



USES

CASES

OR

DON'T USE MY DATABASE

PRESENTATION FOR NOSQL ROADSHOW

HELLO,
MY NAME
IS @dstroyallmodels

IAN PLOSKER

Technical Lead, International Operations at Basho Technologies

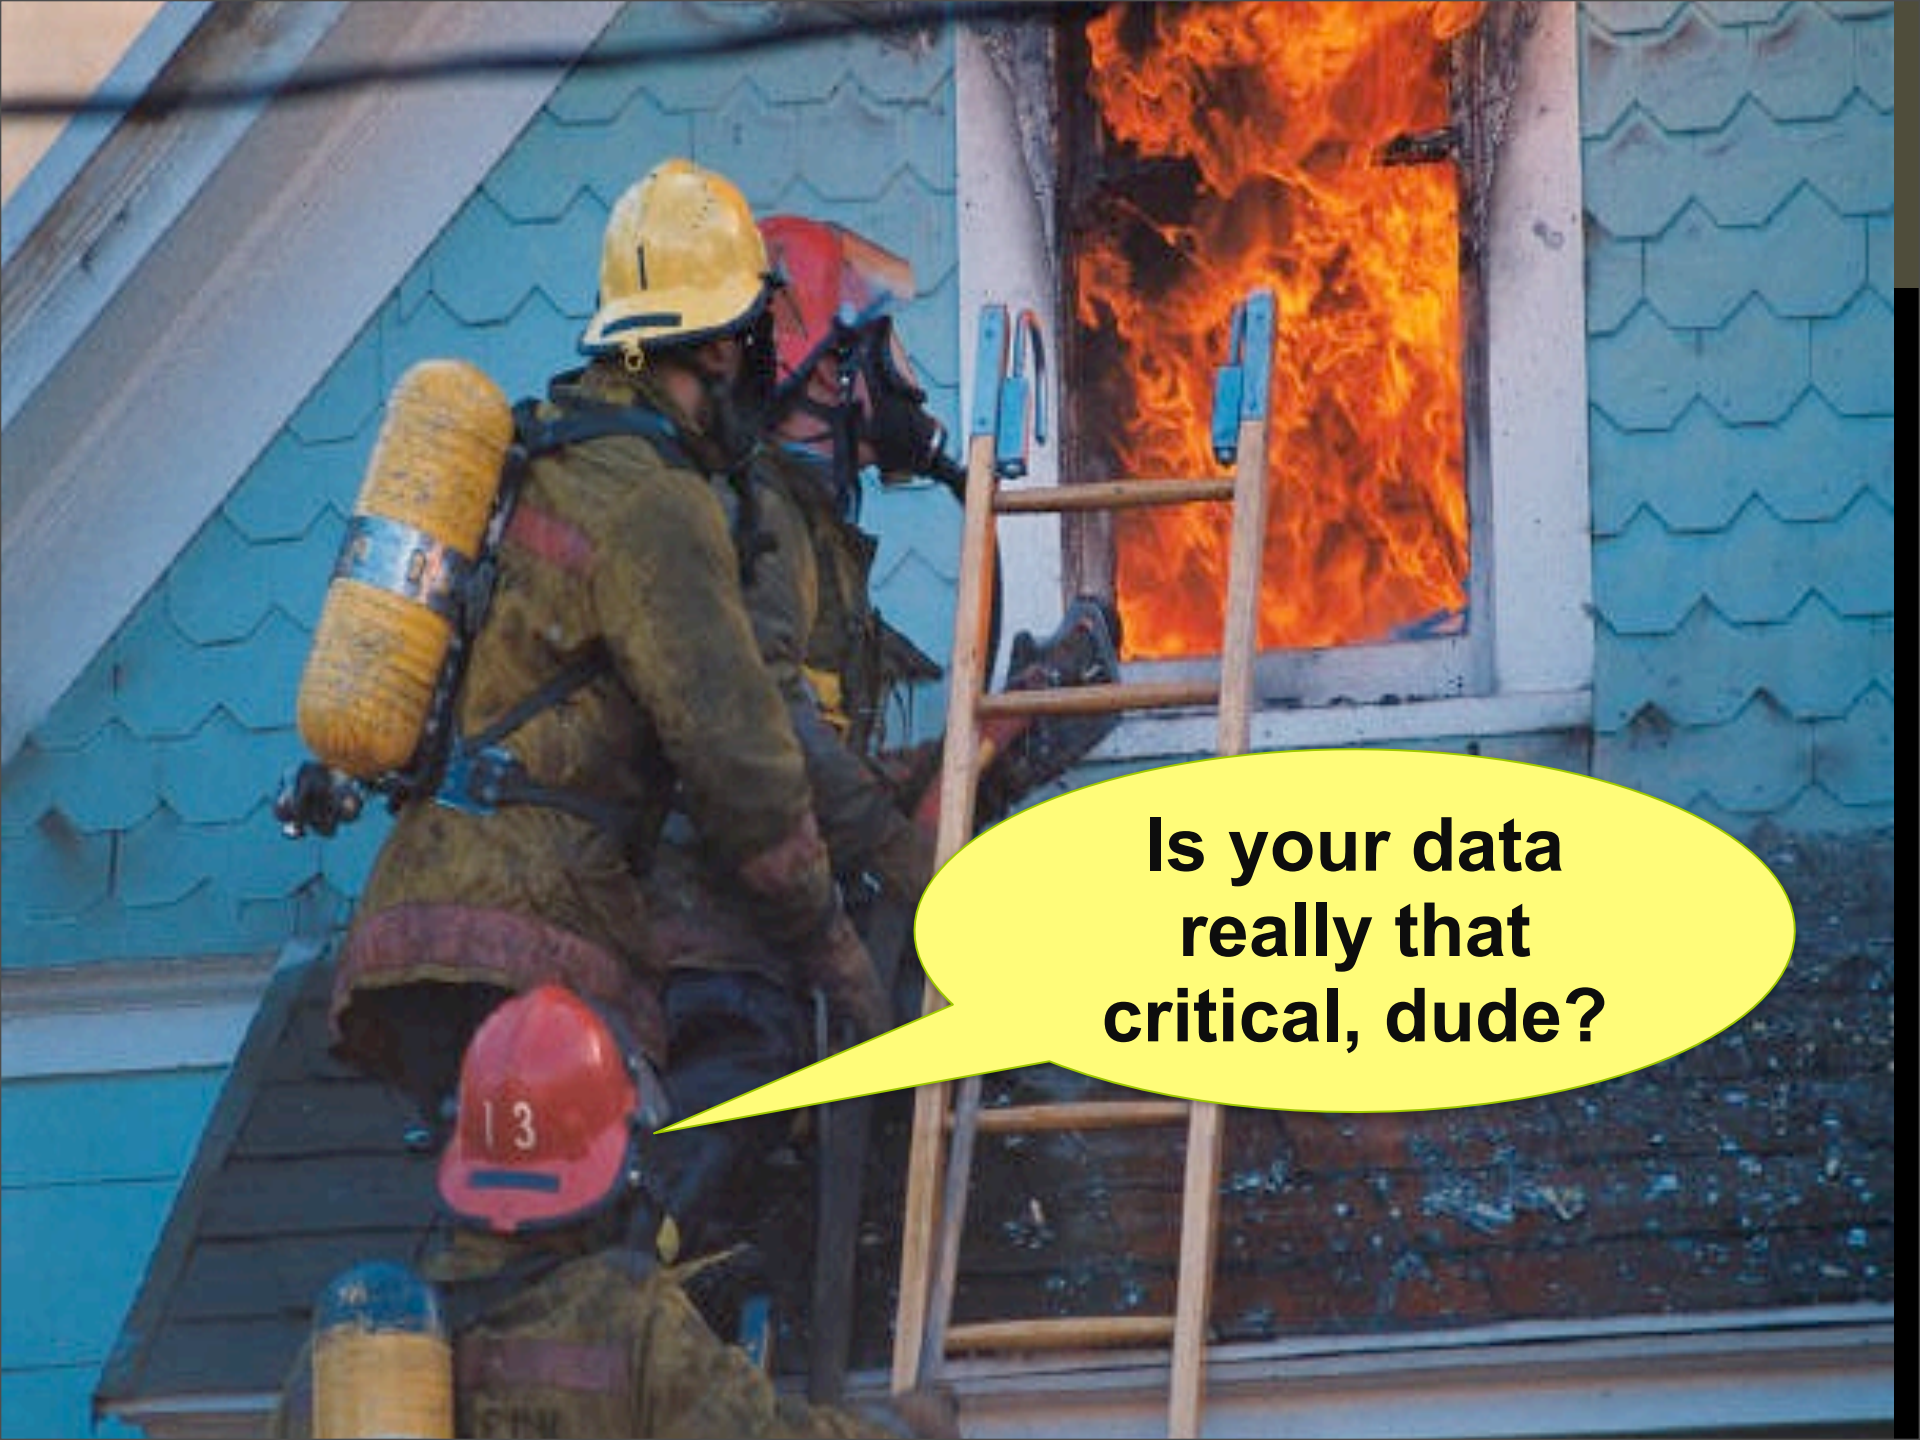


**CAN WE PLEASE NOT TALK ABOUT
BIG DATA?**

LET'S TALK ABOUT

CRITICAL DATA





**Is your data
really that
critical, dude?**

ON OFFER

- what... *is NoSQL? ... Riak?*
- who ... *uses Riak?*
- why ... *should I NOT use Riak?*
(Or other NoSQL for that matter?)

Enter the NoSQL Jungle

TRUTH



YOU ARE HERE



what are you prepared to trade?

Would you trade....

Your current, familiar consistency model
for

a somewhat alien, but perfectly safe, consistency model and
better availability?

Would you trade....

Storage space

for

low latency?

Would you trade....

late night heroics

for

high availability?

Would you trade....

35 years of RDBMS success
for

technology used by a few companies of which
you may have heard?



Linked in.

comcast.



github
SOCIAL CODING

Would you trade....

Two million developers worldwide
for



*"It's a piece of plumbing; it has never been
a root cause of any of our problems."*

Coda Hale, Yammer

Coda Hale, Yammer

**BEFORE YOU ANSWER,
CONSIDER....**





**Can anyone
call RDBMS
a failure?**





IF ANYONE
MAKES THESE
TRADE-OFFS
SEEM **EASY**.....

A low-angle shot of two men in business suits walking past a modern, multi-story office building. The man on the left is Black, wearing a dark suit, light blue shirt, and yellow patterned tie, carrying a briefcase. The man on the right is white, wearing a dark suit, white shirt, and red patterned tie, also carrying a briefcase. A large, bright green speech bubble is superimposed over the upper part of the image, containing the text 'THEY ARE LYING!!!'.

THEY ARE
LYING!!!

Lie # 1:

Distributed Systems: Desirable Properties

- Highly Available
- Low Latency
- Scalable
- Fault Tolerant
- Ops-Friendly
- Predictable

**Any new app must use this
newfangled NoSQL.**

Lie # 2:

Distributed Systems: Desirable Properties

- Highly Available
- Low Latency
- Scalable
- Fault Tolerant
- Ops-Friendly
- Predictable

**Non-distributed databases
provide these properties.**

Lie # 3:

MongoDB

Cassandra

CouchDB

Riak

Neo4J

Voldemort



**Words like “distributed”
have “relative” meaning.**

Lie # 4: The Big One

I say I want:

- 1. safety**
- 2. scale out**
- 3. speed**



When really I want:

- 1. speed**
- 2. speed**
- 3. speed**

The lie developers tell themselves.

REMEMBER

At small scale everything works.

REMEMBER

At small scale everything works.

**At large scale
things start to
break**




RIAK – THE TECHNOLOGY

An open-source, globally distributed database with trade-offs

DISTRIBUTE

Replicate client data to
your Riak clusters
anywhere in the world

Clients run Riak
clusters in
Your data centers



RIAK DETAILS

Dynamo-inspired

modular – core plus storage (or...)

1) bitcask

2) levelDB

Document-oriented AND *key-value*

Tuneable CAP values

Open-source and commercial



BITCASK AND THE LATENCY TRADEOFF

Low Latency: All reads = hash lookup + 1 seek

TRADEOFF: All keys must fit in memory

Who cares about latency?

MARKET SNAPSHOT

U.S.	EUROPE	ASIA
WINNIE: 8,364.39 +10.30	FTSE 100: 6,341.39 +10.30	Nikkei 225: 11,231.80 -0.23%
TOPIX: 717.24 +10.30	DAX: 7,172.40 +10.30	Hang Seng: 18,000.00 +1.20%
EUR: USD 1.3326 0.6564%	Nasdaq 2,441.54 -0.75%	Dow 11,231.80 -0.23%
		US\$ 699.11 +1.10%

Bloomberg Our Company | Professional | Anywhere [Web Your Queue](#) [Sign in](#)

[Home](#) [Quick](#) [News](#) [View](#) [Markets](#) [Personal Finance](#) [Sustainability](#) [TV](#) [Radio](#) [Search News, Quotes and OnlineQ](#)

Related News: [Economy](#) · [Law](#) · [Canada](#) · [U.S.](#) · [Bonds](#) · [Currencies](#) · [Finance](#) · [Insurance](#) · [Real Estate](#)

Want to save this for later? [Add it to your Queue!](#)

Secret Fed Loans Gave Banks Undisclosed \$13B

By Bob Ivry, Bradley Keoun and Phil Kuntz · Nov 27, 2011 4:01 PM PT [SEND TO QUEUE](#)

Bloomberg Markets Magazine

[Recommend](#) [Tweet](#) [Share](#) [Print](#) [Email](#)

The Federal Reserve and the big banks fought for more than two years to keep details of the largest bailout in U.S. history a secret. Now, the rest of the world can see what it was missing.

The Fed didn't tell anyone when banks were in trouble so deep they required a combined \$1.2 trillion on Dec. 5, 2008, their single neediest day. Bankers didn't mention that they took tens of billions of dollars in emergency loans at the same time they were assuring investors their firms were healthy. And no one calculated until now that banks reaped an estimated \$13 billion of income by taking advantage of the Fed's below-market rates, Bloomberg Markets magazine

[Enlarge image](#)

More Stories

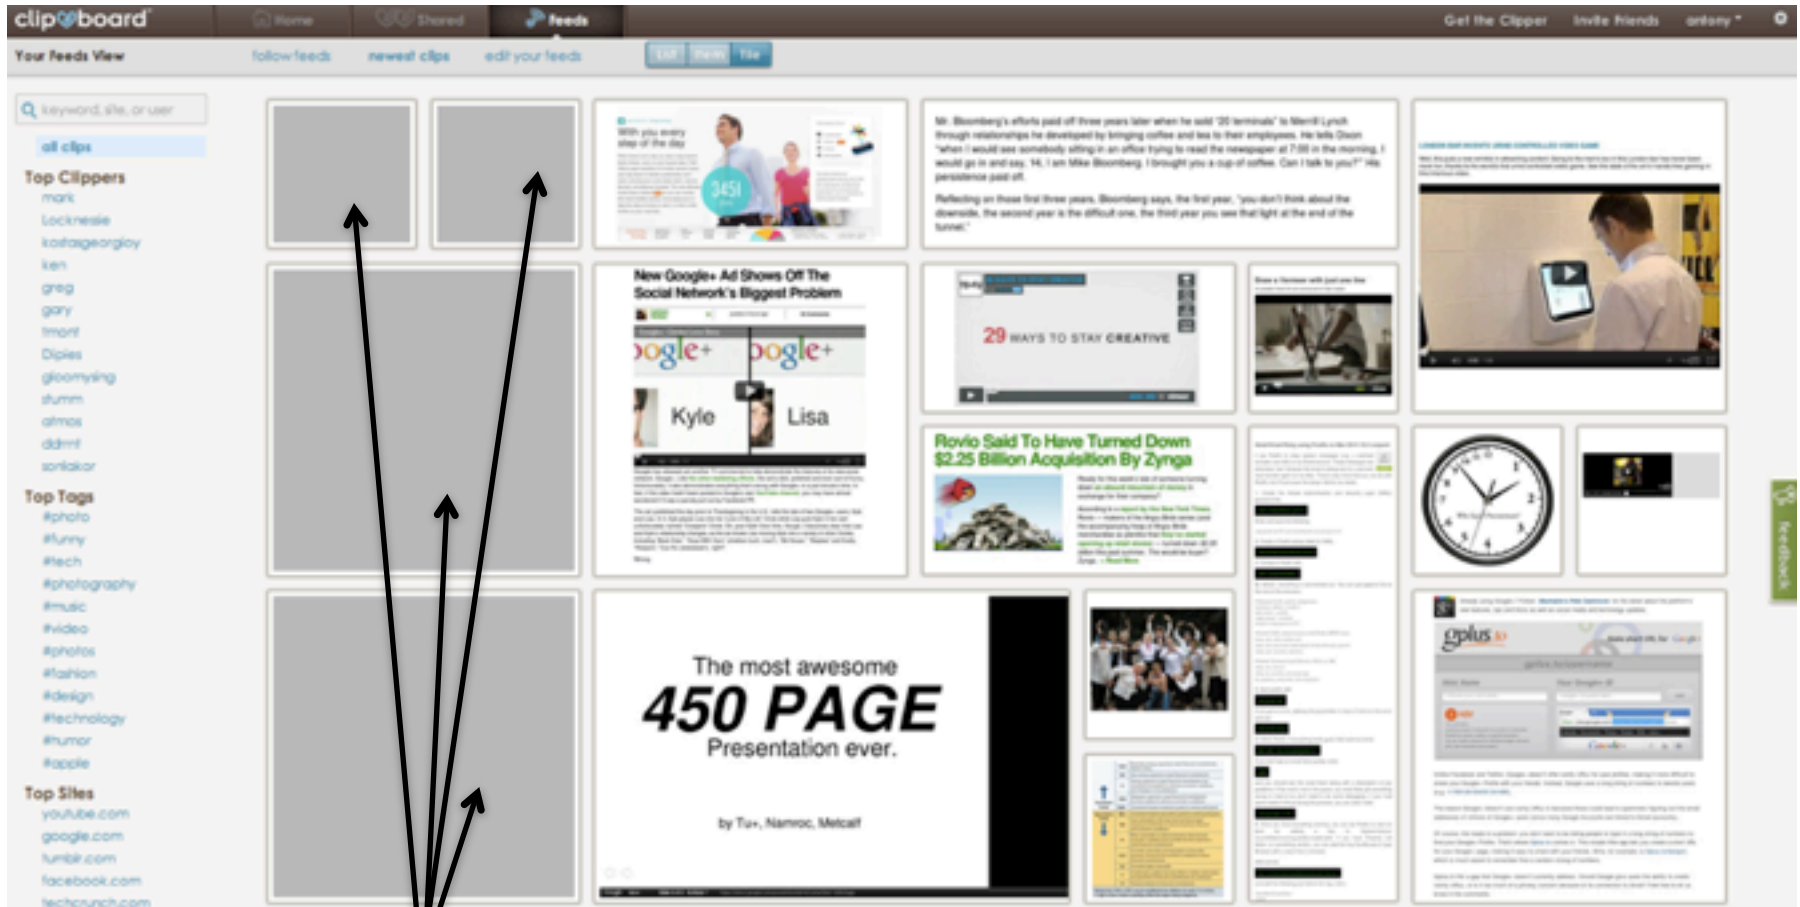
- [U.S. Stock-Index Futures Gain After Report on IMF Planning loan to Italy](#)
- [Texas to Ask Supreme Court for Stay of Mass](#)
- [Euro Advances After Report of IMF Italy Loan Plan: Aussie, Kiwi Strengthen](#)
- [Asia Stocks, U.S. Futures Rally on Italy](#)

[Stories](#) [Rate These More News](#)

Advertisement

SOA

Who cares about latency?



Sometimes high latency looks like an outage to the end user.

Who cares about latency?



USE CASES

"We would not have been able to build this application without node.js and Riak Search."

Clipboard

1. Cloud Storage (S3)
2. Low-latency interactive apps
3. Smartphone apps
4. CDN
5. Medical
6. Session Store
7. Gaming
8. Mobile Infrastructure

STUFF CO.





**Real-time Search
US Retailer Online Inventory**

Mochi – casual interactive gaming



yammer



Social Media Networks

“REAL-TIME” APPS

Advantages: fast, reliable, scalable. Distribute to a *global audience*.




DANISH HEALTH SERVICES



Advantages: Doctors manage prescriptions from highly-available servers and, soon, **smartphones**, 95% cheaper than Oracle, easier to run

Telco distributes content to 10 data centers



MULTI-TENANT OBJECT STORAGE

Advantages: build competitive services to Amazon, reliable storage, easy to operate

Write to Riak
Next...sync with each other



SMARTPHONE APPS

Advantages: better quality of service for write-intensive smartphone users; reduce infrastructure costs



A QUICK LITMUS TEST:
ARE THE TRADEOFFS WORTH IT?

USE RIAK IF YOU NEED...

Predictable latency

Predictable scaling profile

Predictable operations

and if data unavailability costs

thousands of \$\$\$'s/minute

or is even more important.

DON'T USE RIAK IF YOU NEED...

Familiar query patterns

Familiar hiring patterns

Familiar operational problems

and if you

don't have a problem right now.

A full-page photograph of a mountain landscape. In the background, a large, dark, and craggy mountain peak rises steeply against a cloudy sky. The foreground and middle ground consist of rolling hills and valleys covered in vibrant green grass and vegetation. A small, rocky outcrop is visible on the right side of the middle ground. The overall scene is rugged and scenic.

THANK YOU

CREDIT: TONY FALCO