

Key Value Stores

14 November 2017

Ignacio Coterillo

Computer Engineer, CERN

Key-Value Stores (1)

Aggregates are basically a *hash* table

Support for CRUD operations (Create, Read, Update, Delete)

Values are accessed by its key

Examples: Riak, Redis, Memcached, *Dynamo*

Key-Value Stores (2)

Easy scalability


High performance

A pure KV model is pretty limited and doesn't support complex data relations

Most KV oriented stores extend the basic concept blurring the separation line with document stores

Memcached Example

MySQL/Memcached (<https://dev.mysql.com/doc/refman/5.6/en/ha-memcached.html>)



Key-Value Stores (3)

Redis: KV where the Value can be composed:

- Single Value
- Lists
- Sets
- Hashes

Supports complex operations like *range*, *diff*, *union* and *intersection*

Key-Value Stores Limitations (1)

Can't express relations (*Riak*)

No transactions for multiple aggregates (*Redis*)

No bulk operations

Not possible to construct ad-hoc queries based on values, so no possible exploration server side. It needs to be done on the client side

```
Client ----- Q1 -----> DB
Client <----- A1 ----- DB
Client ----- Q2(A1) -----> DB
Client <----- A2 ----- DB
```

Key-Value Typical Applications

- Application caching
- Session caching
- Shopping cart applications

Riak (1)

Inspired by Amazon Dynamo

Dynamo: amazon's highly available key-value store (<https://dl.acm.org/citation.cfm?id=1294281>)

WEB Oriented

Developed in Erlang

HTTP interaction

Riak Basic Concepts (2)

RDBMS	-----	Riak
Database	-----	Cluster
Schema	-----	Bucket
Row	-----	Row
Key	-----	RowID

Riak (4)

You can get a list of the keys in a bucket

- As a batch
- As a stream -> recommended for large buckets

Q. How does this affect the system?

Riak. Links (5)

- One directional
- Metadata for the values
- Can be used in filtered searches (linkwalking)
- Also support for metadata types

Riak. Additional functions (7)

- Stored functions (UDF), written in Javascript
- Map/Reduce (not for real time analysis)
- Pre/Post commit hooks
- Extensions (search)
- Metadata Indexing

Redis (1)

Defined as a *Data Structure* server, but with several functions:

- Data Structure Server
- Queue (Blocking Lists and Pipes)
- Pub/Sub server

Redis Features (2)

- Configurable expiration policies
- Durability levels
- Multiple replication options
- Multivalue operations
- Certain degree of interpretation of the stored values (Integers)

Redis Features (3)

- Transaction support (MULTI, DISCARD)
- Data Types

- Strings
- Lists
- Sets
- Sorted sets
- Hashes (Can't be nested)

- Blocking Lists:

- Like OS pipes
- Consumer/Produce

- Publish/Subscribe

Redis. Durability (4)

- Works on memory by default. Dataset limited by memory size
- Snapshot support

- The whole database is flushed to disk.
- Works by forking the database server process
- The service is interrupted during the operation

- Append Only File

- Used as basis for replication

Redis. High Availability

Single node setup limited by server specifications. The maximum data size is the total memory of the server

- Master-Slave Replication. For Availability, but not size scaling
- Redis Cluster. Support for partitions and partition replication

Optional Consistency

Trade-off between consistency and speed

In a setup with N nodes:

- Write operations to be confirmed by W nodes
- Read operations to be confirmed by R nodes.

W and R can be configured at the system level, per bucket/namespace or per operation

Scaling

Simple partitions by key (either random or uniform) for sharding

Other KV stores

Embeddable

- LevelDB
- RocksDB

Cloud

- Amazon Dynamo

Use case examples: CERN

Riak used as back-end for the Teigi *ecosystem*:

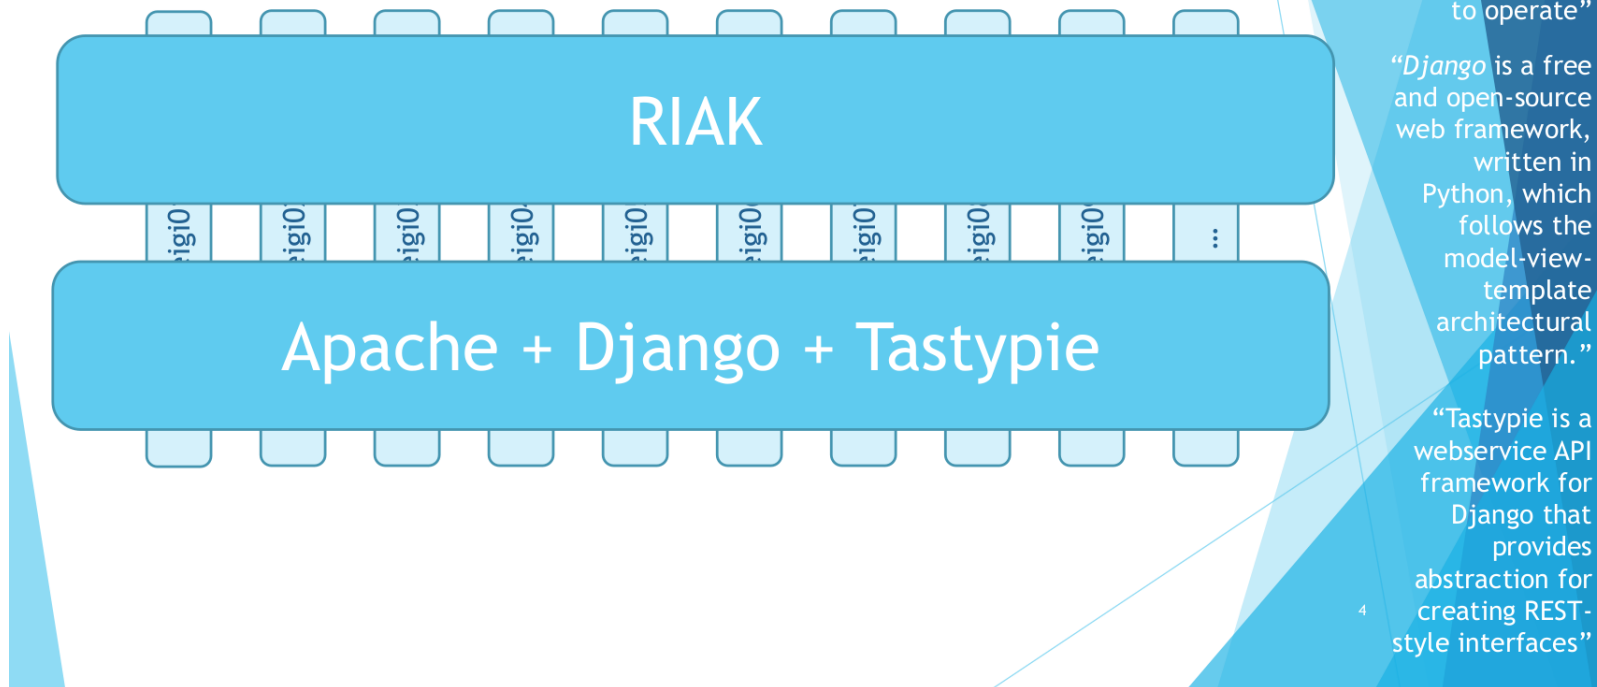
- Tbag
- Roger
- Pwn
- Tellme

9 Node Riak Cluster as back-end

Use case examples: CERN

Credits: CERN, Zhechka Toteva


Current architecture



Use case examples: CERN

Credits: CERN, Zhechka Toteva

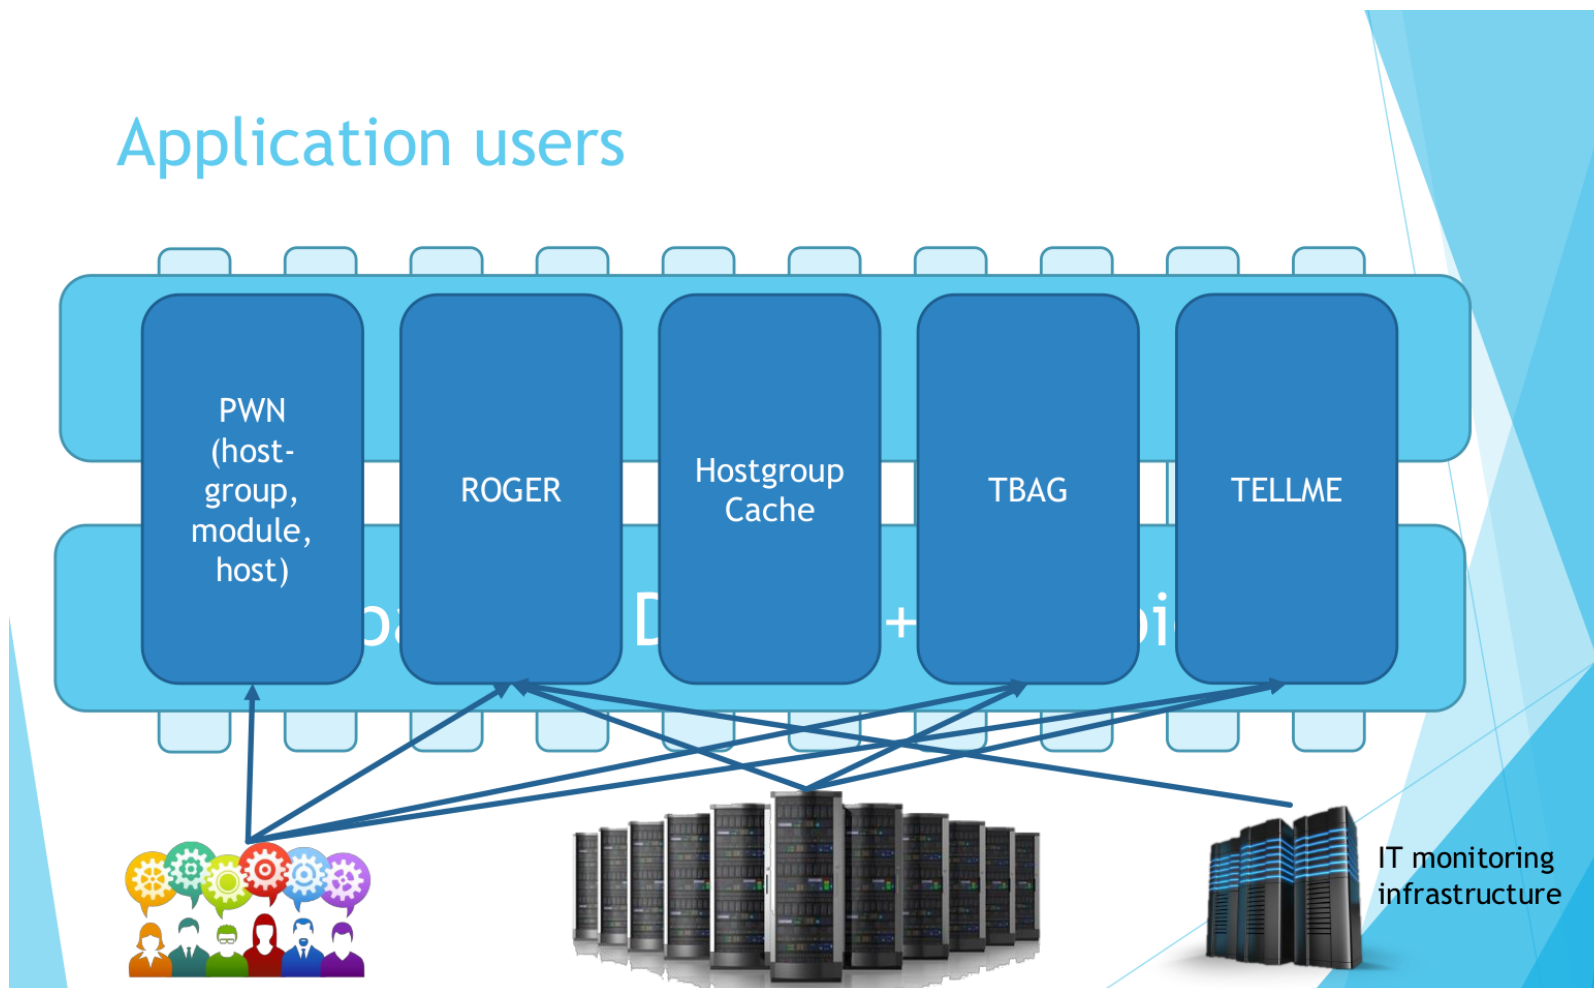
Applications



Use case examples: CERN

Credits: CERN, Zhechka Toteva

Application users



Bet365



ACADEMY

DOWNLOAD RIAK

DOCS

CONTACT

SEARCH



PRODUCTS

INTEGRATIONS

SERVICES

INDUSTRIES

USE CASES

RESOURCES

COMPANY

BLOG

Can't afford to gamble on your database infrastructure?

WHY DID BET365 CHOOSE RIAK? TURN UP THE VOLUME AND FIND OUT NOW!

Can't afford to Gamble on your database infrastructure Why bet365 chose Riak



Can't afford to Gamble on your
database infrastructure?

Why bet365 chose Riak.

0:00 / 1:01:00

CC BY-NC-SA YouTube

Bet365

FW: Riak and the demise of Basho

Martin.Cox at bet365.com [Martin.Cox at bet365.com](#)

Thu Aug 24 11:17:43 EDT 2017

- Previous message (by thread): [Call to action](#)
- Next message (by thread): [Riak and the demise of Basho](#)
- **Messages sorted by:** [\[date \]](#) [\[thread \]](#) [\[subject \]](#) [\[author \]](#)

Hi

I have been asked to forward the below message from Martin Davies, the CEO of Technology for bet365.

Kind Regards

Martin Cox

From: Martin Davies
Sent: 24 August 2017 16:11
To: Martin Cox
Subject: Riak and the demise of Basho

Hi,


I have been wanting to make you aware for a few weeks now that we have reached an agreement, in principle, to buy all of Basho's remaining assets (except support contracts) from the receiver. Up until this afternoon, I was constrained by some confidentiality needs of the receiver and was unable to speak.


We have agreed a price for the assets and are almost at the end of sorting out the legal agreement. Once this is complete, this will then need to be processed through the courts which, I am advised, should take a week or so.



It is our intention to open source all of Basho's products and all of the source code that they have been working on. We'll do this as quickly as we are able to organise it, and we would appreciate some input from the community on how you would like this done.

Martin Davies
Chief Executive Officer - Technology
Hillside (Technology) Limited
bet365.com


Bet365

 Search Hacker News

 basho

by  algolia 

Search Stories ▼ by Date ▼ for All time ▼

278 results (0.004 seconds) 

Riak and the demise of **Basho**
207 points | bryanasmussen | a month ago | 96 comments | (http://lists.basho.com/pipermail/riak-users_lists.basho.com/2017-August/019500.html)

Riak acquired by bet356
3 points | tschern | a month ago | 0 comments | (http://lists.basho.com/pipermail/riak-users_lists.basho.com/2017-August/019500.html)

Bet365 to buy all of **Basho's** remaining assets
2 points | baq | 2 months ago | 0 comments | (http://lists.basho.com/pipermail/riak-users_lists.basho.com/2017-August/019500.html)

Riak lives. bet365 buys **Basho's** assets
6 points | auxbuss | 2 months ago | 0 comments | (<https://twitter.com/ErlangSolutions/status/900754876971585536>)

Court puts **Basho** into receivership
1 points | thehodge | 3 months ago | 0 comments | (https://www.theregister.co.uk/2017/07/31/end_of_the_road_for_basho_as_court_puts_biz_into_receivership/?mt=1501618708071)

Show HN: Find serendipitous haiku on Twitter
7 points | bevan | 3 months ago | 0 comments | (<https://basho.me>)

Is **Basho** finally dead?
3 points | mprev | 3 months ago | 0 comments | (https://www.theregister.co.uk/2017/07/13/will_the_last_person_at_basho_get_the_lights_oh_too_late/)

Will the last person at **Basho** please turn out the lights?
11 points | sushimako | 3 months ago | 2 comments | (https://www.theregister.co.uk/2017/07/13/will_the_last_person_at_basho_get_the_lights_oh_too_late/)

Is **Basho** dead?
67 points | okket | 4 months ago | 9 comments | (https://twitter.com/_jared/status/878036763805048832)

Basho CTO Dave McCrory: The Future of the Network
2 points | majke | 4 months ago | 0 comments | (<https://blog.mccrory.me/2017/05/22/the-future-of-the-network/>)

Thank you

Ignacio Coterillo

Computer Engineer, CERN

ignacio.coterillo.coz@cern.ch (mailto:ignacio.coterillo.coz@cern.ch)

