


**Basho**



**DISTRIBUTED SYSTEMS EXPERTS**



riidk



# Basics

*“We use the Erlang/OTP programming language in building our products here at Basho. We made that choice consciously, believing that it would be a tradeoff – significant benefits balanced by a handful of costs. I am often asked if we would make the same choice all over again. To answer that question I need to address the tradeoff we thought we were making”*

*“The single most compelling reason to choose Erlang was the attribute for which it is best known: extremely high availability”*



# CP

**Consistency**

 **riak**

# AP

**Availability**

 **riak**

**Partition-tolerance**

Riak is the  
**ops-friendly**  
database

# Masterless



a cluster of nodes

# Distributed

each node is identical



# Highly Available





# Multiple Platforms



**symmetrical**

# **Masterless**

**no Single Point of Failure**



# Scalable




add or remove nodes





# Scalable

linear scalability





# Predictable Performance





# Data Model

{“key”: “value”}

Keys are namespaced  
into **“Buckets”**

|     |       |
|-----|-------|
| key | value |

**GET** /buckets/**bucket**/keys/**key**

**PUT** /buckets/**bucket**/keys/**key**



**DEL** /buckets/**bucket**/keys/**key**

**Riak  
Search**



**Map/Reduce**

**2i**




**etc...**

Riak is  
*distributed*


# Consistent Hashing

hash("bucket/key")

# hash ring



# tokenize it





node 0

node 1


node 2

hash(key) →





Replicas are stored to the N  
**Scaling out**  
 1 contiguous partitions





# Vector Clocks

*establish temporality*

# Multi-Datacenter Replication

# Full Sync

# Real-Time





# Riak 2.0

# Strong Consistency

riakcs

# Daily Mail COMMENT

**VASTLY** expensive, utterly unfit for purpose, the NHS computer scheme will go down in history as the most egregious of countless examples of the last Labour Government's incompetence and waste.

It was only our money, after all.

How bitterly telling it is that the £12billion ministers squandered on this vanity

project is exactly the size of the black hole identified this week in the Chancellor's deficit reduction plan.

And how starkly this saga underlines yesterday's rare moment of insight from Nick Clegg, when he told delegates at the Lib Dem conference: 'Never, ever, trust Labour with our economy again.'



**ORACLE®**

“NHS tears out its Oracle Spine in favour of open source”





The background is a solid orange color. It features several abstract elements: three light orange circles of varying sizes and three light orange lines of varying thicknesses that connect the circles in a network-like pattern. The lines and circles are positioned in the upper and left portions of the frame.

Q & A