



## You need Riak

Stuart McCaul, Basho EMEA, London, UK



[stuart@basho.com](mailto:stuart@basho.com)

Stuart\_McCaul

@\_stu\_

[www.xing.com/profile/](http://www.xing.com/profile/)

basho

## Computing has changed


- Commodity hardware
- Concurrent access for billions of devices
- Volume of data in exabytes
- Velocity of data growth
- Variety of unstructured data types

→ Scaling Up is expensive

# Addressing the Data Deluge

A Tough Storage Problem Lies Ahead

**Data Creation Outpaces Growth:  
According to IDC**



JSON  
Video  
Social media  
time series  
SMS  
gaming  
Cloud Storage  
geospatial

email backup  
Audit/Compliance  
sensor data  
bio informatics  
M2M  
GRAPH  
medical records  
search  
Mobile

**Most of these “new” data types are unstructured, semi-structured and non-transactional**

# Why Distributed Data is Critical

In New Era: Latency, Availability and Scalability Matter

## Latency Matters

How one second delay could cost Amazon \$1.6 Billion?

**FASTCOMPANY**

Slowing Google by four tenths of a second costs eight million searches

**Google**

## Availability Matters

Without access to data, automotive production lines shutdown in 15 minutes

**AIAG**

Tech managers often underestimate impact of data loss; 39% say loss would have significant costs

**PCWORLD**

## Scalability Matters

46 million mobile apps downloaded per month from App Store

**KP CB** 2012

By 2020, 20 billion Things will be connected to the Internet

**CISCO**

*More users – human and machine – are **doing more online** from more places with more apps on more devices, **generating more data** than ever before, **requiring more speed, reliability and performance**. And, the trends are only accelerating.*




basho



# Our Era

## The Rise of Distributed Systems



Through 2014, 20% of enterprise data projects will add distributed processes into their production  
– Gartner



# We Bring Amazon Technology Standards to the World

amazon.com

Dynamo: Amazon's Highly Available Key-value Store  
Drew Akers, David Schlegel, Nathan Marz, Dinesh Gopalakrishnan, Michael J. Franklin, Aram T. Pham, Sridhar Ramakrishnan, Peter Vassilakis, and Warren Vogels

**ABSTRACT**  
This paper describes the design and implementation of Dynamo, a highly available key-value store that is used by Amazon.com to power its core services. Dynamo is designed to be highly available, scalable, and fault-tolerant. It is implemented as a distributed system that uses a combination of replication and consistency to ensure that data is available and consistent even in the face of failures. The paper discusses the challenges of designing and implementing a system like Dynamo, and the solutions that were chosen. It also discusses the performance and scalability of Dynamo, and the lessons learned from its development and deployment.

**Categories and Subject Descriptors**  
D.4.7 [Operating Systems]: Storage Management  
H.2.4 [Database Management]: Distributed Databases  
I.2.4 [Database Management]: Distributed Databases

**General Terms**  
Algorithms, Languages, Performance, Reliability, Design

**1. INTRODUCTION**  
This paper describes the design and implementation of Dynamo, a highly available key-value store that is used by Amazon.com to power its core services. Dynamo is designed to be highly available, scalable, and fault-tolerant. It is implemented as a distributed system that uses a combination of replication and consistency to ensure that data is available and consistent even in the face of failures. The paper discusses the challenges of designing and implementing a system like Dynamo, and the solutions that were chosen. It also discusses the performance and scalability of Dynamo, and the lessons learned from its development and deployment.

Amazon Dynamo Specifications Paper  
Released October 2007

Enabled tens of millions of  
simultaneous users on the World's  
most popular e-commerce site

- Amazon has become the defacto standard for modern technology architecture
- In 2007, Amazon published the **Amazon Dynamo** paper
- Dynamo paper outlined requirements not addressable by **any commercial database**
- In 2008, inspired by Dynamo, former Akamai engineers set out to build a **distributed database** and founded Basho Technologies
- After two years of community testing, Basho released **Riak 1.0** in Sept 2011
- In March 2012, Basho launched **Riak CS**, to enter **cloud storage** software market



basho

6

An industry defining milestone was created by pioneer Amazon in 2007.

In 2007, Amazon released the Amazon Dynamo paper. It outlined the business and technical requirements to achieve reliability at massive scale for Amazon.com. The paper presented the design and implementation of Dynamo, a highly available key-value storage system that some of Amazon's core services use to provide an "always-on" experience.

To achieve reliability at massive scale, Amazon was forced to build its own data store, as traditional databases, such as Oracle or IBM DB2, were not built for Internet scale applications.

A group of Akamai engineers and executives with a distributed networking background were inspired by the Dynamo paper. They believed they could radically affect the scale of data stored on the Internet, much as they had accomplished with the Web layer at Akamai. With this goal in mind, they founded Basho in January, 2008.

After nearly three years of development, 500,000+ lines of code, and two years of community testing, Basho released its first product, Riak, in September 2011. Riak is our Dynamo-inspired distributed database.

In 2011, Basho began work on its second product addressing cloud storage and paralleling Amazon's S3 cloud storage service. This compatibility makes migration from Amazon to Riak CS simple. That product, called Basho Riak CS, was selected in late 2011 by AT&T and Verizon/Terremark. It will be officially released in March for broad adoption.

# About Basho

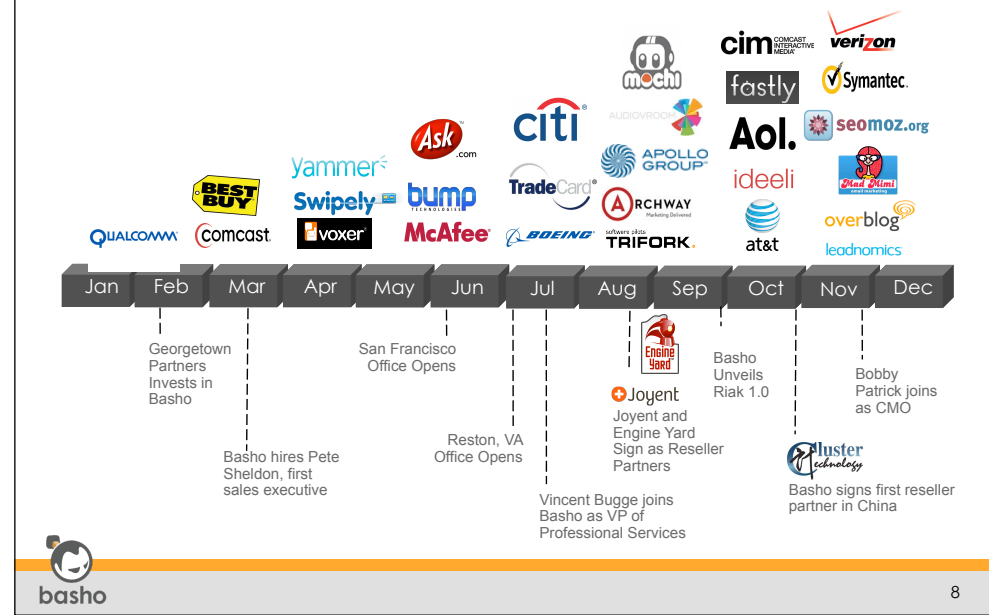
Our Mission is to Be The Leader in Distributed Systems



- Founded January 2008
- 109 employees worldwide
- Basho Makes Riak, used by some of the largest companies in the world and powering the 4th and 5th most popular mobile apps
- Headquartered in USA with regional offices in London and Tokyo
- Board members include Eric Brewer, father of the CAP Theorem
- Strategic partners include Citrix, IDC Frontier, Yahoo! Japan, and Microsoft

# Basho 2011: Launch Year

Basho Picked Up Considerable Speed in 2011

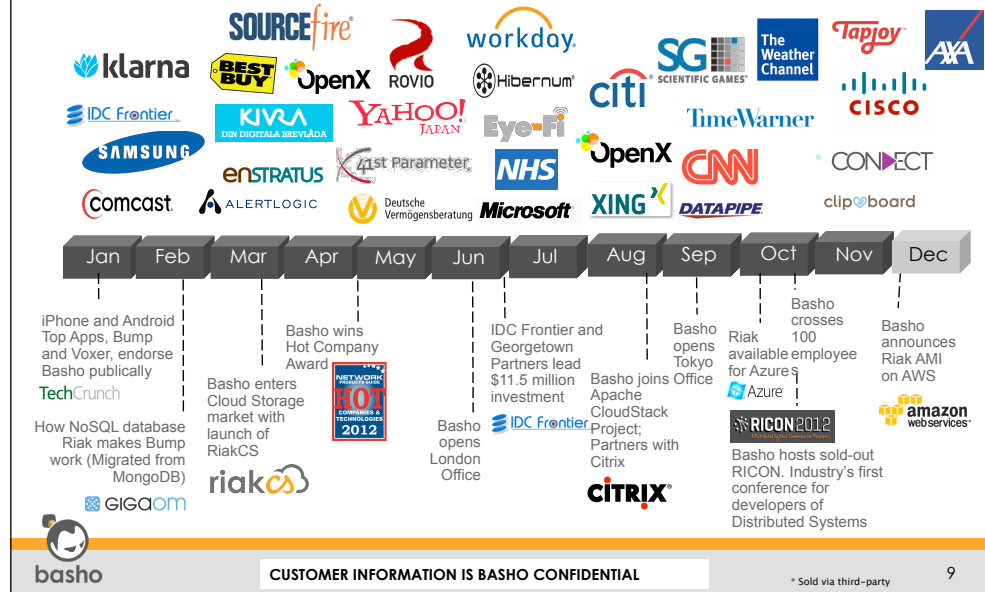


Basho picked up speed in 2011. Fueled by fresh funding from Georgetown Partners in February, the company began expanding its operations:

- Basho quickly established a sales organization, hiring Pete Sheldon, formerly of Verizon and Cybertrust, to start up Basho sales.
- In June 2011, Basho established its first non-shared office facility in San Francisco, CA
- Don Rippert, former CTO of Accenture, joined as president and CEO
- The company began establishing channel partners to expand its sales efforts, recruiting Joyent, Engine Yard and Cluster Technology in China.
- In September, Basho released the first official release of Riak at v1.0 after years of development and community testing.
- During the year, new customer adoption accelerated, culminating with great new customers, such as Symantec and Verizon in Dec.
- The company formalized a marketing organization late in the year with the recruitment of Bobby Patrick, formerly head of marketing at e-commerce leader GXS and hosting leader Digex.

# Basho 2012: Momentum Year

Enterprise-Class Customer Base and Pipeline Growing



CUSTOMER INFORMATION IS BASHO CONFIDENTIAL

\* Sold via third-party

9

## This talk is about

- Distributed Systems
- High Availability
- Non-relational databases
- High Speed / Low Latency



## Distributed Systems

- Clustering
  - Make the most of multiple machines
- Trade offs:
  - Availability vs Consistency
- Start with 5 nodes
  - 3 copies of your data
  - 1 node in maintenance
  - 1 node can fail

## High Availability

- Any sufficiently large system is always in a partial state of failure
- Riak 100% availability\* for years
  - Across Riak versions
  - Across generations of hardware
- Masterless
  - No SPOF

\* What does availability mean?  
Always accepts writes  
No outages for upgrades  
Can survive DC/AZ outages





## High Availability

**enStratus:** Cloud infrastructure management and monitoring company

- Replaced Oracle MySQL with Riak in 2012
- Eliminated the one big remaining single-point of failure in the enStratus environment
- Moved from relational to non-relational system required a new data model approach
- Basho.com/blog

enSTRATUS



## Non-Relational

- No Schema
- No Schema Migrations
- Supports Rapid Application Development
  
- No business logic in the database
- No locks
- No joins
- No referential integrity
  
- Riak is better for concurrent access



## Non-Relational

### Alert Logic

Fast-growing leading Managed Security Services Provider

- Replaced MySQL with Riak in 2012
- 35k ops/sec/server
- 5TB a day
- Data doubling every year
- 100,000s of edge endpoints
- Massive performance over MySQL
- Basho.com/blog



## High Speed / Low Latency

- Consistent speed
- Even distribution of load
- No need for in-memory cache

## High Speed / Low Latency

**Best Buy:** Leading electronics retailer selected Riak for bestbuy.com, including its online product catalog

- Switched from relational Oracle database to Riak for bestbuy.com performance
- Uses multi-data centre replication to ensure 100% availability
- Added 70 Amazon Web Services nodes for holiday expansion (peak load)
- Basho.com/blog



## DevOps Friendly

- Scaling +/- nodes is fast and simple
- Backup is easy and zero-impact
- Monitoring is easy
- Flexible topology
- Riak Control GUI
- Puppet/Chef/Fabric

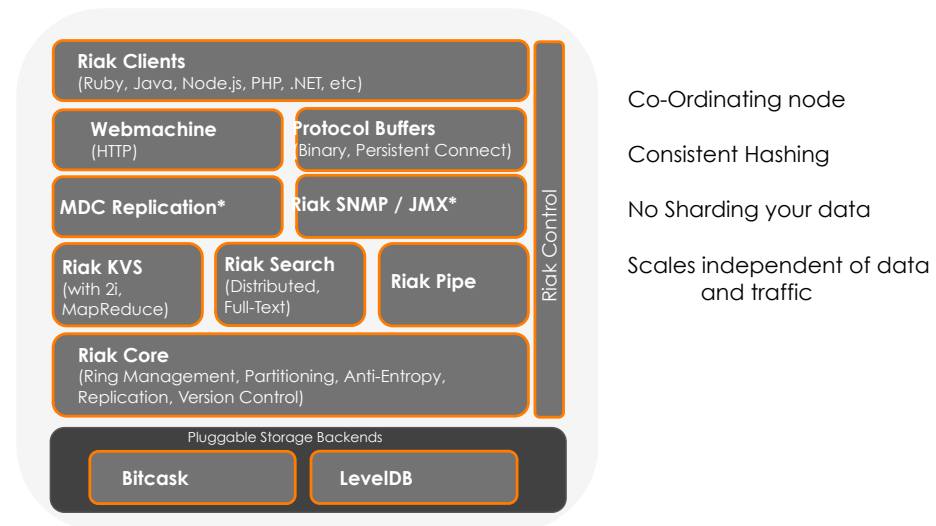


## DevOps Friendly

- Mad Mimi:** Large email marketing service provider with over 100,000 customers
- Replaced MySQL to scale quickly and easily without manual sharding
  - Riak currently stores nearly 5 billion keys, and adds between 10-20 million keys each day
  - Gained 100% availability since moving to Riak
  - Uses Riak Secondary Indexing to make retrieving data easier and faster than with MySQL
  - Basho.com/blog




# Riak Architecture





# Basho Family of Products

Basho's Distributed Systems Technologies



Open Source  
Distributed Database

- Always-available, scalable, low-cost NoSQL database
- Over 35,000 Downloads per Month
- Thousands of users worldwide
- Version 1.0 unveiled September 2011



Commercial  
Distributed Database

- Adds multi-data center replication & Basho 24x7 support
- Enables Secondary Search and MapReduce Clusters
- Requires commercial contract and secure download
- Version 1.3 coming January 2013



Distributed Cloud  
Storage Platform

- Expands with multi-tenancy, large object support, metering and Amazon S3 API
- Enables public and private cloud storage
- Developer trial license available



basho

McAfee BOEING



Comcast

IDC Frontier at&t

# Yahoo! JAPAN Cloud Storage

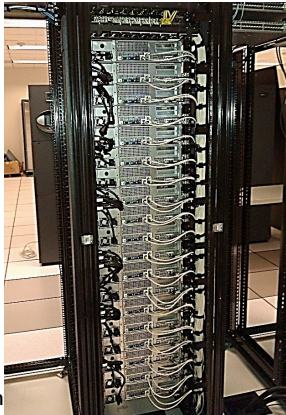
Powered by Riak CS




**Yahoo! Japan  
Launches  
Amazon S3  
Competitor for  
Japan market  
– January 2013**


# Riak CS 'in the Cloud'


## AT&T's Largest Cloud Platform and Growing







+T BSS/OSS  
& Service  
Orchestration










+T Network Orchestrator



at&t

format

SilverLining is AT&T's

conducing

Today

Third Outline Level

Fourth Outline Level

Fifth Outline Level

Sixth Outline Level

Seventh Outline Level

Eighth Outline Level

Ninth Outline Level

Click to edit

Master



Third level

Fourth level

Third level

Fourth level

# mHealth Platform-as-a-Service



- AT&T unveiled mHealth at DevSummit in January 2012
- Building a Platform-as-a-Service ecosystem connecting developers, apps and devices
- 3rd Party Devices include Withings, FitBit and MyZeo
- Key values stored in Riak include Users, Tokens, Applications, and Health Metrics
- HIPAA and HL7 compliant

*Clinical care teams, insurance companies, pharmacists, employers will benefit from a more holistic view of someone's current state of health.*




# Cloud-Based Meter Management

Basho in Utilities and Smart Metering



- Internet delivered MDM with revenue-critical data
  - Final Repository for all Meter Readings
- Mobile Devices, Back Office Users
- "Reliability and Availability trumps all"
  - SQL to Riak Migrations
    - Split binary and meta
- "Ops is Easy / Nodes down Unnoticed"

Hosted Meters in Temetra



# Mobile Walkie-Talkie App

Voxer – Serious Scale with Text, Audio and Video




Meteoric growth powered by basho

- Switched to Riak in April 2011
- #4 most popular Apple App Store *behind Facebook, Skype and Twitter*
- Scaled 10x between Nov 2011 and Jan 1, 2012
- Over 100 nodes today supporting 2 billion operations / day
- Growing by 5 TB every day



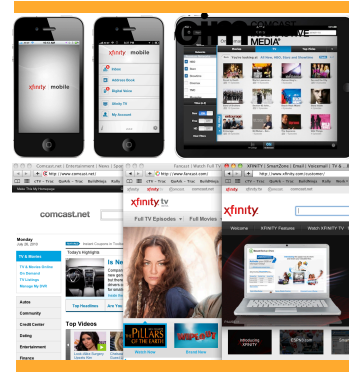
# Mobile-to-Mobile Content Store



- Switched to Riak in August 2011
- #7 Most Downloaded iPhone App
- 800 million pieces of structural data in Riak, including Photos, Chats, and Contact Cards.
- 10 million active users
- 77 million downloads to date

# Video Delivery Platform

Comcast – Riak Expands from One Project to the Next



- Originally used Amazon S3; challenged by latency
- Deployed Riak initially to store user related data for DVRs
- Expanded Riak to manage streaming premium content and remote DVR controls for xfinitytv.comcast.net, as well as XfinityTV iPhone and Android
- In 1Q12, Riak launched to support Comcast StreamPix. Comcast's new product against Netflix
- Expanded to multi-data center in Q112



28

Comcast on NoSQL: "A KV store is flat out simpler to use and manage than comparable, more general RDBMS technologies, i

Comcast on Operations with Riak: "On a day-to-day basis, other than figuring out how much capacity we want to add, we have no Riak administrator— we don't need one."

Jonathan Moore  
Technical Fellow Comcast Interactive Media





# Product Information Repository

## Best Buy – Multi-Center Replication for Availability



- Best Buy selected Riak to be the on-line product catalog for bestbuy.com and retail stores
- 2 Million SKUs stored in Riak
- Accessed by 30 internal applications, such as customer service enterprise returns app
- Integrates with 3rd Party APIs to Amazon.com and Pricegrabber.com
- Increasing geographic diversity implementing Riak's Multi-Data Center replication in 1Q12

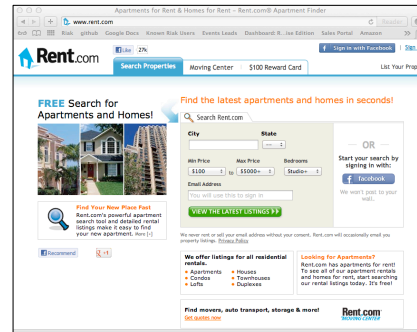


Mirror of Riak is available at [code.bestbuy.com](http://code.bestbuy.com).

# Riak Powering The Interest Graph

## PRIMEDIA

The Targeted Media Specialists



Primedia is a property search web site and search engine.

Primedia uses **Riak** as the database to house, **personalized consumer searches**. Riak provides Primedia with a **fast read/write** capability to fully capture search criteria and enable a per consumer interest graph.

# Summary

Riak is:

- Distributed
- Highly Available
- Schemaless
- Fast and Scalable
- Developer Friendly
- Ops Friendly




# What We are Working On

Riak v1.3 release this week:

- Active Anti-Entropy: automatic, self-healing properties that repair entropy on an ongoing basis.
- Improved Riak Enterprise's multi-datacenter replication performance: more TCP connections and easier configuration.
- Improved graphical user experience.
- Expanded IPv6 support.
- Improved MapReduce.
- Simplified log management.



# The 1<sup>st</sup> Distributed Systems Conference



# Do you need Riak?

- Request a Tech Talk for your team
  - [bit.ly/RiakTechTalk](http://bit.ly/RiakTechTalk)





---

**Questions?**

**Feedback to @\_stu\_**

Presentation Includes Basho Confidential Information