Architecting for the Cloud

Horizontal Scalability via Transient, Shardable, Share-Nothing Resources

Adam Wiggins

QCon 2009

heroku

Heroku is home to over 40,000 applications

Automatically scaling apps - without code changes

Enabling factors:

Enabling factors:

- Virtualization

Enabling factors:

- Virtualization
- Cloud (virtualization as a service)

Cloud is about
horizontal scalability

Scale out instead of up to avoid the ceiling of Moore's law

Taking advantage of cloud: shardable resources

Resources, aka software infrastructure:

- Database
- Caching
- HTTP router
- Message bus

The father of modern shardable resources:

memcached

{ }

“hashtable in the sky”

Built by one of the first web-scale produts: LiveJournal

Facebook:
800 memcached servers supplying 28 terrabytes of memory

http://www.facebook.com/note.php?note_id=39391378919

Pricinples

- Transient
- Shardable
- Share-nothing

Transient

Any node in the cluster can be lost

Shardable

Client lookup by hashring

12
1
2
3
4
5
6
7
8
9
10
11

Share nothing

Nodes are unaware of each other

Is memcached is cheating?

CouchDB

CouchDB

- Document database

CouchDB

- Document database
- Eventual consistency

CouchDB

- Document database
- Eventual consistency
- MVCC

Eventually consistent

Multiversion concurrency control instead of locking

Comparable to a distributed source control system

Transient

No master server

Shardable

Clients can go to any server

Share nothing

Nodes communicate only when asked to replicate

CouchDB: The Definitive Guide

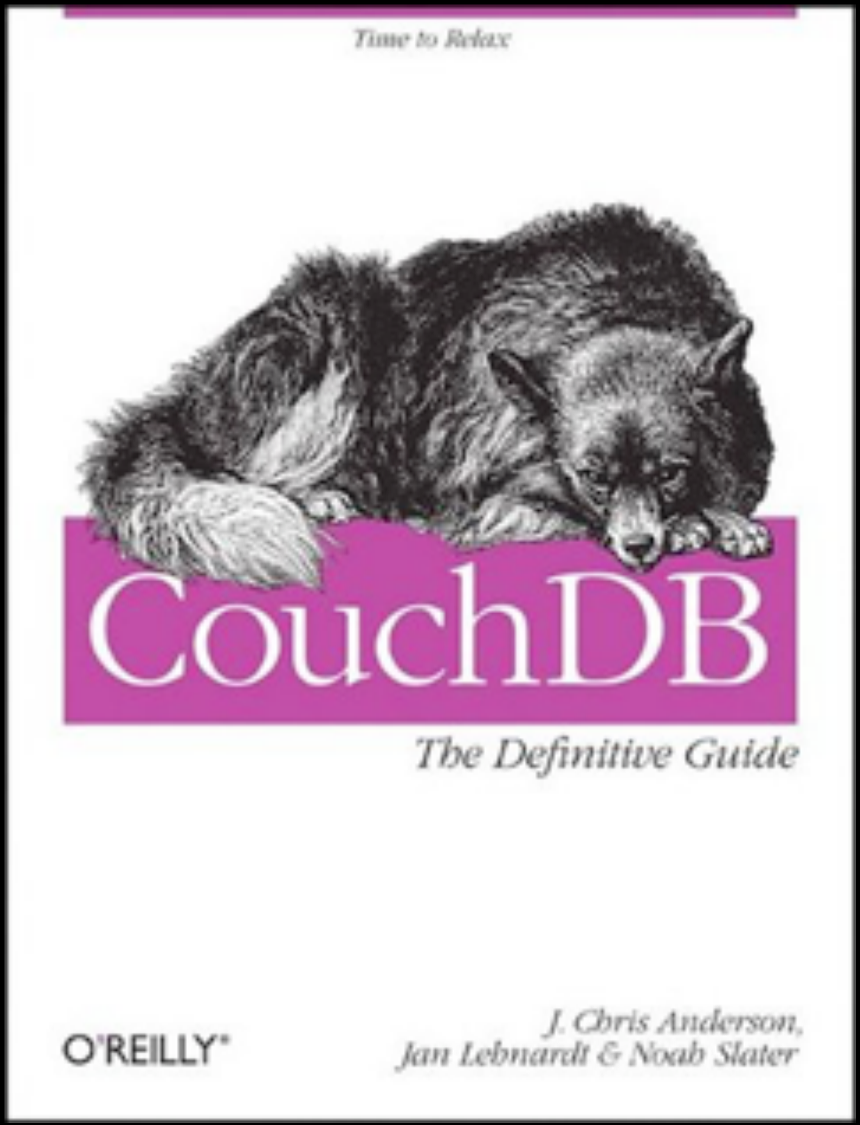

http://books.couchdb.org/relax/

Hadoop

big data processing

- MapReduce
- Cut big data into small chunks
- Cut big work into distributable jobs

http://hadoop.apache.org/

Redis

key-value store

- Like memcached with persistence
- Shards with hashring
- Lists and sets
- Extremely fast and lightweight

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

Varnish http cache

- Like Squid, but horizontally scalable
- Combine with `ngx_http_upstream_consistent_hash` for hashring-style access

http://varnish.projects.linpro.no/

http://wiki.nginx.org/NginxHttpUpstreamConsistentHash

RabbitMQ message bus

- Job queueing
- Cluster broadcast via exchanges
- Cross-language

http://www.rabbitmq.com/

Erlang

functional language

- High concurrency
- No mutable variables
- Lightweight processes

http://www.erlang.org/

Horizontal scalability promises to shatter the glass ceiling of vertical scale

...but only if we architect resources to be transient, shardable, and share-nothing

The End.

http://heroku.com

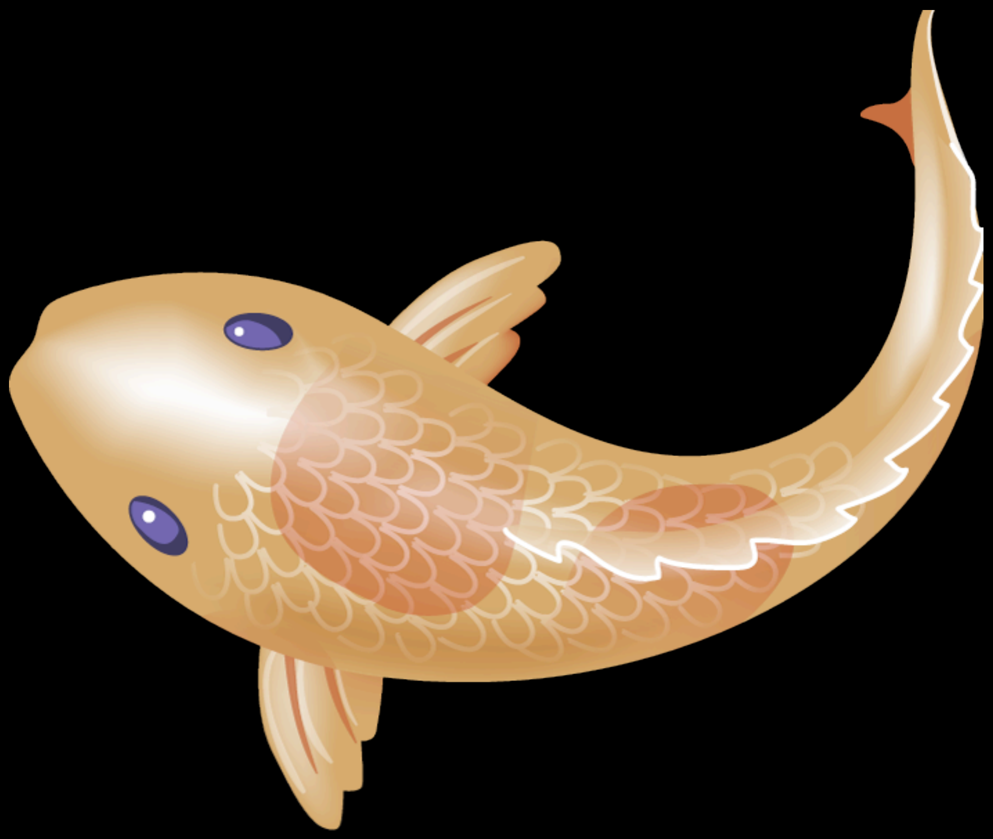

Adam Wiggins

http://adam.blog.heroku.com