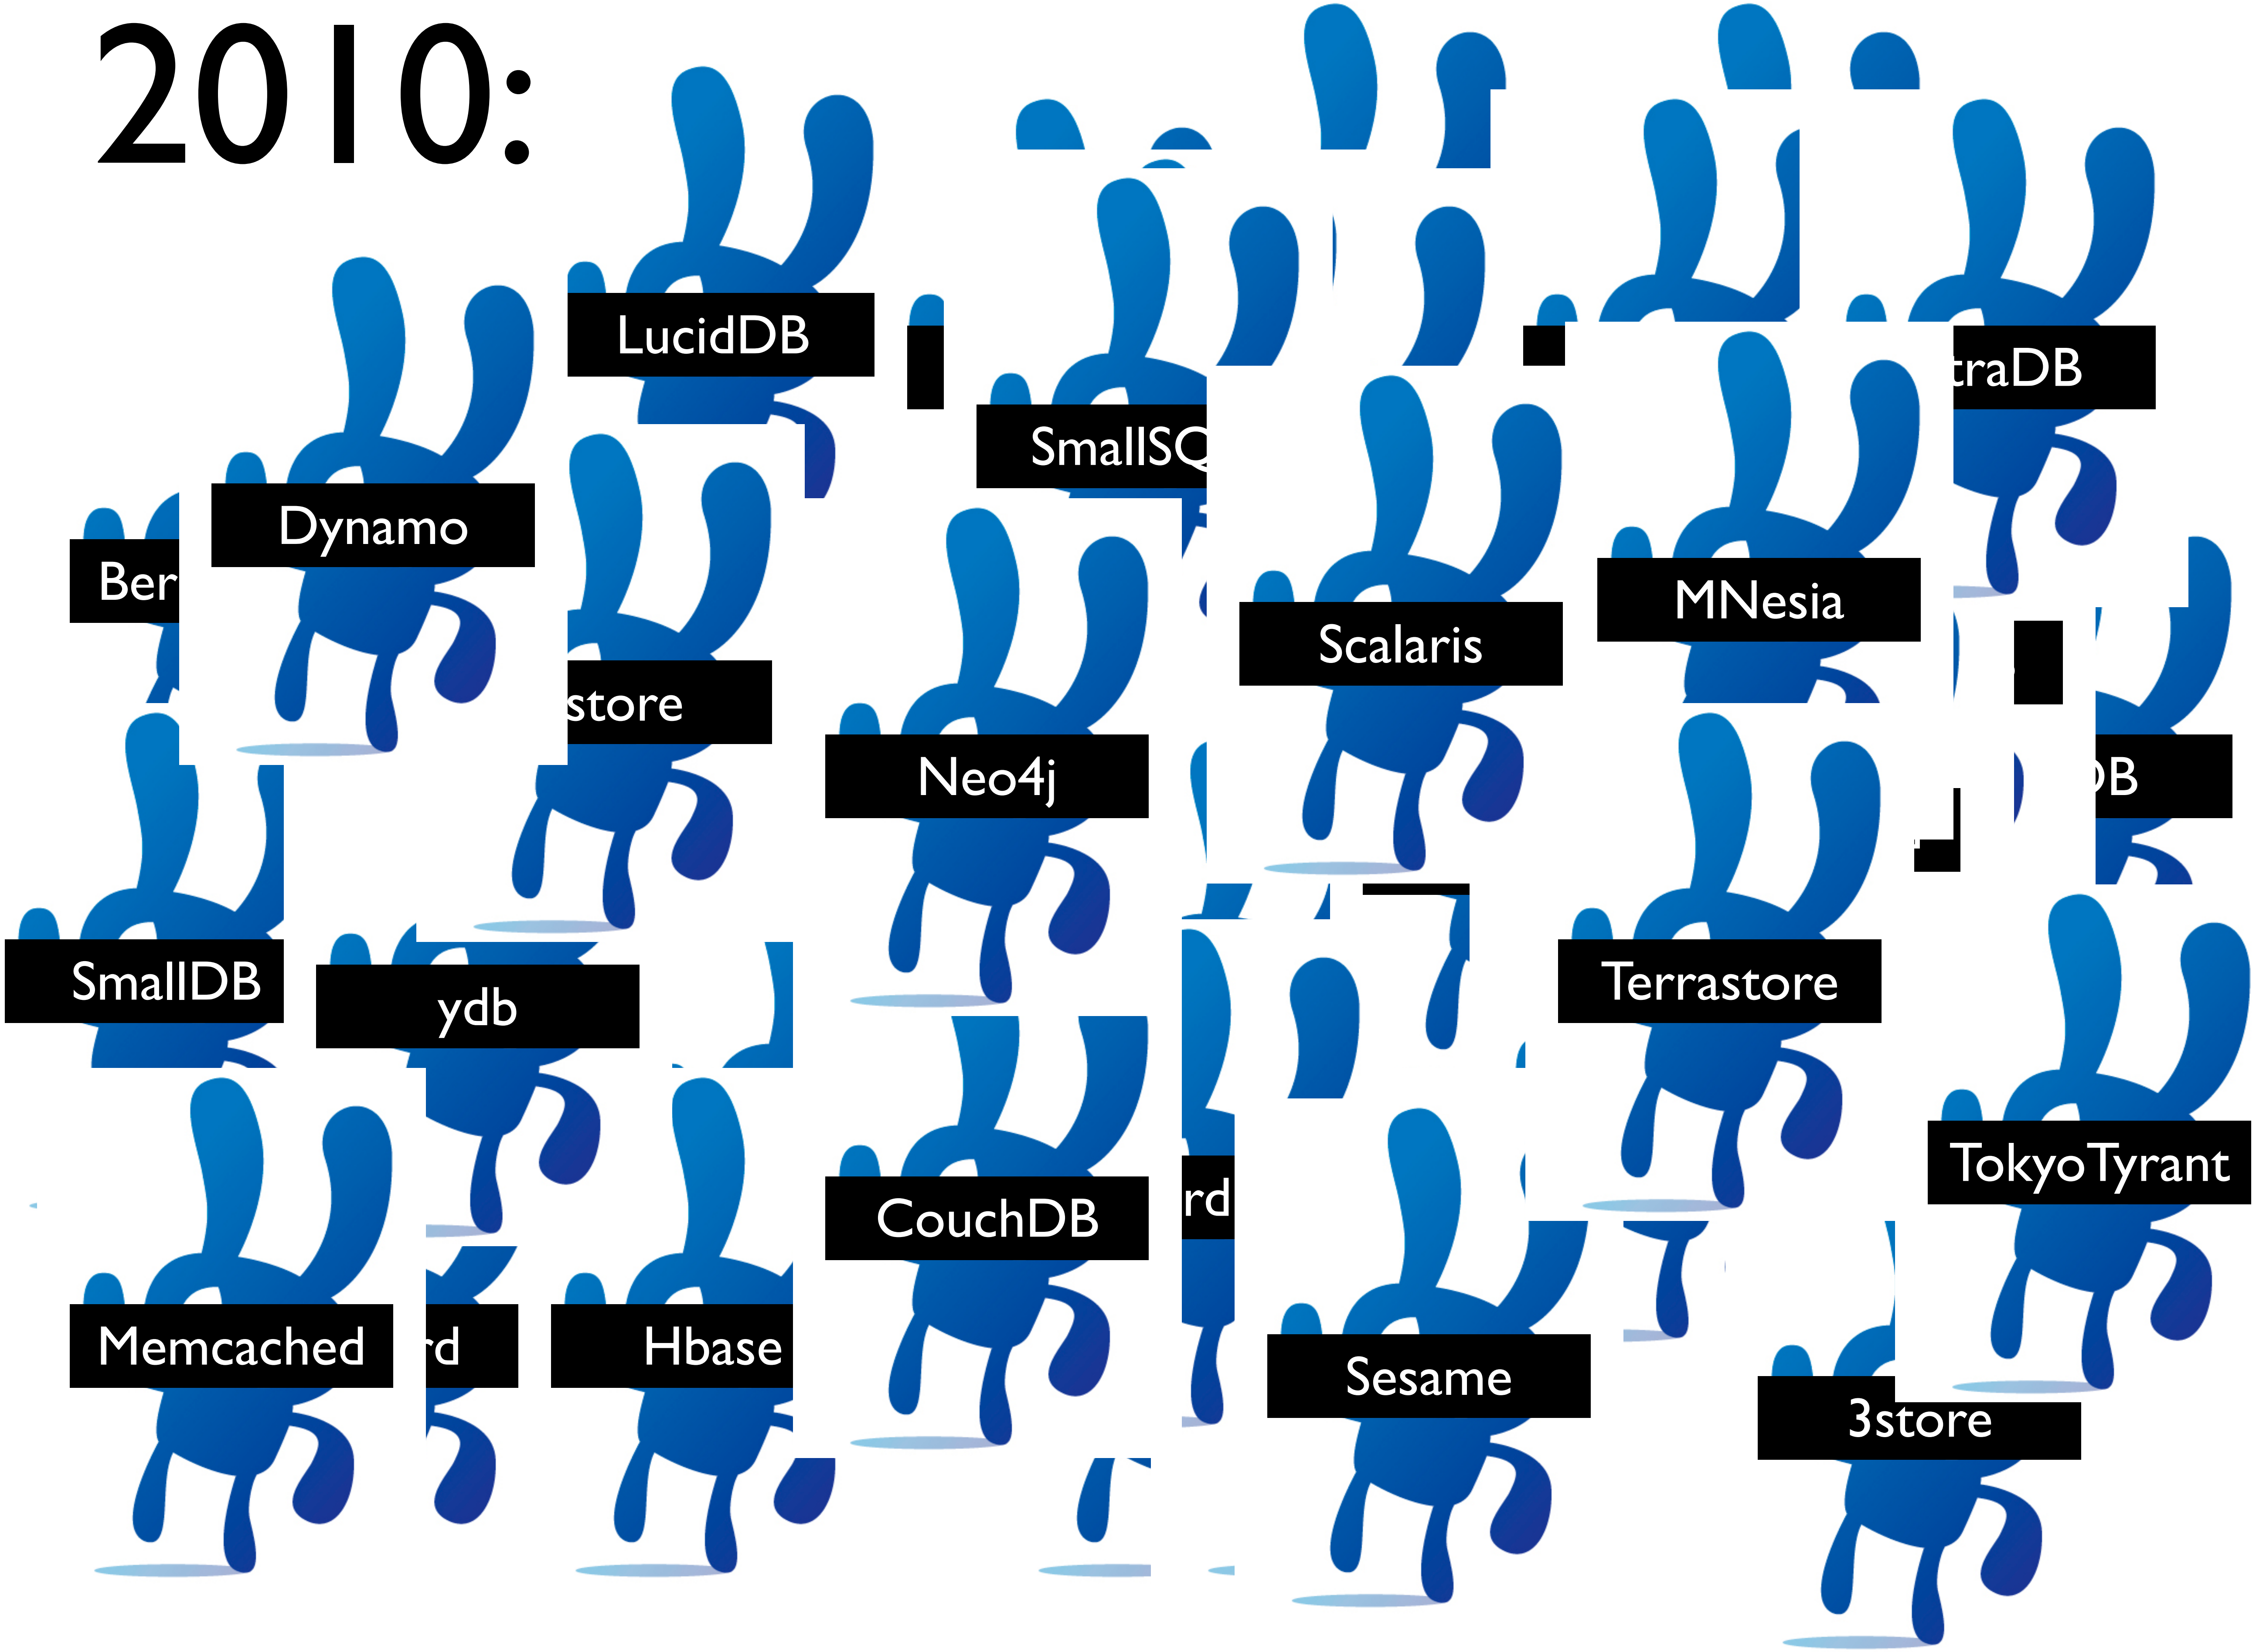
LucidDB
SmallSQ
traDB
Ber
Dynamo
store
Scalaris
MNesia
Neo4j
DB
SmallDB
ydb
Terrastore
TokyoTyrant
rd
CouchDB
Memcached
rd
Hbase
Sesame
3store

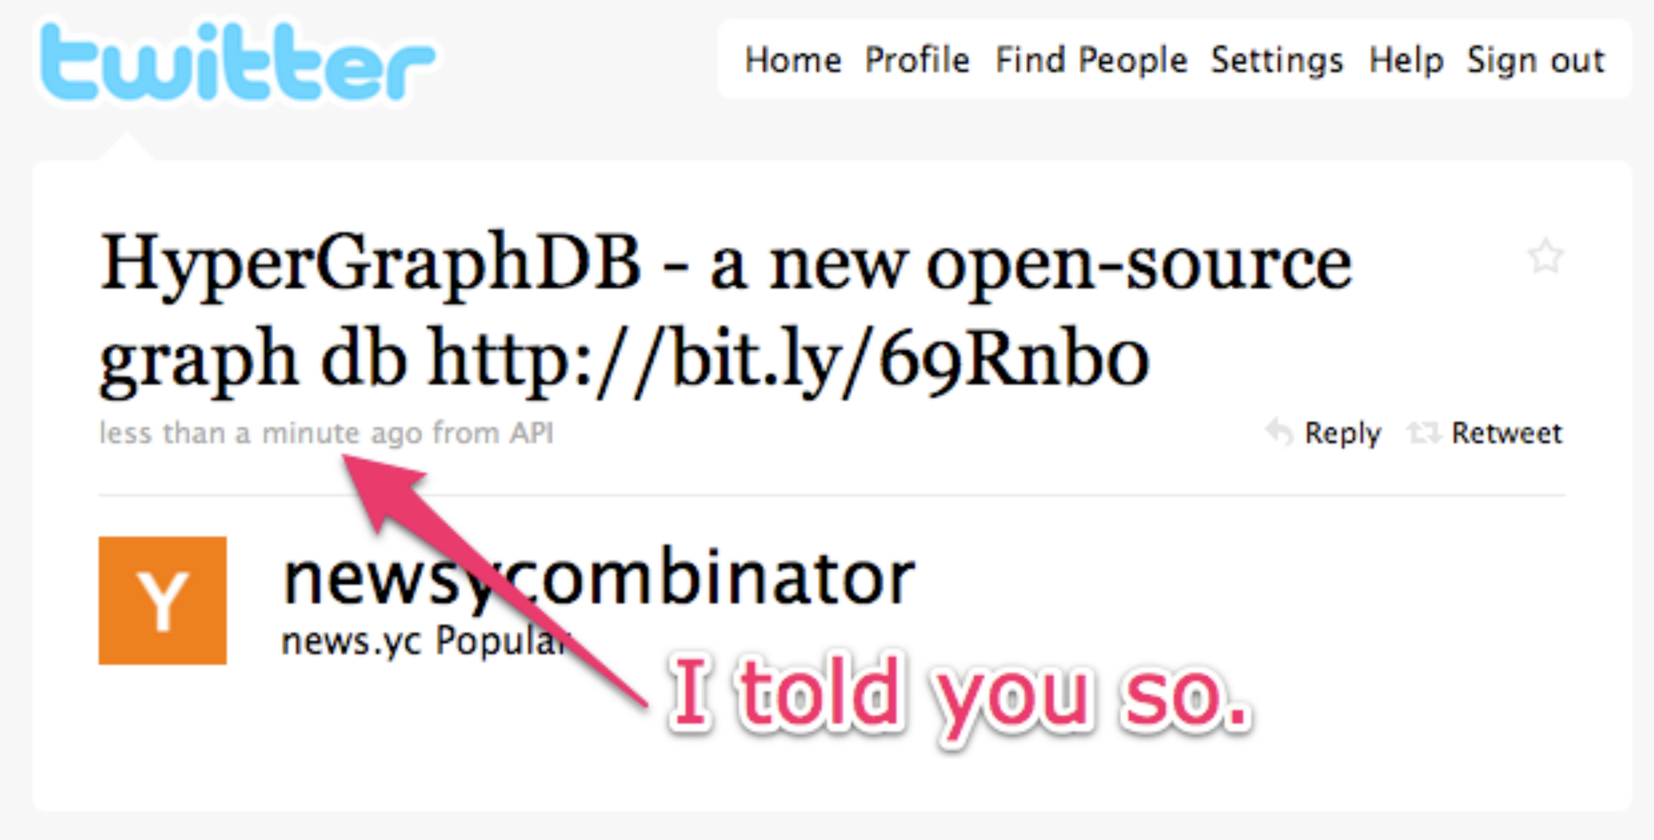

HyperGraphDB - a new open-source graph db http://bit.ly/69Rnbo

less than a minute ago from API

newsycombinator

news.yc Popular

I told you so.

2am on Monday morning.

Which open source database should I use?

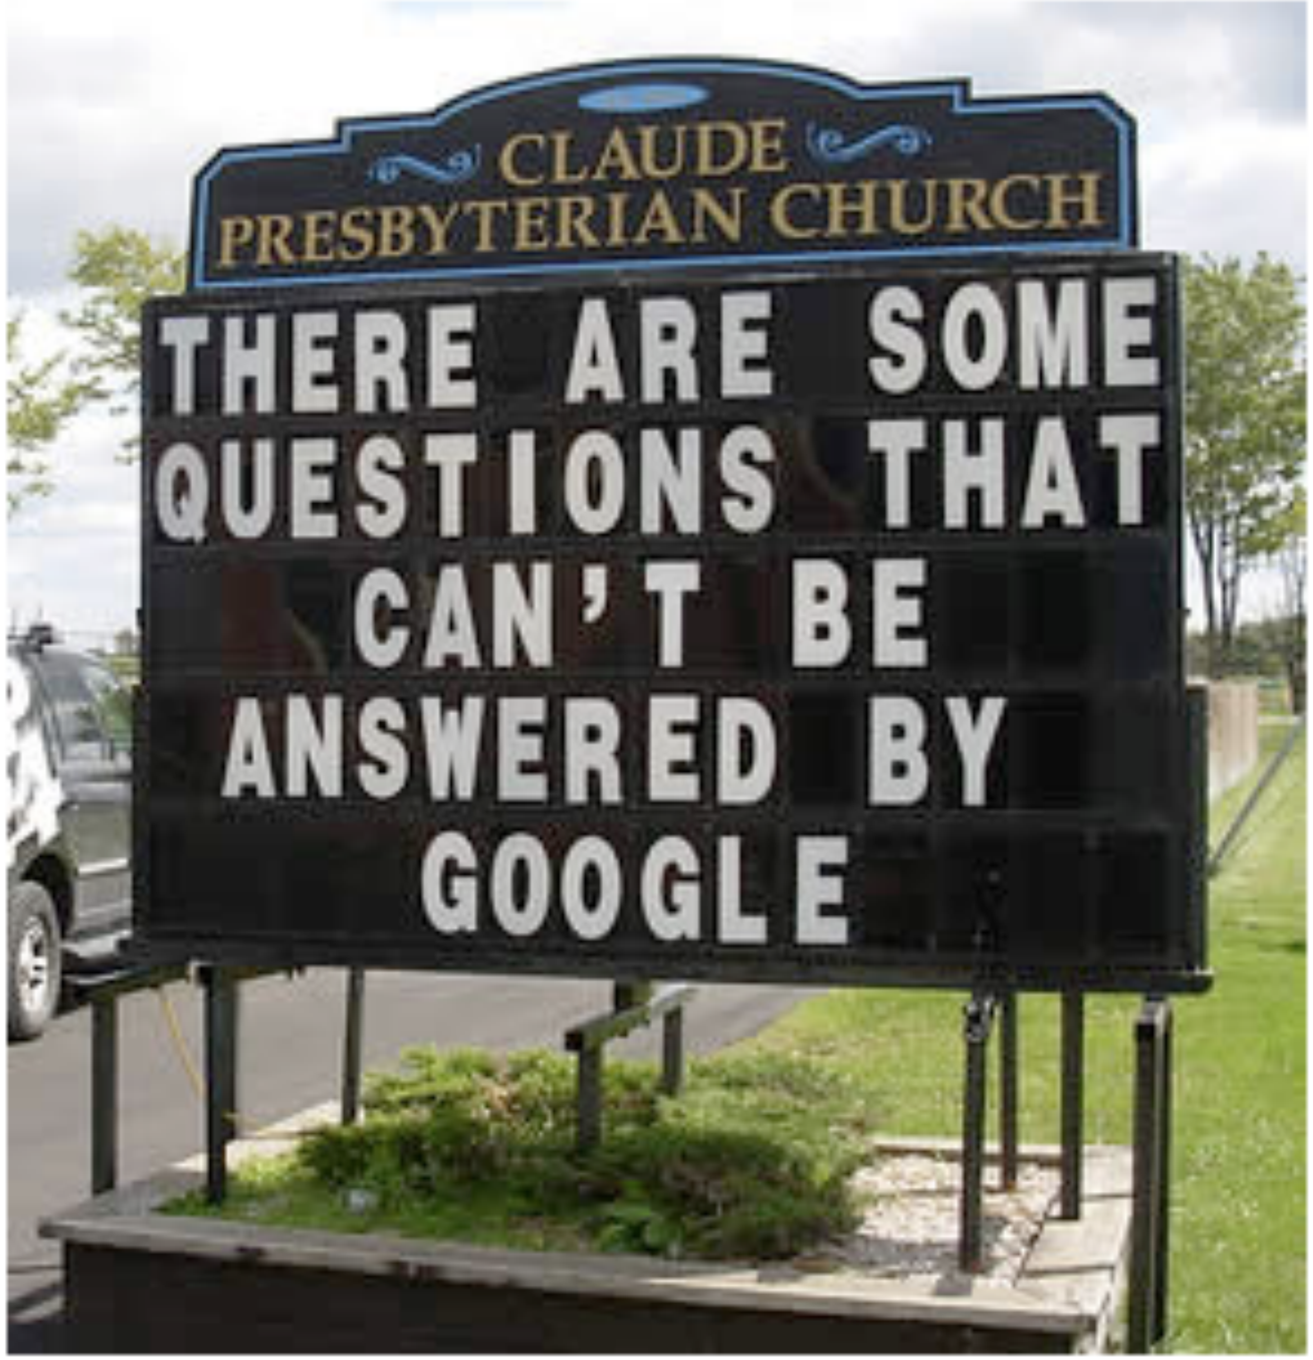

via http://www.oddee.com/item_86516.aspx

MySQL vs PostgreSQL

LCA 2010

ENDpoint

What problem are you trying to solve?

Some problems:

I need to store and manipulate GIS data.

LCA 2010

END point

I need a database for my blog.

LCA 2010

END point

I have ONE BILLION users to store and analyze data from.

Define your problem.

ORACLE
SOFTWARE POWERS THE INTERNET
Priceline's House Of Cards
In Search of Tom Peters • Working Past 90
FORTUNE
The Next Richest Man In The World
Larry Ellison is sick of playing second banana to Bill Gates. So what's his plan to become No. 1? Make Oracle the Microsoft of the Internet.
aol keyword: fortune
www.fortune.com

Which problems are important?

performance

your use case.

test with real data.

interoperability

can I get my data in/out?

how painful is it?

sustainability

how is the software made?

Which databases solve my problem?

LCA 2010

ENDpoint

🟨 Free and Open Source Databases*

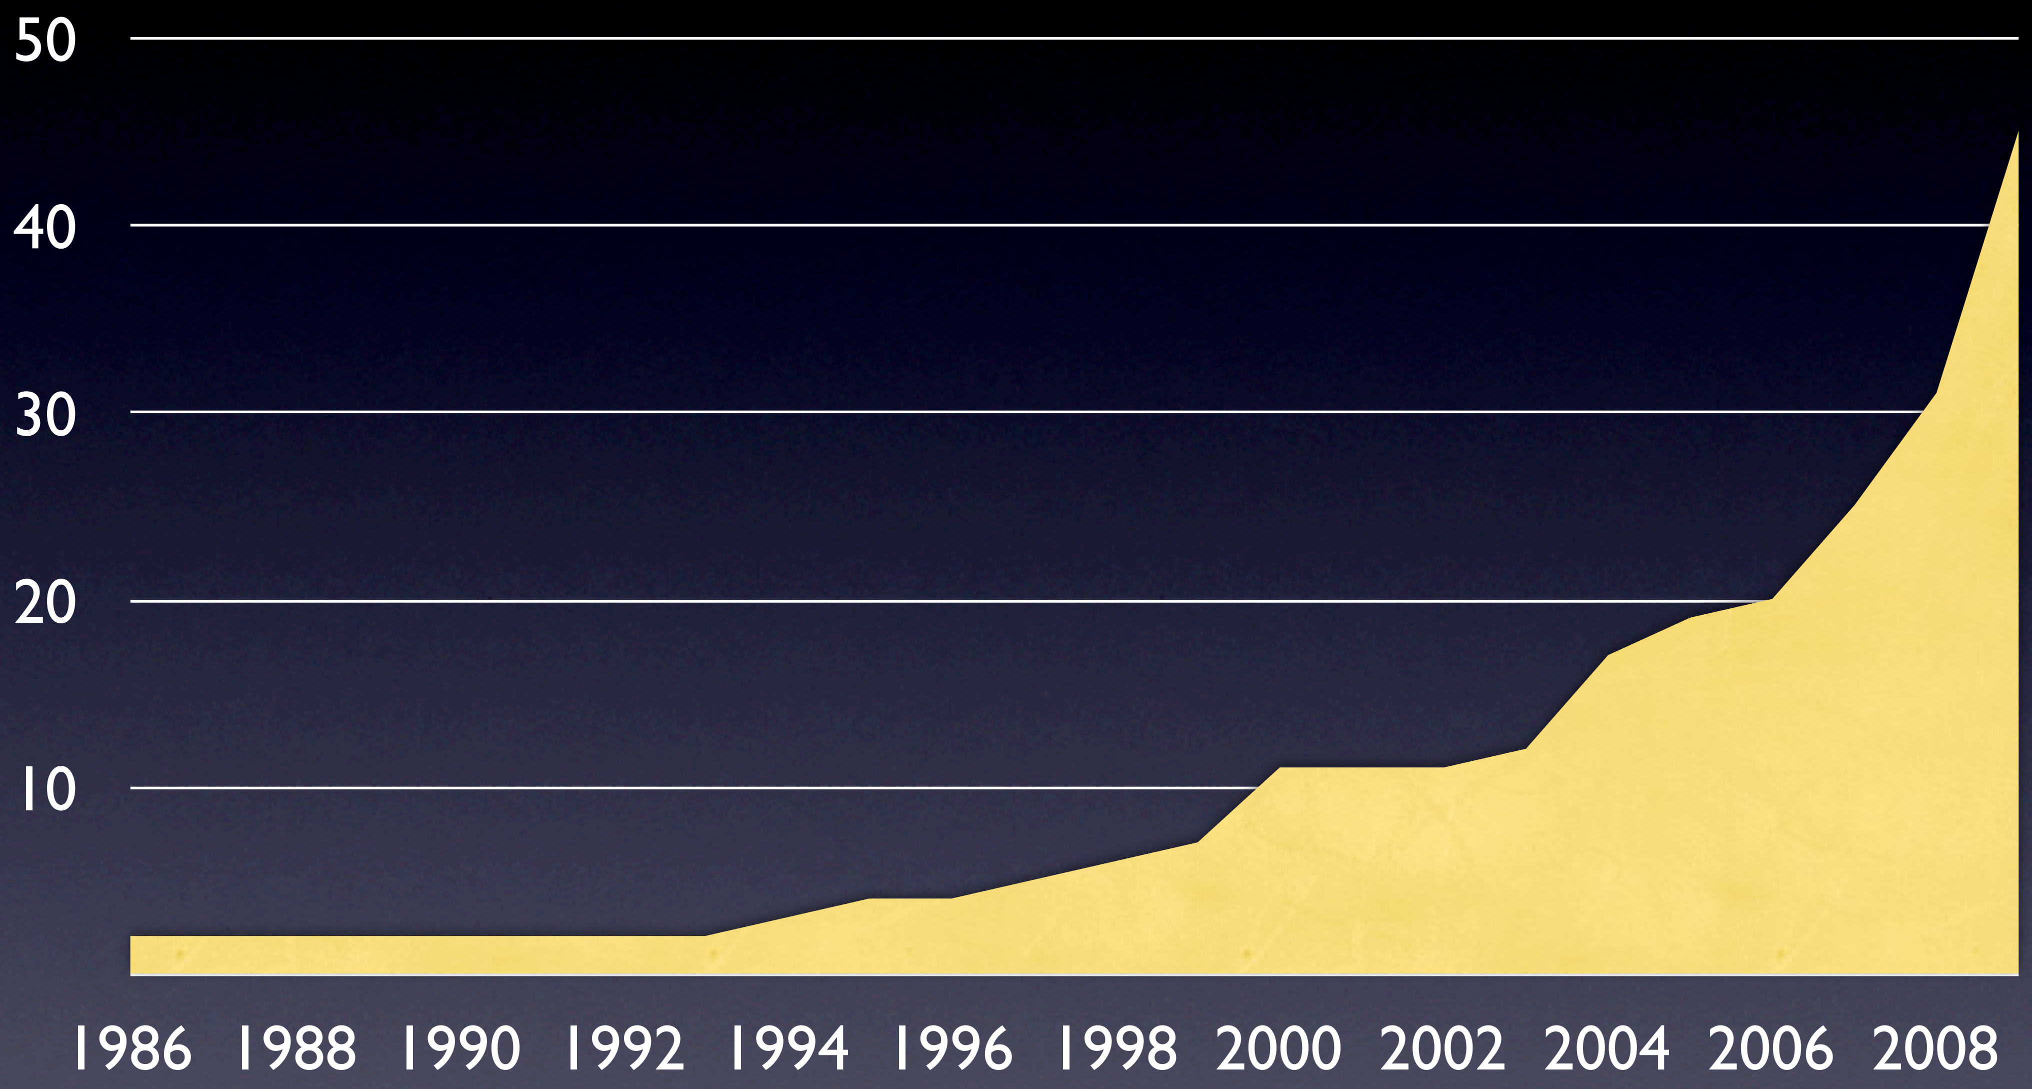

* That I can find information about

LCA 2010

The Survey

- Wasn't perfect.
- Contacted 25 projects, 12 responses.
- Will try again with different questions, cooler website.

The questions:

- What is the name of your project?
- How would you describe your software and what it does in a sentence or two?
- Who is the target user or audience for your database? Do you have any case studies to share?
- Is there a proprietary work-alike or equivalent to your open source database?
- What's the best mailing list for users of your database to subscribe to?
- What's the best mailing list for developers of your database to subscribe to?
- What's the best document for new developers to read if they want to get involved?

- What revision control system does your project primarily use?
- What motivated you to create a new project, rather than join an existing project?
- Do you have a roadmap for the next year? If so, what is it?
- Does anyone provide commercial support for your software?
- What languages are drivers available in, and/or what protocols does your database support? Are they up to date?
- Do you need help with any particular drivers?
- Is there some question I should have asked?
- What feature(s) sets your project apart from your peers?

And I did my own research...

Means of comparison

Database model

Infrastructure features

Development style

Models: defining what operations you'll likely perform on the data

Relational Database models

OLTP: Transaction-oriented

Embedded: Bundling, simplicity, testing

Column: Data warehouses

MPP: Massively Parallel

Streaming: Query streams, not storage

Relational Database Models

OLTP	Embedded	Column-store
CUBRID MySQL (InnoDB) PostgreSQL	H2 HSQLDB SQLite	MonetDB LucidDB C-store/Vertica (Cassandra Hbase)

non-Relational Database models

Flatfile: See Tin (http://tr.im/KNFp)

Key-value: map-reduce, fault-tolerance, caching

Multi-value: Multi-dimensional - GT.M

Graph/Triple-store: Relationship queries

Document-oriented: Semi-structured data

non-Relational Database Models

Key-value	Graph/ Triple-store	Document
BerkeleyDB Cassandra Hbase Memcached Riak Redis TokyoCabinet ydb	Neo4j 4store Parliament	CouchDB BerkeleyDB-XML MongoDB

LCA 2010

infrastructure features:

“distributed”

memory

HA

“Distributed”

Partitioning/ Sharding	Replication	
Cassandra Hbase Voldemort Riak MySQL	BerkeleyDB CouchDB Cassandra MySQL PostgreSQL Riak	Scalaris Voldemort HyperTable HBase Memcached MNesia

Memory vs Disk

In-memory*	Configurable	Disk
Memcached Scalaris Redis	Cassandra Hbase HyperTable MNesia	Everyone else

*This is databases existing solely in memory and being unable or never persisting to disk.

High Availability

Node failover
Cassandra HBase Riak

Otherwise, use one or more of: heartbeat, DBRD, filesystem replication, etc.

Sustainable open source development is code + community.

Code
Development Model

Core + modules	Monolithic	Infrastructure
Drizzle LucidDB PostgreSQL	GT.M Ingres CUBRID	Memcached Redis Scalaris

Community Development Model

Benevolent Dictator	Feature driven	Small Group	A mix
Redis XtraDB MckoiDDB	Apache Derby InfiniDB SmallSQL	CouchDB MonetDB Riak	LucidDB Drizzle H2 PostgreSQL

Plans for the data

- Attempt to update Wikipedia
- Talk to people who write real surveys
- Contacting more projects
- http://ossdbsurvey.org

The Future!

SO YOU'RE SAYING PEOPLE WILL "TWEET" WHAT THEY'RE EATING FOR BREAKFAST?
AND "UPLOAD" PICTURES OF THEIR BREAKFASTS TO A "FACEBOOK"?
AND OTHER PEOPLE WILL LOOK AT THE BREAKFASTS AND MAKE COMMENTS?
SORRY TO BURST YOUR BUBBLE, DUDES, BUT YOU ASKED. YES, THAT'S THE FUTURE.
NO OFFENSE, FUTURE MAN, BUT IS EVERYONE IN YOUR TIME RETARDED?

Protocols

How client/server communication happens

LucidDB, H2 -> PostgreSQL protocol

Sphinx -> MySQL protocol

Tokyo Cabinet / Tyrant -> memcached protocol

Verification

- ‘memcapable’ certifies memcached implementations
- Need automated, repeatable tests for complex systems (Cucumber?)
- More people connections between projects

Databases.
Talking to each other.

Thrift -> ThruDB

http://code.google.com/p/thrudb/

LCA 2010

ENDpoint

Thanks go to:

- Sheeri Cabral
- Josh Berkus
- Brian Aker
- Monty Taylor
- Stewart Smith
- Mark Atwood
- J Chris Anderson
- Jan Lehnardt
- Rick Hillegas
- Salvatore Sanfilippo
- Martin Kersten
- Robin Schumacher
- Vadim Tkachenko
- Justin Sheehy
- Nicholas Goodman, John Sichi, Joseph A. di Paolantonio
- Jay Pipes
- Tobias Downer
- Thomas Mueller
- Scott Deckelmann

LCA 2010

END point

Questions?

This work by Selena Deckelmann is licensed under a Creative Commons Attribution-Noncommercial-Share Alike 3.0 United States License.

LCA 2010

ENDpoint