

Embracing Concurrency at scale (it's about time!)

With a guest appearance by Riak.



Justin Sheehy
justin@basho.com

Concurrency Matters



"The free lunch is over."

- Herb Sutter, 2005

Concurrency Matters



You got a free lunch!?



"The free lunch is over."

- Herb Sutter, 2005

New Problems, Old Solutions

Distributed Systems matter now more than ever,
and we must learn from the past to build the future.

New Problems, Old Solutions, New Systems

Distributed Systems matter now more than ever,
and we must learn from the past to build **the future**.



Don't do what I say. (yet)



Working at scale isn't just "more." It is different.

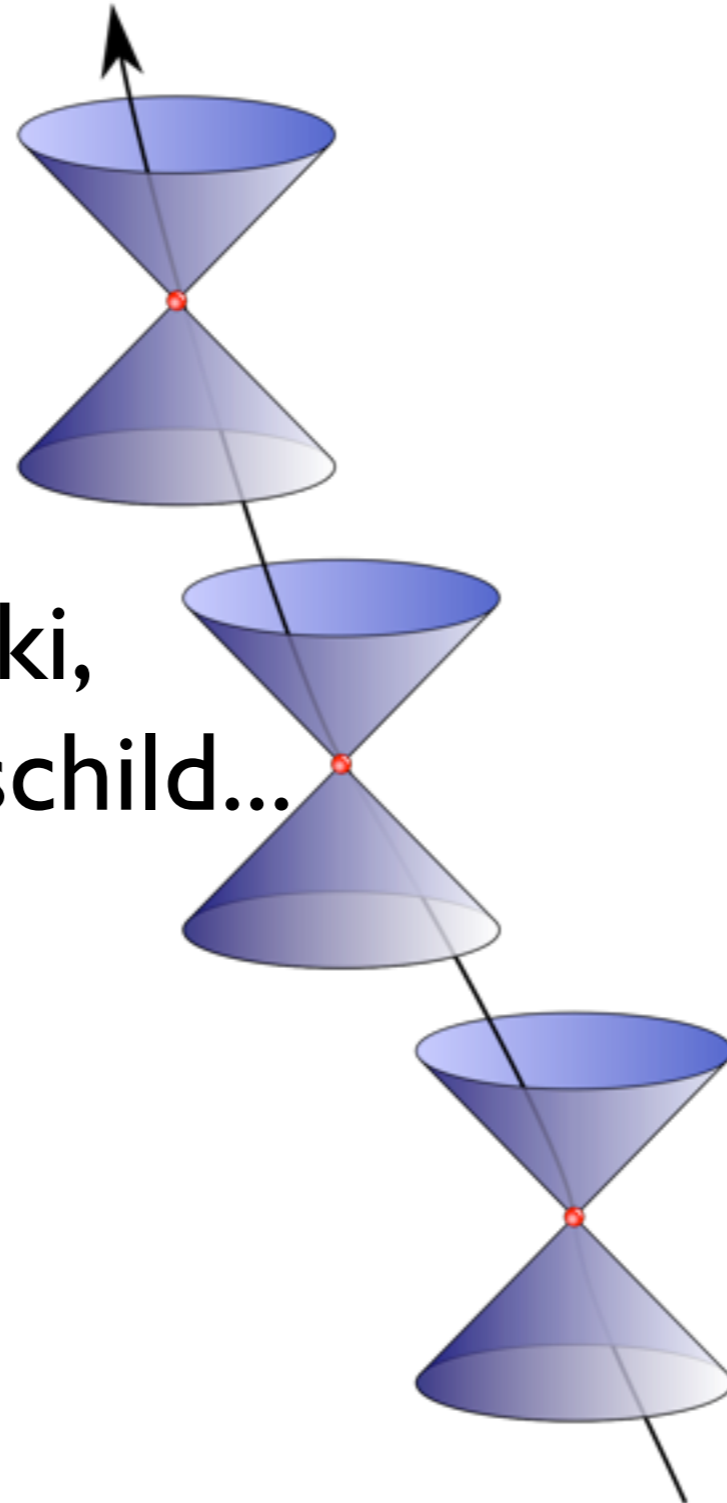
What is Concurrency?

concurrent: occurring at the same time

concurring: agreeing with others


Time is a Hard Problem

Einstein,
Minkowski,
Schwarzschild...



Time in Computing

Lamport, 1978 -- gave us “happened before”



Mattern, 1989 -- closer to Minkowski causality

Time is a Hard Problem

In computing, we like to pretend it's easy.

This is a trap!

Distributed Computing is Asynchronous Computing

Synchrony (distributed transactions) throws away
the biggest gains of being distributed!

There is no “Global State”

You only know about the past -- deal with it!

This sadly often means giving up on ACID.
(globally, not locally)

This is going to hurt!

Atomicity
Consistency
Isolation
Durability

~~Atomicity~~

Consistency

~~Isolation~~

Durability

~~Atomicity~~

Consistency

~~Isolation~~

Durability

Basically
Available

~~Atomicity~~

Consistency

~~Isolation~~

~~Durability~~

Basically
Available

~~Atomicity~~

Consistency

~~Isolation~~

~~Durability~~

Basically
Available
Soft State

~~Atomicity~~
~~Consistency~~
~~Isolation~~
~~Durability~~

Basically
Available
Soft State

~~Atomicity~~
~~Consistency~~
~~Isolation~~
~~Durability~~

Basically
Available
Soft State
Eventually-Consistent

This is a real tradeoff -- if you make it, understand it!

(Eric Brewer, 1997)

Basically
Available
Soft State
Eventually-Consistent

CAP tradeoffs

Consistency

Availability

Partition-Tolerance

Pick any two?

CAP tradeoffs

Consistency

Availability

Partition-Tolerance

This is all about failure.

~~Pick any two?~~

CAP tradeoffs

Consistency

Availability

Partition-Tolerance

Distributed Transactions

(on any real network, this fails)

CAP tradeoffs

Consistency

Availability

Partition-Tolerance

Quorum Protocols &
typical Distributed Databases
(nodes outside the quorum fail)

CAP tradeoffs

Consistency

Availability

Partition-Tolerance

Sometimes allow stale data...

...but everything can keep going.

CAP tradeoffs

“Half your nodes can be down and you can still write to the system.”

Matt Ranney
CTO, Voxer

Sometimes allow stale data...
...but everything can keep going.

CAP tradeoffs

Consis
Availab
Partitic

“Given Riak’s high availability and fault tolerance, we knew that we would not lose any individual’s data, nor lose the ability to offer up a medical history simply because some servers were down.”

Kresten Krab Thorup
CTO, Trifork

Sometimes allow stale data...
...but everything can keep going.

CAP tradeoffs

Consistency

Availability

Partition-Tolerance

This is where BASE leads us.

This is a real tradeoff -- if you make it, understand it!

Basically
Available
Soft State
Eventually-Consistent

Know How You Degrade

Plan it and understand it before your users do.

You might prevent whole system failure if you're lucky and good, but what happens during **partial failure?**

Know How You Degrade

Plan it and understand it before your users do.

You think you know
which parts will break.

Know How You Degrade

Plan it and understand it before your users do.

You think you know
which parts will break.

You are wrong.



Harvest and Yield

harvest: a fraction

data available / complete data

yield: a probability

queries completed / q's requested

in tension with each other:

(harvest * yield) ~ constant

goal: failures cause known linear reduction to one of these

Harvest and Yield

global ACID demands 100% **harvest**
but success of modern applications is
often measured in **yield**

plan ahead, know when you care!



This is a real tradeoff -- if you make it, understand it!

Basically
Available
Soft State
Eventually-Consistent

Eventually-Consistent

It means that you can change some state instead of failing.

It also doesn't mean slow.

Sometimes you go "eventual" in order to go fast.

Eventually-Consistent doesn't mean “not consistent”!

**It means that you can change
some state instead of failing.**

It also doesn't mean slow.

Sometimes you go "eventual" in order to go fast.

**You said you were going to
talk about**

Scalability!

RPC is a scaling antipattern.

Treating remote communication like local function calls is a fundamentally bad abstraction.

- Network can fail after call “succeeds”.
- Data copying cost can be hard to predict.
- Tricks you by working locally.
(and then failing in a real dist sys)
- Prevents awareness of swimlanes.
(and thus causes cascading failure)

Protocols vs. APIs

- Explicit understanding of boundaries.
(trust boundaries, failure boundaries...)
- Better re-use and composition.
(unintuitive but true in the large)
- Asynchronous reality, described accurately.
(see Clojure or Erlang/OTP libraries)

Successful Protocols

Kings of the Web: DNS & HTTP

What do they have in common?

Successful Protocols

Kings of the Web: DNS & HTTP

What do they have in common?

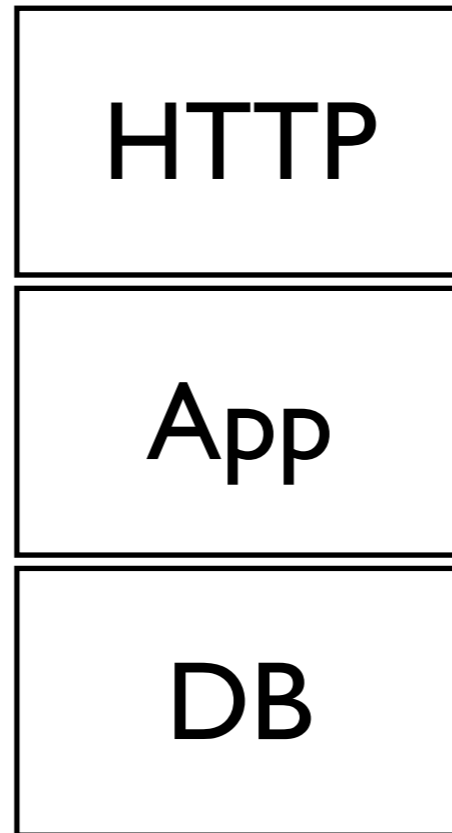
B
A
S
E

The Web

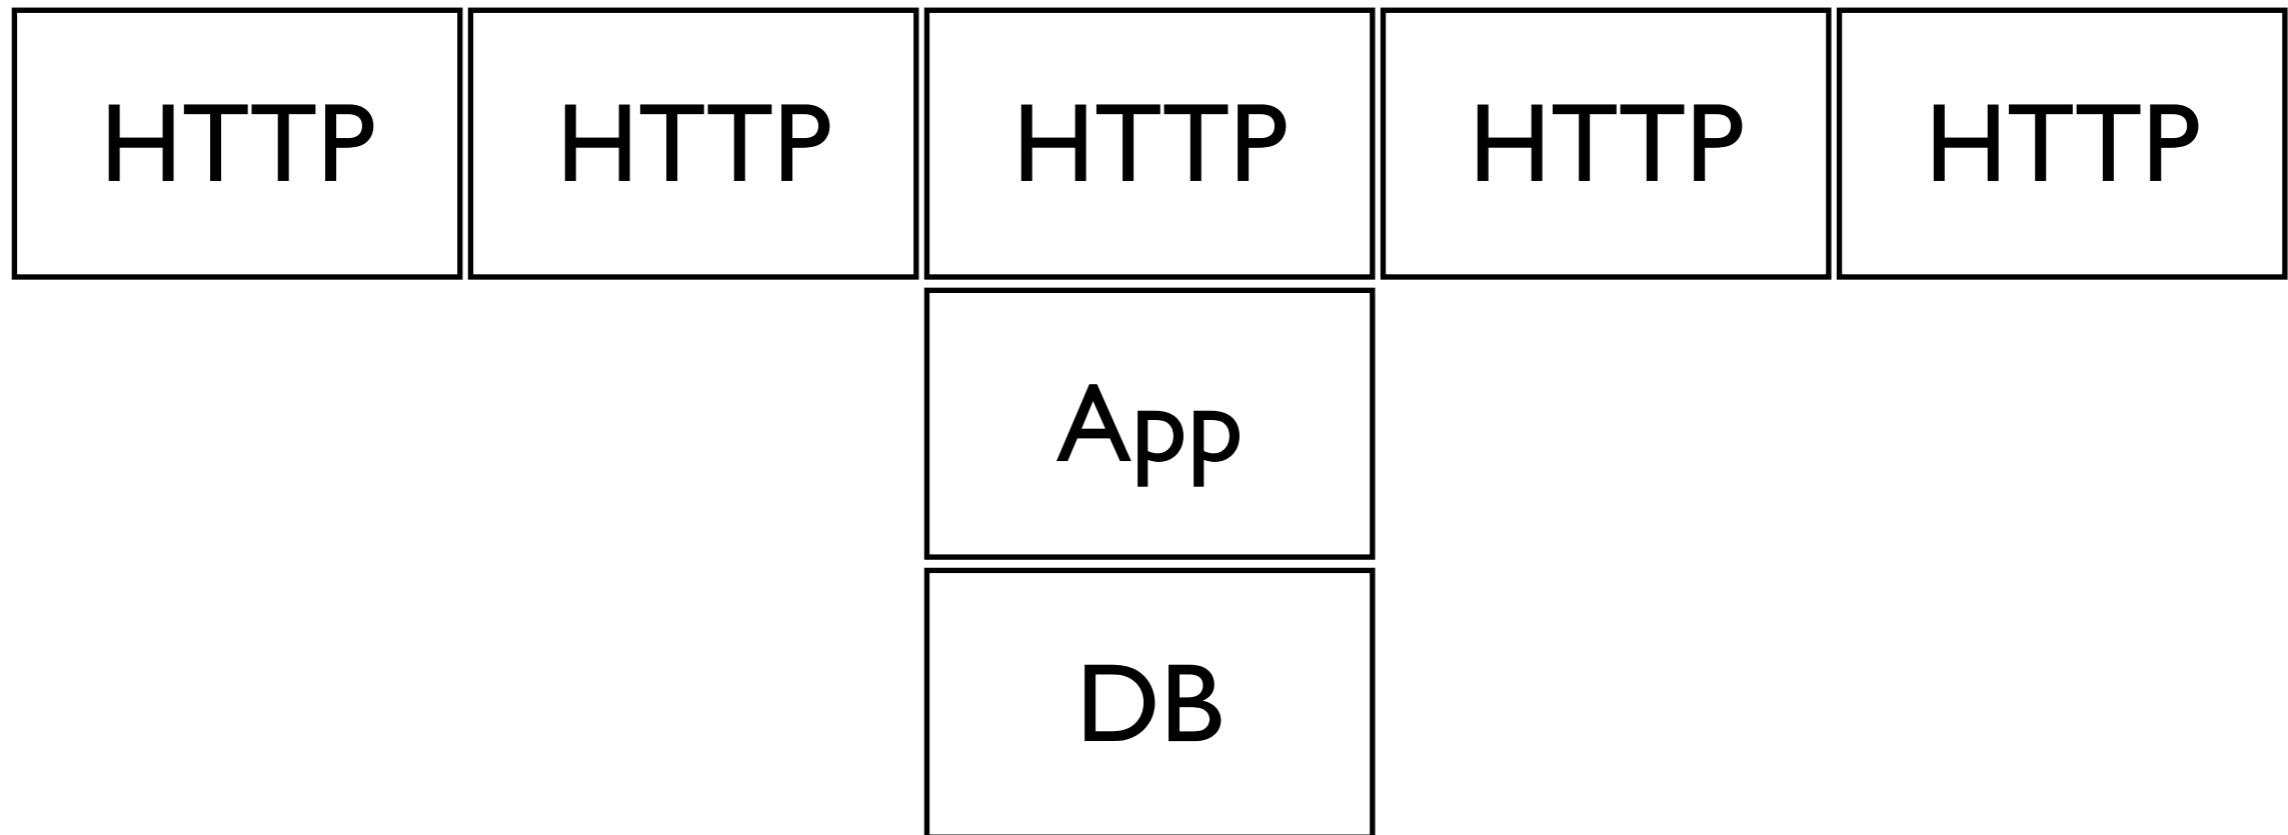
(the second most successful distributed system ever)

- no global state (closest: DNS root & MIME)
- well-defined caching for eventual consistency
 - idempotent operations!
- loose coupling
 - links instead of global relations
 - no must-understands except HTTP

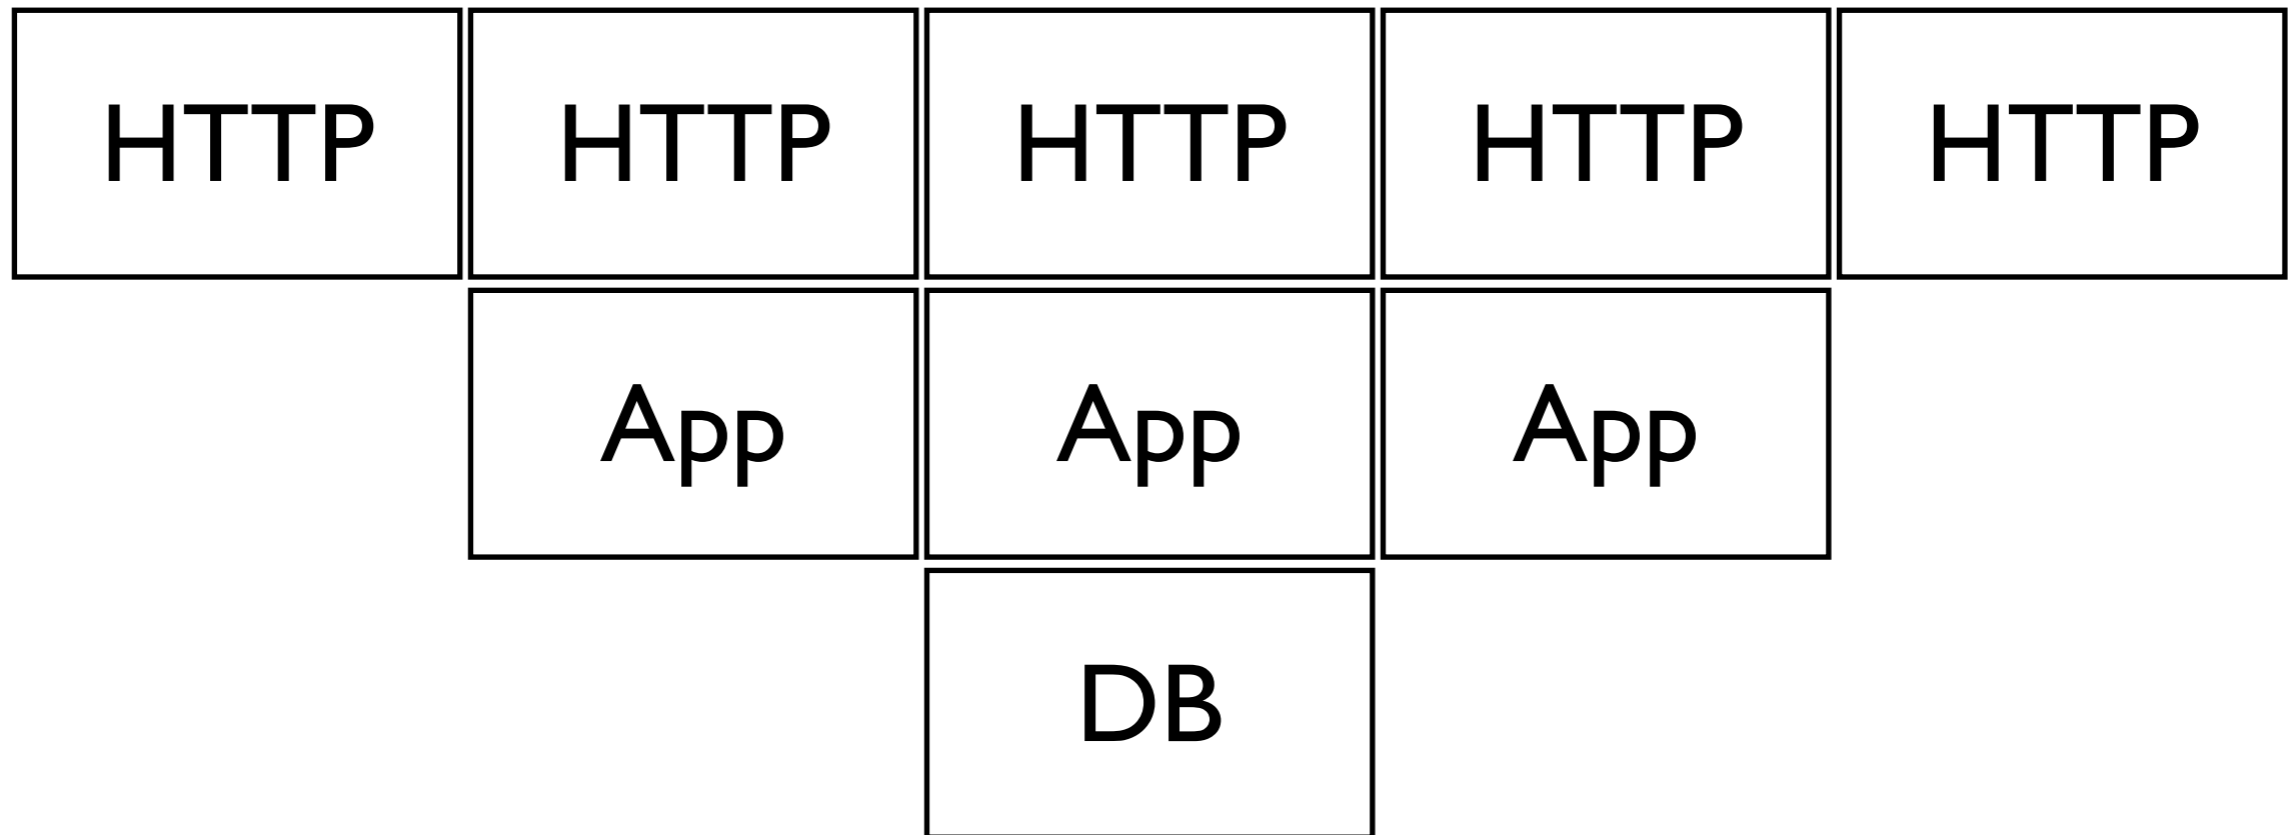
History of Scaling The Web



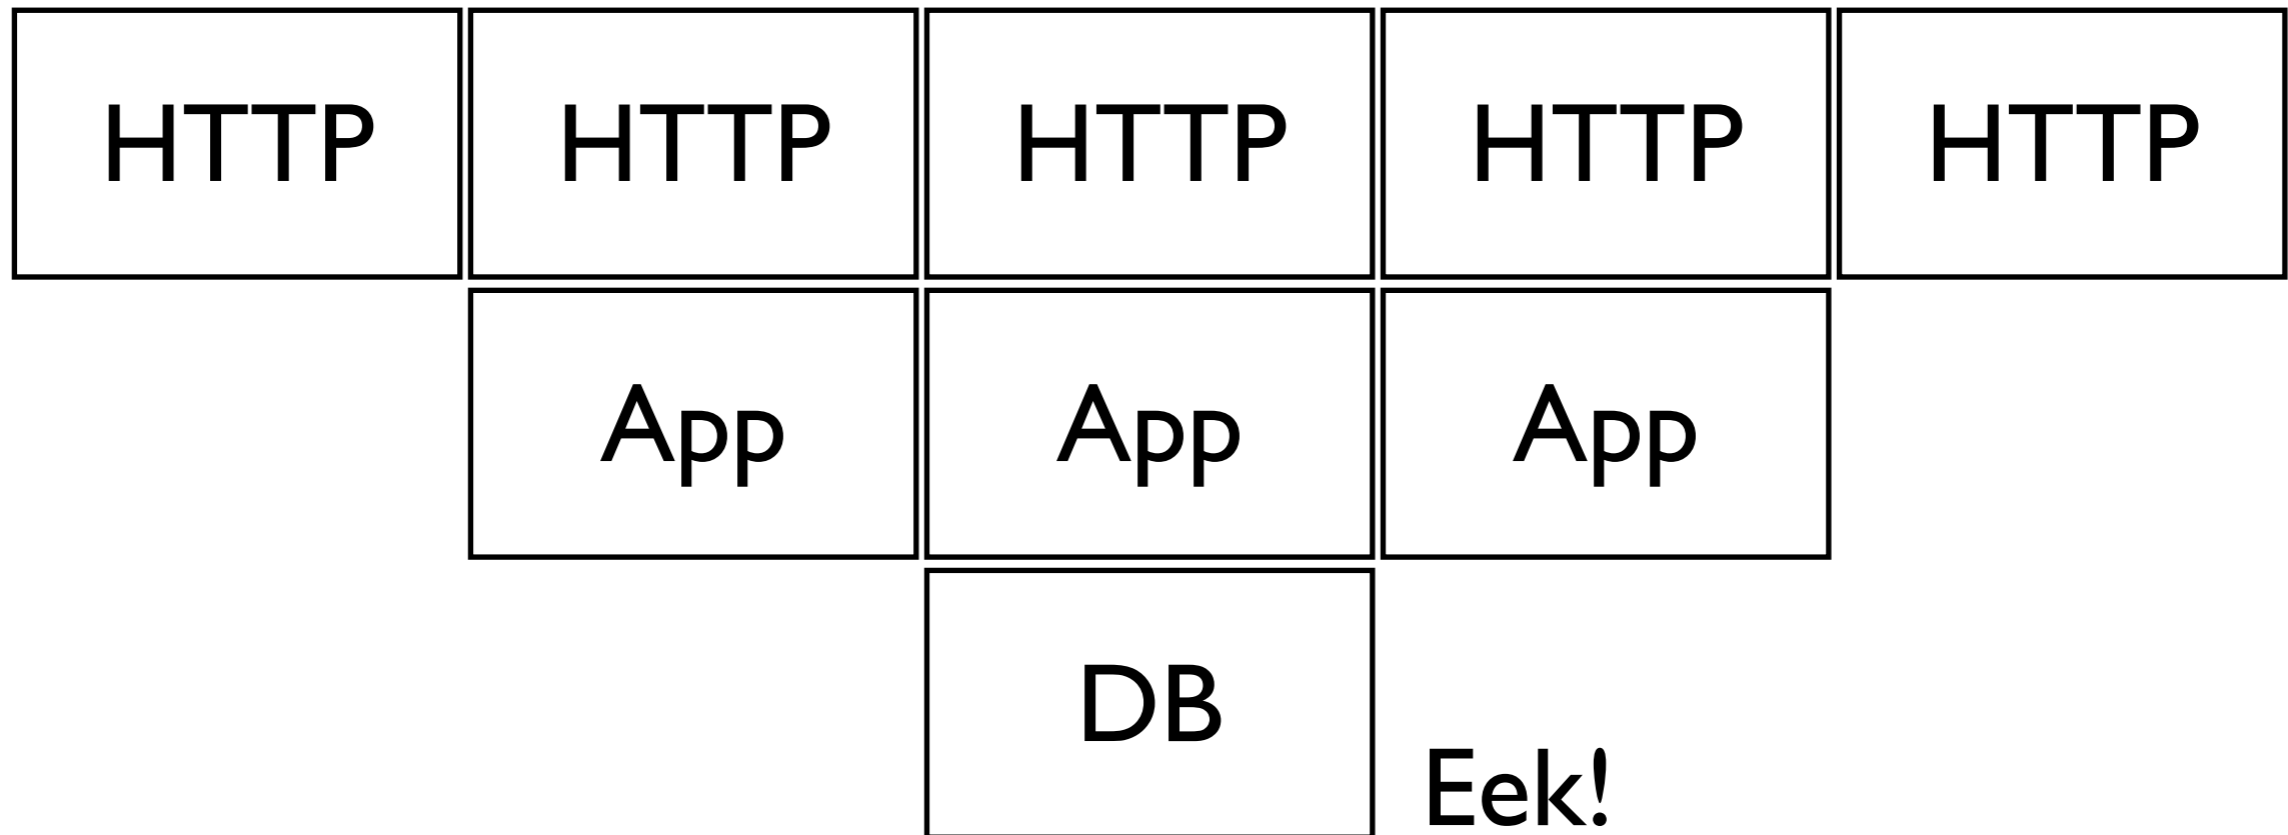
History of Scaling The Web



History of Scaling The Web








History of Scaling The Web

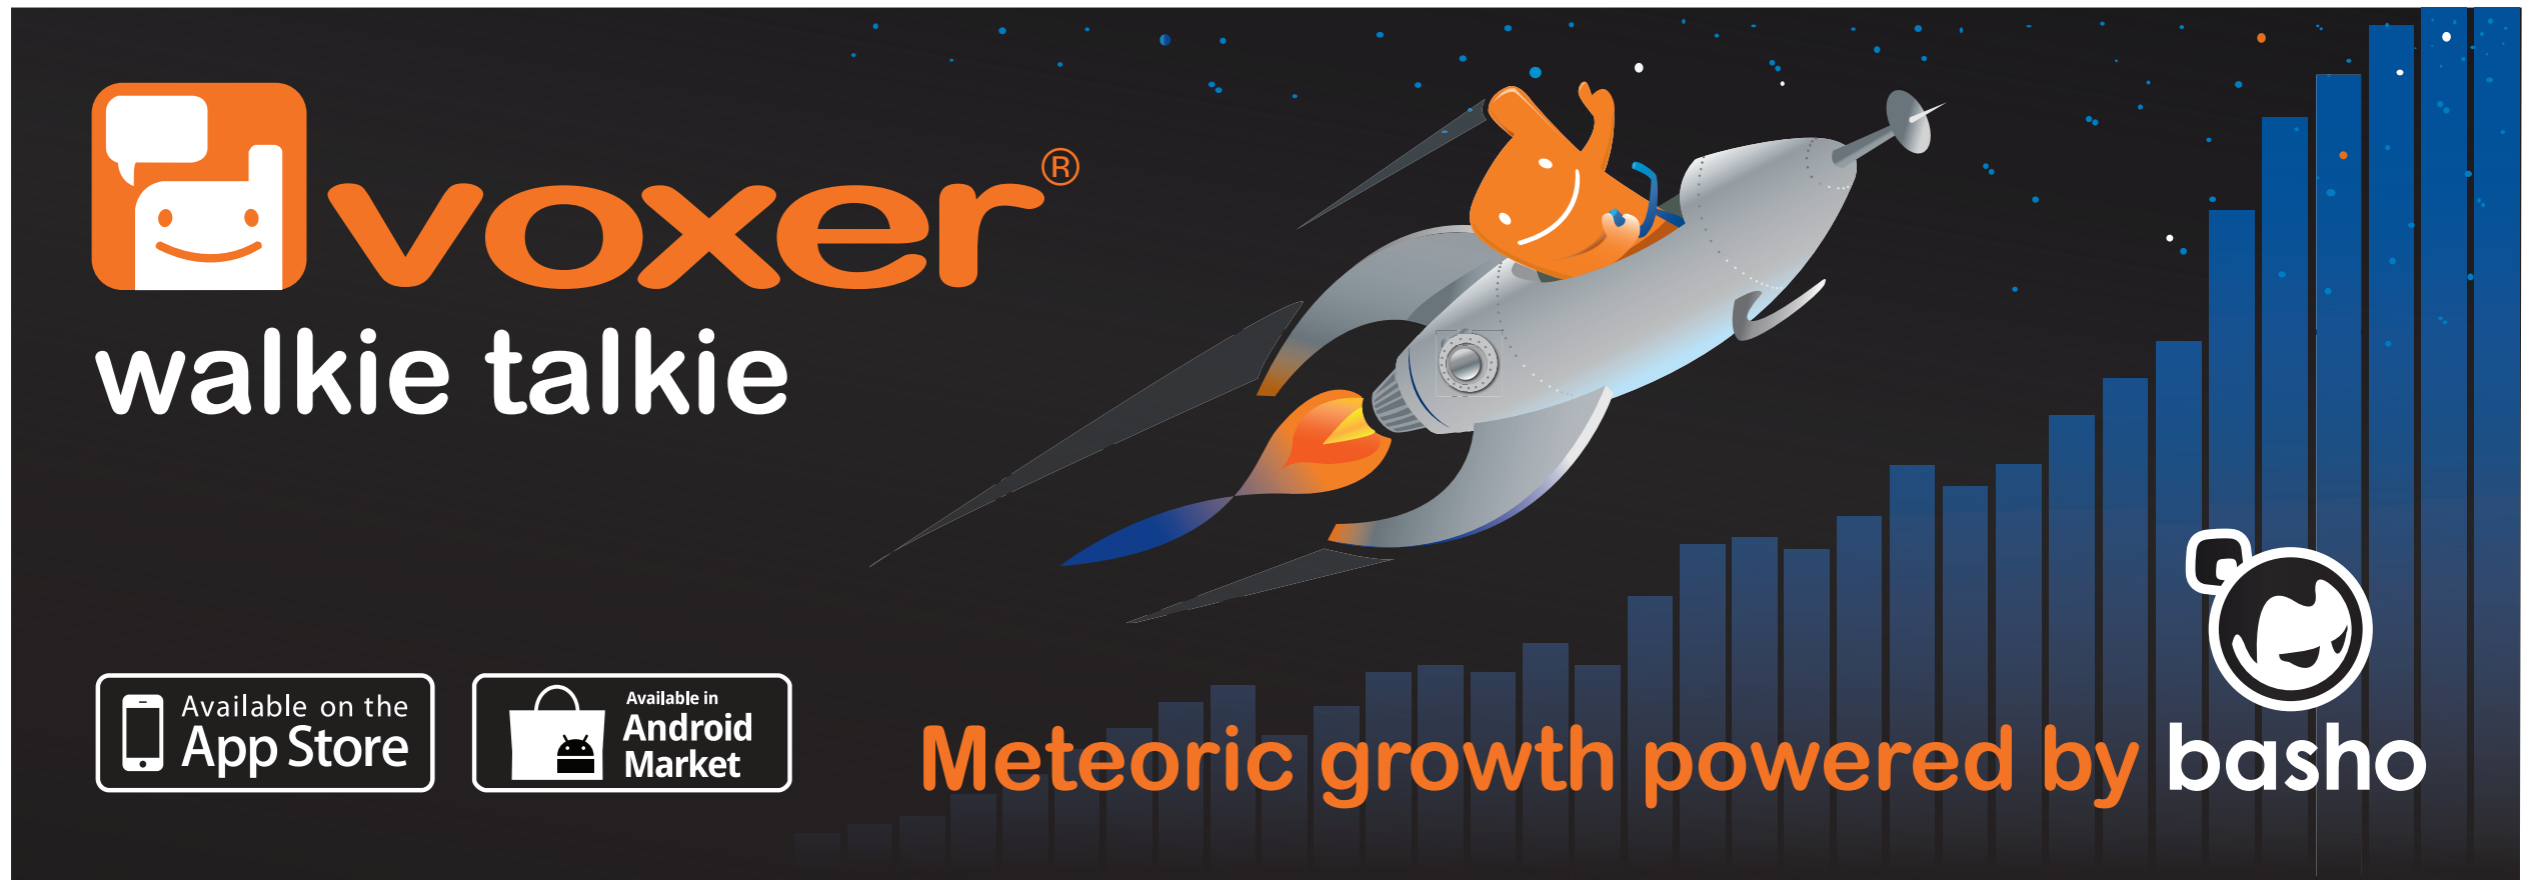


Help from "NoSQL"?

History of Scaling The Web

HTTP	HTTP	HTTP	HTTP	HTTP
App	App	App	App	App
 riak	 riak	 riak	 riak	 riak

History of Scaling



voxer[®]
walkie talkie

Available on the App Store Available in Android Market

Meteoric growth powered by basho

The advertisement features a dark blue background with a rocket ship carrying an orange character, a bar chart showing growth, and the Basho logo.



riak riak riak riak riak

A horizontal row of five Riak logos, each consisting of a stylized network icon followed by the word 'riak'.

Scalable

"I can add twice as much X to get twice as much Y ."

Scalable

computers



"I can add twice as much X to get twice as much Y."

Scalable

computers




"I can add twice as much X to get twice as much Y."

write-throughput!

storage capacity!

map/red power!



Linearly Scalable

computers

"I can add twice as much X to get twice as much Y."

write-throughput!

storage capacity!

map/red power!

Linearly Scalable

computers

"I can add twice as much X to get twice as much Y."

write-throughput!

storage capacity!

map/red power!



Resilient

Assume that **failures will happen.**

At scale, they are **ALWAYS** happening.

Designing whole systems and components with individual failures in mind is a plan for **predictable success.**

Things fail.

Plan it and understand it before your users do.

If you think you can prevent failure, then you aren't developing your ability to respond. - Paul Hammond

Being able to recover quickly from failure is more important than having failures less often. - John Allspaw

Things fail.

Plan it and understand it before your users do.

Applications built for scale can make survival either easier or harder. You get to choose.



Embracing Concurrency at scale

(it's about time!)



Justin Sheehy
justin@basho.com