My Path to NoSQL CouchDB, Redis, Elasticsearch

Karel Minařík

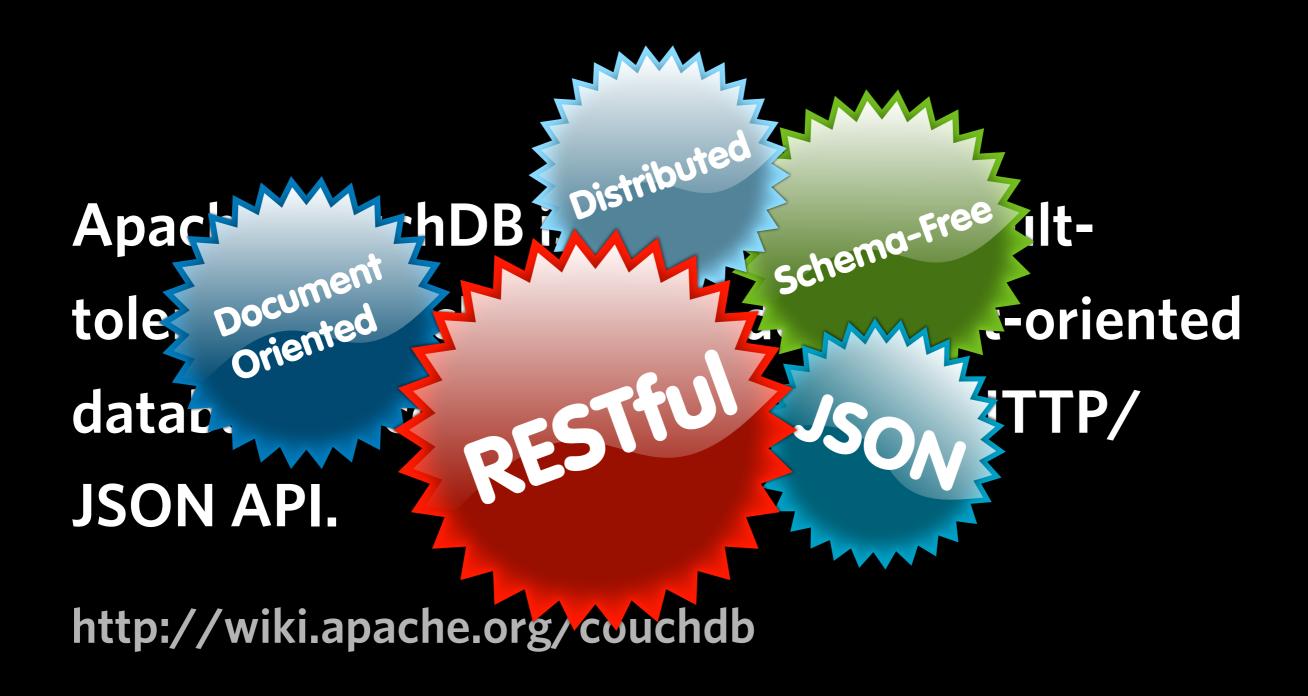
- → Software developer at Elasticsearch.com
- → Previously: Adobe Flash Developer, Art Director, Information Architect, Rails programmer, ...
- → Graduated as Philosophy major
- → @karmiq at Twitter

→ **karmi.cz**



Apache CouchDB is a distributed, faulttolerant and schema-free document-oriented database accessible via a RESTful HTTP/ JSON API.

http://wiki.apache.org/couchdb



R.I.P, CouchDB



DAMIEN KATZ JUST RELAX. NOTHING IS UNDER CONTROL.



The Future of CouchDB

What's the future of CouchDB? It's Couchbase.

Huh? So what about Apache CouchDB? Well, that's a great project. I founded it, coded the earliest versions almost completely myself, I've spent a huge amount of blood, sweat and tears on it. I'm very proud of it and the impact it's had. And now I, and the Couchbase team, are mostly moving on. It's not that we think CouchDB isn't awesome. It's that we are creating the successor to it: Couchbase Server. A product and project with similar capabilities and goals, but more faster, more scalable, more customer and developer focused. And definitely not part of Apache.

With Apache CouchDB, much of the focus has been around creating a consensus based, developer community that helps govern and move the project forward. Apache has done, and is doing a good job of that. But for us, it's no longer enough. CouchDB was something I created because I thought an easy to use, peer based, replicating document store was something the world would find useful. And it proved a lot of the ideas were possible and useful and it's been successful beyond my wildest ambitions. But if I had it all to do again, I'd do many things different.

If it sounds like I'm saying Apache was a mistake, I'm not. Apache was a big part in the success of CouchDB, without it CouchDB would not have enjoyed the early success it did. But in my opinion it's reached a point where the consensus based approach has limited the competitiveness of the project. It's not personal, it's business.

And now, as it turns out, I have a chance to do it all again, without the pain of starting from scratch. Building on the previous Apache CouchDB and Membase projects, throwing out what didn't work, and strengthening what

About Me me@damienkatz.com

Follow @damienkatz

WHERE AM 1?

You are at the blog of Damien Katz, who once stole the goatee from a billy goat. He kept it in a Skoal can on his nightstand for years. Then one day, it was gone. Vanished.

Or maybe it was lost during a move.





Database denormalization at Digg

```
SELECT `digdate`, `id` FROM `Diggs`
WHERE `userid` IN (1, 2, 3, 4, ... 1000000)
AND itemid = 123 ORDER BY `digdate` DESC, `id` DESC;
```

"A full query can actually clock in at 1.5kb, which is many times larger than the actual data we want. With a cold cache, this query can take 14 seconds to execute."

"Non-relational data stores reverse this model completely, because they don't have the complex read operations of SQL. The model forces you to shift your computation to the writes, while reducing most reads to simple operations – the equivalent of SELECT * FROM `Table`."

http://about.digg.com/blog/looking-future-cassandra

"Relational Data"



Home Profile Find People Settings Help Sign out

OH: "The world is relational!!!"

17 minutes ago via Tweetie for Mac Retweeted by 10000 people

♠ Reply 13 Retweet





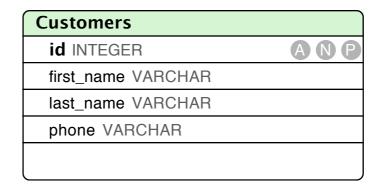
Design a customer database. People have names, e-mail, phone numbers, ...



Many phone numbers?

Relational Databases 101

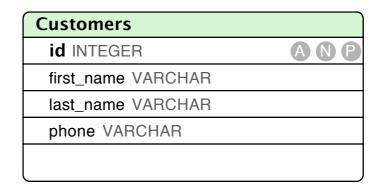
Customer ID	First Name	Surname	Telephone Number
123	Robert	Ingram	555-861-2025
456	Jane	Wright	555-403-1659
789	Maria	Fernandez	555-808-9633



Now. What about multiple phone numbers?

The "solution", Pt. 1

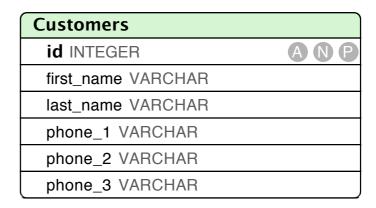
Customer ID	First Name	Surname	Telephone Numbers
123	Robert	Ingram	555-861-2025
456	Jane	Wright	555-403-1659, 555-776-4100
789	Maria	Fernandez	555-808-9633



"We will use the database only from the application, anyway."

The "solution", Pt. 2

Customer ID	First Name	Surname	Tel. No. 1	Tel. No. 2	Tel. No. 3
123	Robert	Ingram	555-861-2025		
456	Jane	Wright	555-403-1659	555-776-4100	555-403-1659
789	Maria	Fernandez	555-808-9633		



"This is clearly better design!"

Alright. Then, please answer these questions:

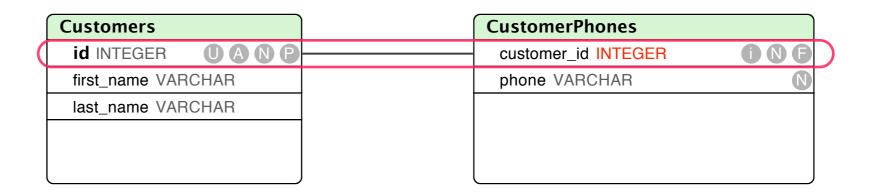
- How do you search for a customers given a phone number?
- Which customers have the same phone number?
- How many phone numbers a customer has?

Then, please add the ability to store four phone numbers. Thanks.

The Right Solution

Customer ID	First Name	Surname
123	Robert	Ingram
456	Jane	Wright
789	Maria	Fernandez

Customer ID	Telephone Number
123	555-861-2025
456	555-403-1659
456	555-776-4100
789	555-808-9633



```
mysql> SELECT * FROM Customers WHERE id = 1;
+---+
| id | first_name | last_name |
  1 John Smith
+---+
mysql> SELECT phone FROM CustomerPhones WHERE customer_id IN (1);
+----+
phone
 123
456
+----+
```

Structured data

But, damn!, I want something like this:

```
{
  "id" : 1,

  "first_name" : "Clara",
  "last_name" : "Rice",

  "phones" : ["0000 777 888 999", "0000 123 456 789", "0000 314 181 116"]
}
```

"No problem, you just iterate over the rows and build your object. That's the way it is!"

"If this would be too painful, we will put some cache there."

Ephemeral data

Not everything needs to be done "right". Right?

```
class User < ActiveRecord::Base
  serialize :preferences
end</pre>
```

"Consistency"

Does the "Right Way" sometimes fail? Hell yeah.

EXAMPLE

When designing an invoicing application, you store the customer for the invoice the "right way", via foreign keys.

Then, the customer address changes.

Did the address on the invoice also changed?

Documents in the Real World



Documents in the Real World

```
: "clara-rice",
"_id"
"_rev" : "1-def456",
"first_name" : "Clara",
"last_name" : "Rice",
"phones" : {
  "mobile" : "0000 777 888 999"
 "home": "0000 123 456 789",
 "work" : "0000 314 181 116"
},
"addresses" : {
 "home" : {
   "street" : "Wintheiser Ports",
   "number" : "789/23",
   "city" : "Erinshire",
   "country" : "United Kingdom"
 },
},
"occupation" : "model",
"birthday" : "1970/05/01",
"groups" : ["friends", "models"],
"created at" : "2010/01/01 10:00:00 +0000"
```



Built "Of the Web"

Django may be built for the Web, but CouchDB is built of the Web.

I've never seen software that so completely embraces the philosophies behind HTTP.

Jacob Kaplan-Moss, Of the Web (2007)

HTTP API

```
HOST=http://localhost:5984
curl -X GET $HOST
# {"couchdb":"Welcome", "version": "0.11.0b22c551bb-git"}
curl -X GET $HOST/my-database
# {"error":"not_found","reason":"no_db_file"}
curl -X PUT $HOST/my-database
# {"ok":true}
curl -X PUT $HOST/my-database/abc123 -d '{"foo":"bar"}'
# {"ok":true, "id": "abc123", "rev": "1-4c6114c65e295552ab1019e2b046b10e"}
curl -X GET $HOST/my-database/abc123
# {" id": "abc123", " rev": "1-4c6114c65e295552ab1019e2b046b10e", "foo": "bar"}
curl -X DELETE $HOST/my-database/abc123?rev=2-d179f665eb01834faf192153dc72dcb3
# {"ok":true, "id": "abc123", "rev": "1-4c6114c65e295552ab1019e2b046b10e"}
```

Easy To Wrap

```
require 'rubygems'
require 'ostruct'
require('restclient')
                                         HTTP library
                                         JSON library
require('json'
class Article < OpenStruct</pre>
  def self.db(path='')
    RestClient::Resource.new "http://localhost:5984/blog/#{path}",
                              :headers => { :content type => :json, :accept => :json }
  end
  db.put '' rescue RestClient::PreconditionFailed
  def self.create(params={})
    new db.post(params.to json)
  end
  def self.find(id)
    new JSON.parse( db(id).get )
  end
  def destroy
    self.class.db(self. id + "?rev=#{self. rev}").delete
  end
end
```

What is "RESTful"?

REST is a set of principles that define how Web standards, such as HTTP and URIs, are supposed to be used. (...) In summary, the five key principles are:

- Give every "thing" an ID
- Link things together
- Use standard methods
- Resources with multiple representations
- Communicate statelessly

Stefan Tilkov, A Brief Introduction to REST

What is "RESTful"?

The basic idea is even more simple, though.

HTTP is not just a "transfer protocol".

It is the interface for interacting with "things" themselves.

Playing Tricks with HTTP



playing http tricks with nginx

Karel Minařík
October 7, 2014

One of the defining features of Elasticsearch is that it's exposed as a (loosely) RESTful service over HTTP.

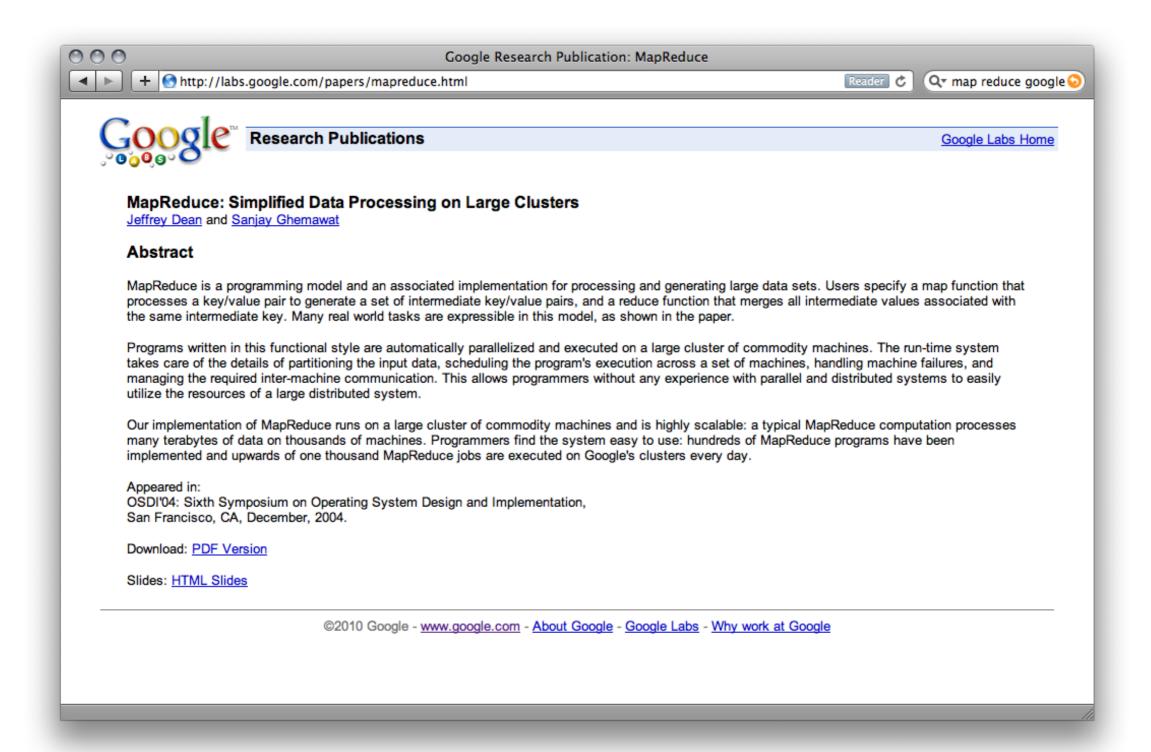
The benefits are easy to spell out, of course: the API is familiar and predictable to all web developers. It's easy to use with "bare hands" via the curl command, or in the browser. It's easy to write API wrappers in various programming languages.

Nevertheless, the importance of the HTTP-based nature of Elasticsearch is rooted deeper: in the way it fits into the existing paradigm of software development and architecture.

http as an architectural paradigm

Display a menu

The Google Paper



The Concept

```
module Enumerable
  alias :reduce :inject unless method_defined? :reduce
end

(1..3).map { | number | number * 2 }

# => [2, 4, 6]

(1..3).reduce(0) { | sum, number | sum += number}

# => 6
```

The Simplest View

```
function(doc) {
   if (doc.last_name && doc.first_name) {
      emit( doc.last_name + ' ' + doc.first_name, doc )
   }
}
```

The Simplest View

```
function(doc) {
  if (doc.last_name && doc.first_name) {
    emit( doc.last_name + ' ' + doc.first_name) doc) }
}
OUTPUT

KEY

VALUE
```

The Result of Map

Key	Value
"Armstrong Lottie"	<pre>_id: "lottie-armstrong", _rev: "2-fcb71b26096957b3ff3ffd2970f3c933", addresses: { home: { city: "Murphyville" } }, first_name: "Lottie", last_name: "Armstrong", occupation: "programmer",</pre>
"Bailey Kaelyn"	<pre>_id: "kaelyn-bailey", _rev: "1-2e25e6c9448520fa796988894423a23b", addresses: { home: { city: "Lake Dedric" } }, first_name: "Kaelyn", last_name: "Bailey", occupation: "supermodel"</pre>
•••	•••

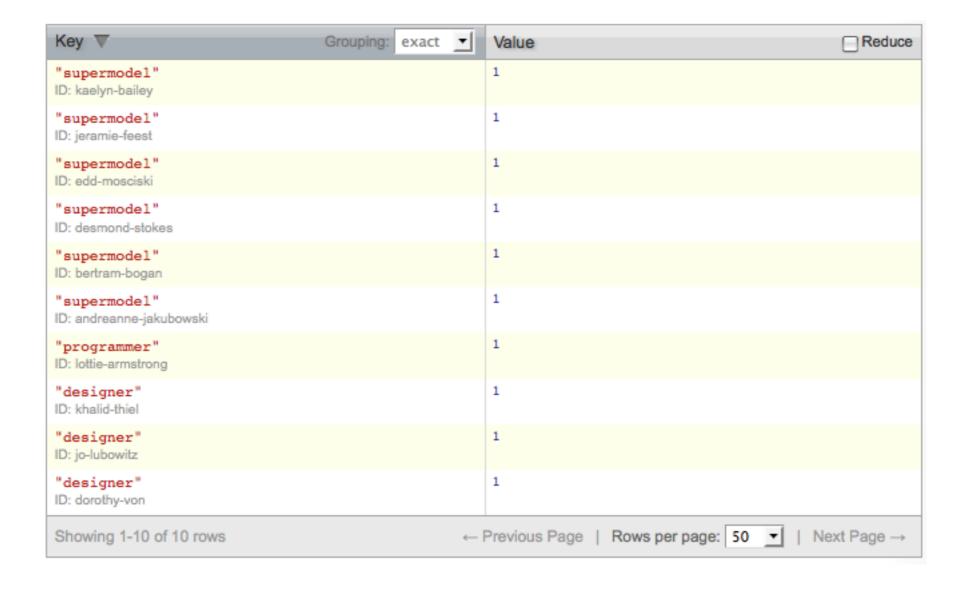
The Result of Map

```
000
                                            http://localhost:5984/addressbook/_design/person/_view/by_name?reduce=false
                                                                                                                           Google Google
                                                                                                                                                         Q)
                           http://localhost:5984/addressbook/_design/person/_view/by_name?reduce=false
{delicious} [rack::bug] localhost:3000 localhost:4567 futon [doc] [GData API]
🔼 Apache CouchDB – Futon: Brow... 🖸 🔼 http://localhost:5984/address... 😧 🛨
   total_rows: 10,
   offset: 0,
  - rows: [
     - {
           id: "lottie-armstrong",
           key: "Armstrong Lottie",
         - value: {
               _id: "lottie-armstrong",
               rev: "2-fcb71b26096957b3ff3ffd2970f3c933",
             + addresses: { ... },
               first name: "Lottie",
               last_name: "Armstrong",
               occupation: "programmer",
             + phones: { ... },
             + groups: [ ... ],
              birthday: "1952/10/24",
             + _attachments: { ... }
       },
           id: "kaelyn-bailey",
           key: "Bailey Kaelyn",
         - value: {
               _id: "kaelyn-bailey",
               rev: "1-2e25e6c9448520fa796988894423a23b",
             + addresses: { ... },
               first_name: "Kaelyn",
               last_name: "Bailey",
               occupation: "supermodel",
                  cell: "1-758-583-3506 x57386",
                  home: "(738)356-9599 x5321"
               },
             + groups: [ ... ],
               birthday: "1970/07/09"
Done
```

Even Simpler View

```
function(doc) {
  emit(doc.occupation, 1);
}
```

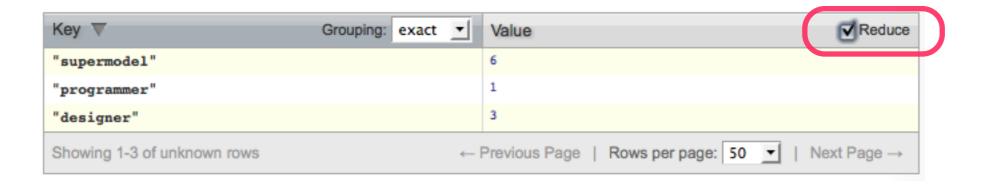
Result of Even Simpler View



A Simple Reduce

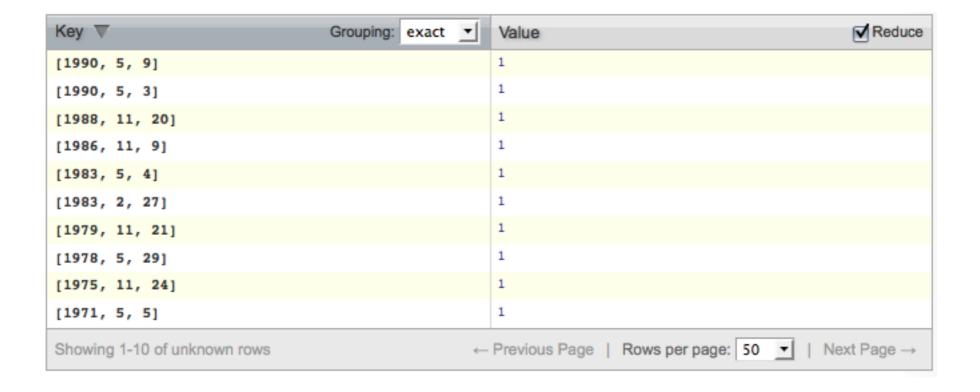
```
function(keys, values) {
  return sum(values)
}
```

Result of a Simple Reduce

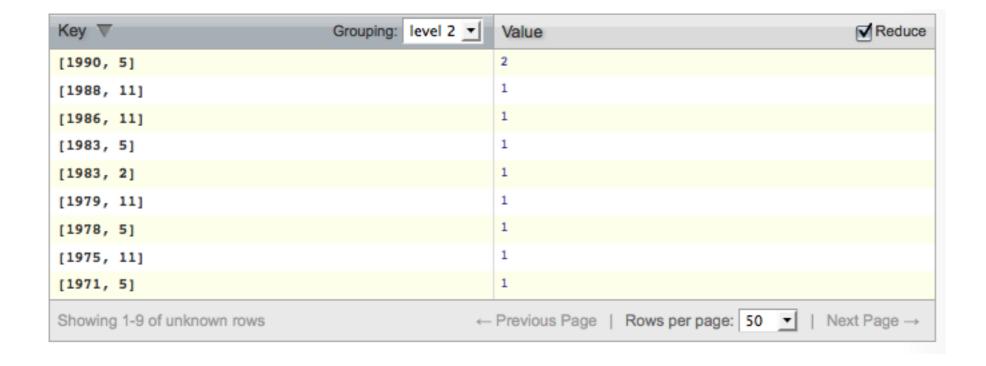


Group Levels

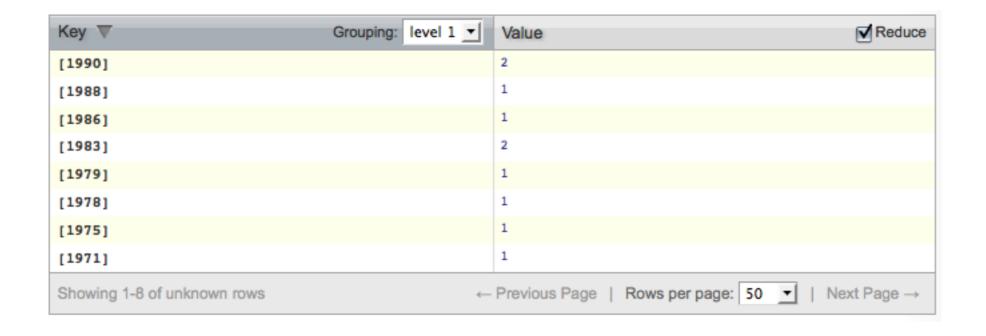
Group Level Exact



Group Level 2



Group Level 1



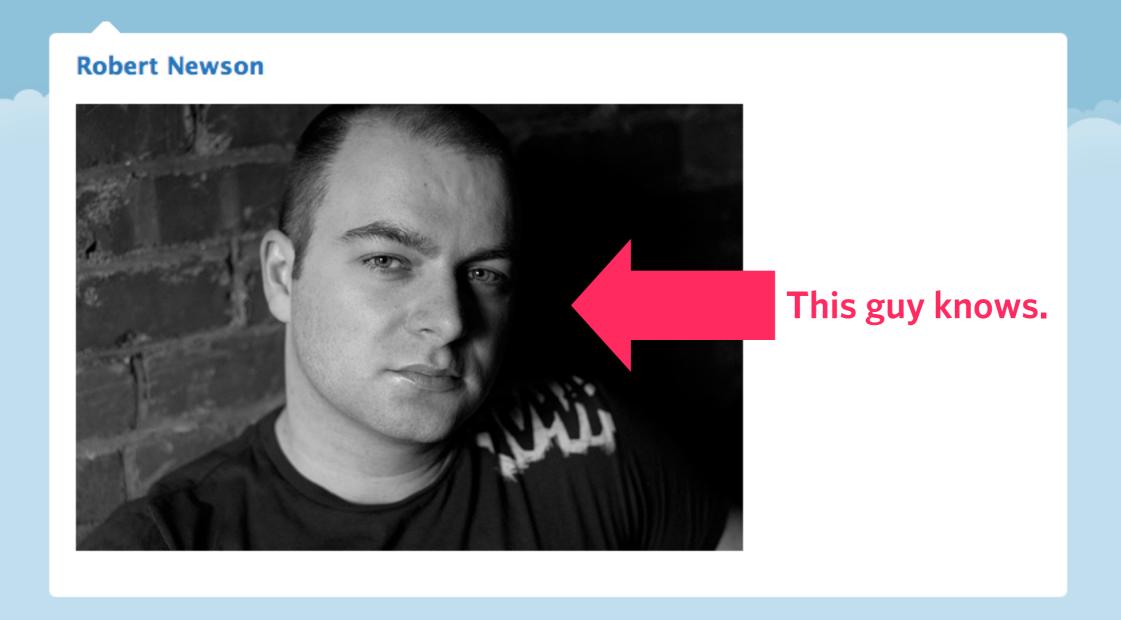
Complex Queries

So... What if you need something like:

Show me all designers who live in Beckerborough.

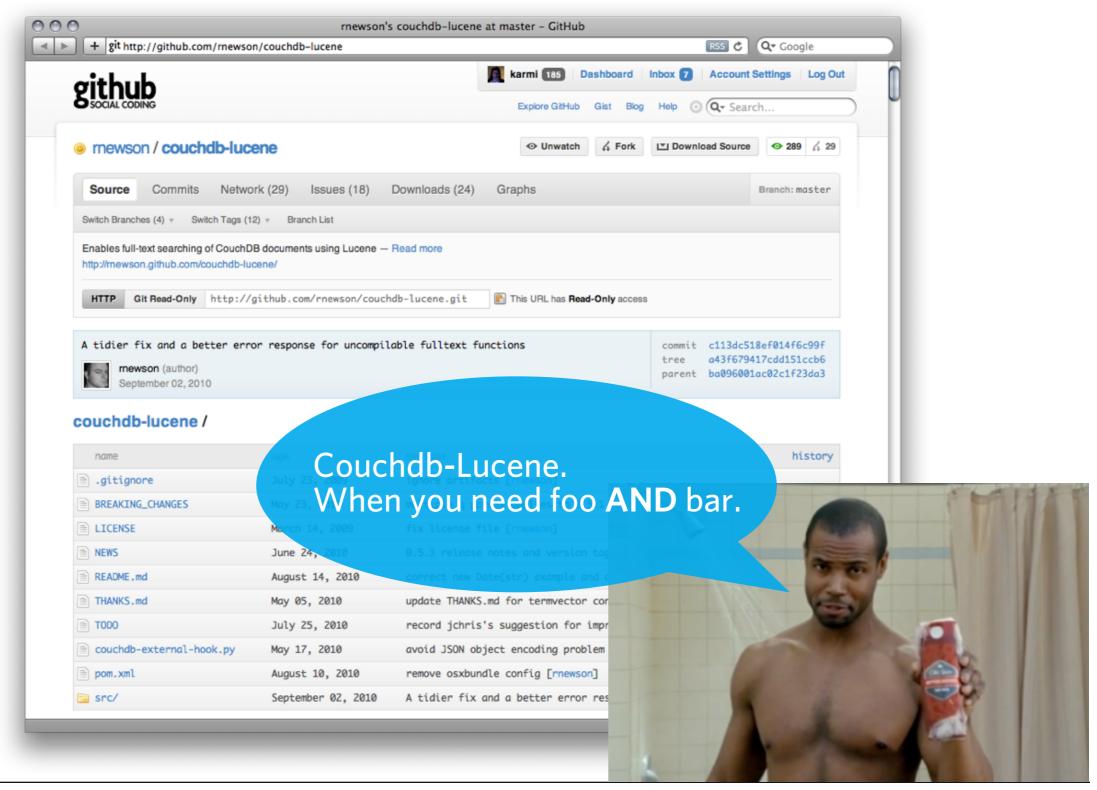
Out of luck?

CouchDB-Lucene

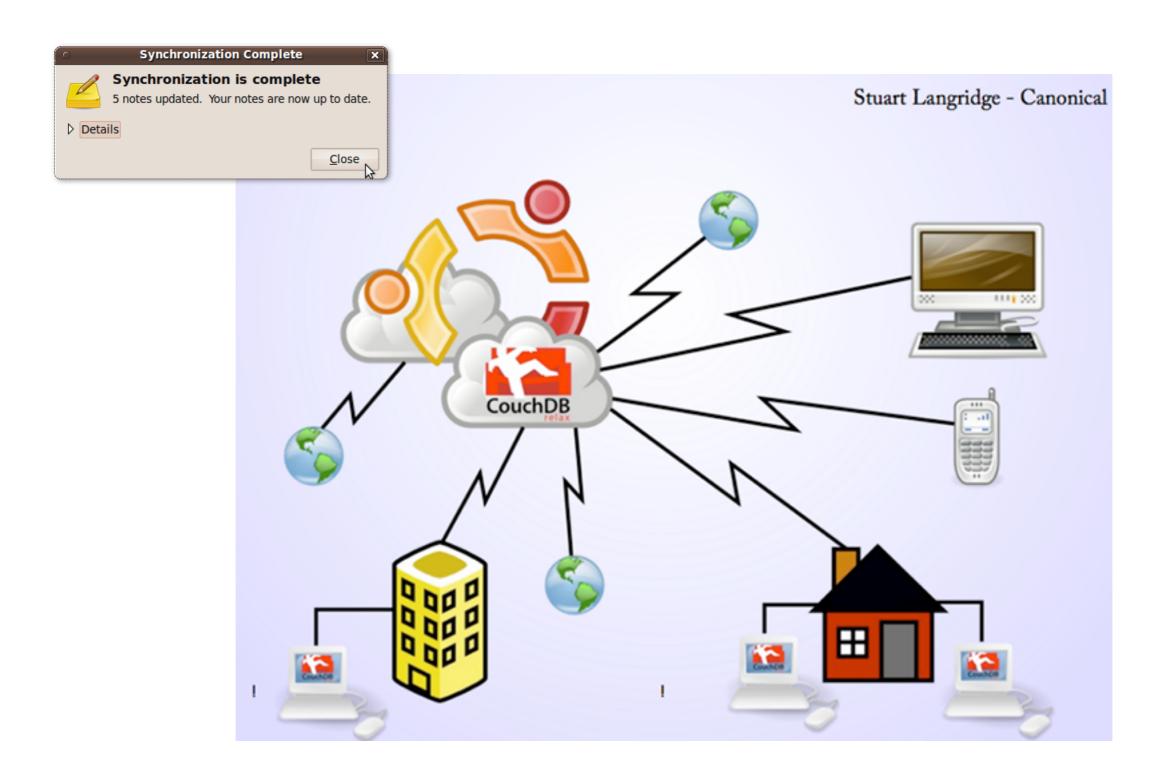


Show me all designers who live in Beckerborough.

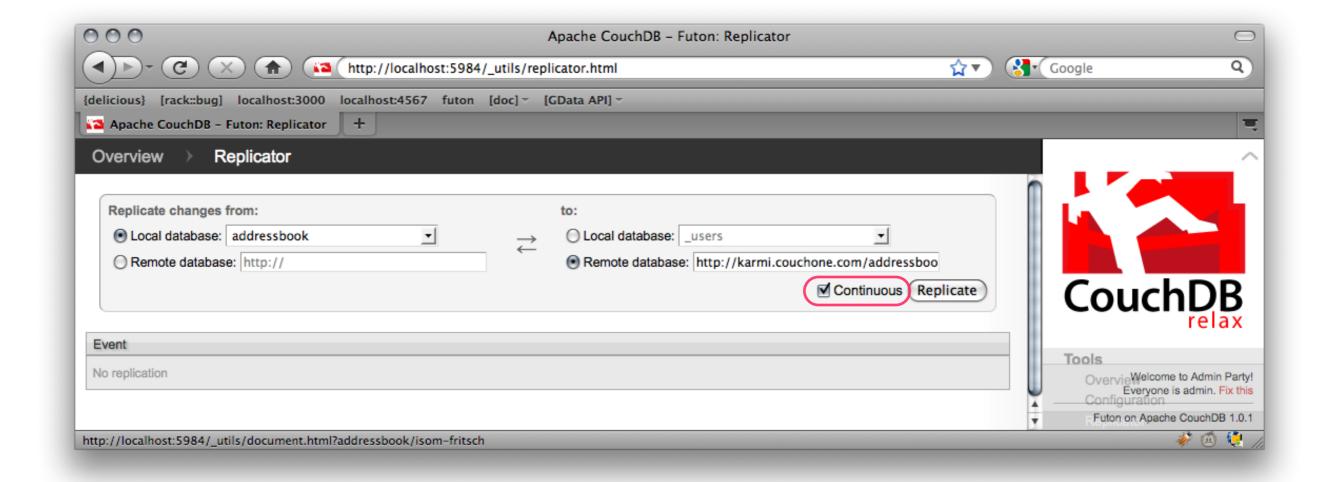
foo AND bar



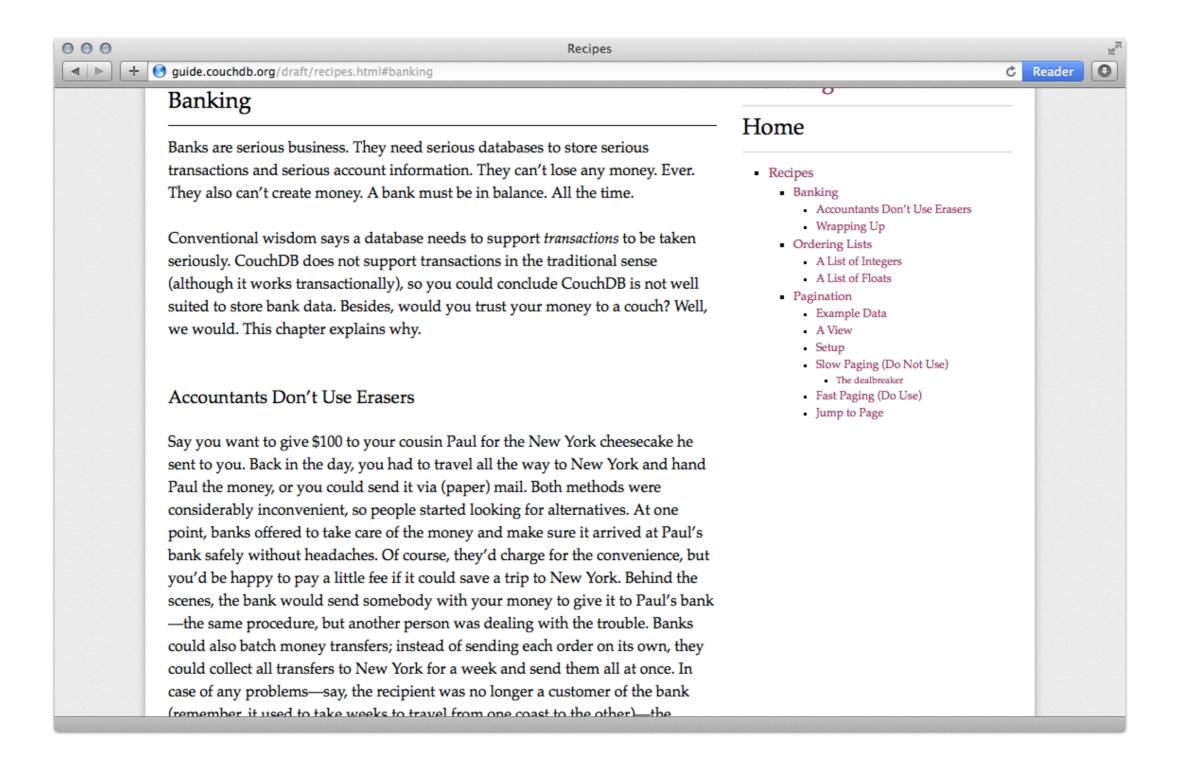
Ubuntu One



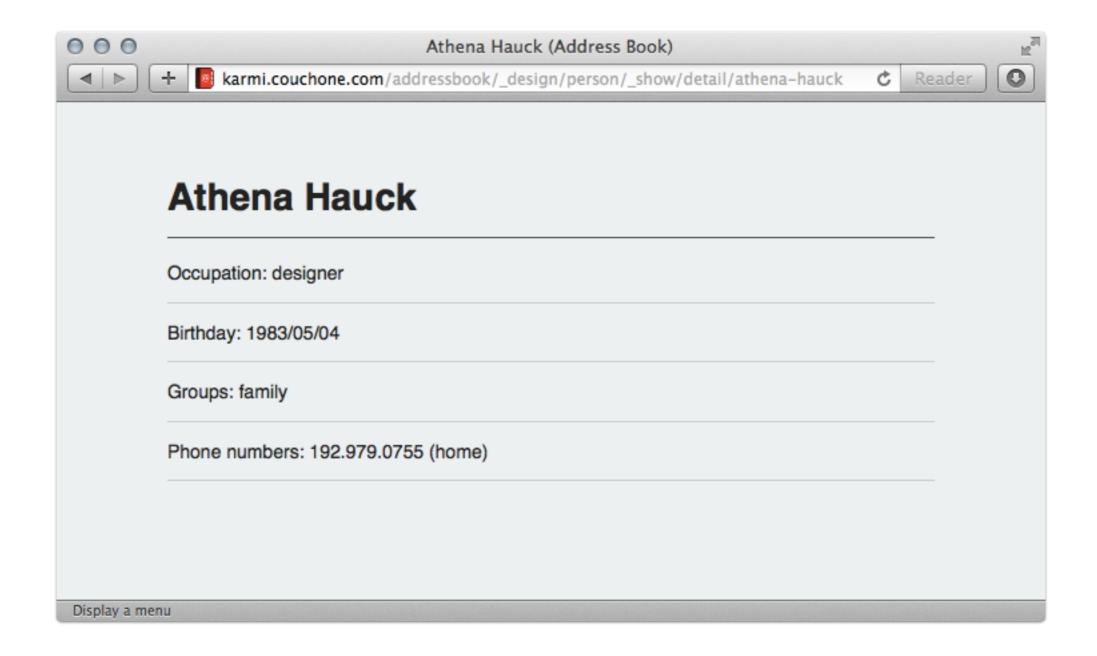
Replication



"Accountants don't use erasers"



Example "CouchApp"



SOURCE CODE: http://github.com/karmi/couchdb-showcase





AK-47

noun

a type of assault rifle, originally manufactured in the Soviet Union.

ORIGIN acronym for Russian *Avtomat Kalashnikova 1947*, the designation of the original model, designed in 1947 by Mikhail T. Kalashnikov (b. 1919).



http://www.youtube.com/watch?v=J6c3DLIM9KA#t=8m32s

AK-47

Designed at the end of WWII by Mikhail Kalashnikov

Driven by changes in military strategy and tactics

Extremely simple design

Designed for mass production & low-quality manufacturing

Extreme reliability, at the cost of accuracy

Reliability



 $\begin{array}{l} sim \bullet plic \bullet i \bullet ty \mid sim ^l plisit \bar{e} \mid \\ noun \end{array}$

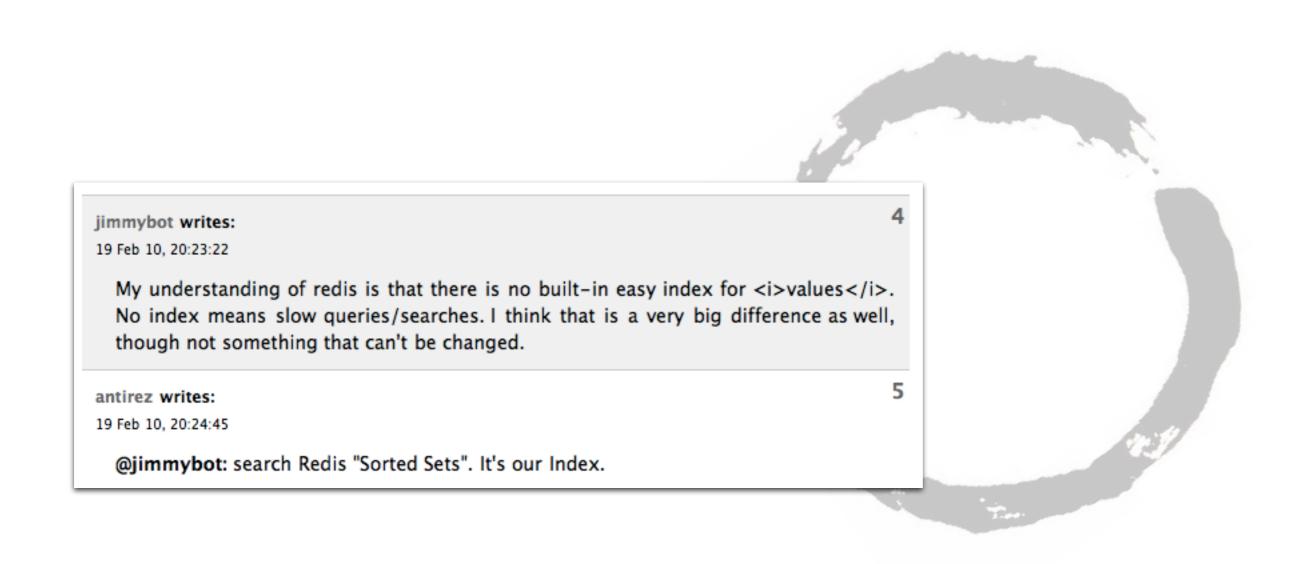
Reliability Speed

Simplicity Versatility

Predictability Low Footprint



The Zen of Redis



(...) what Redis provides are data structures (...)

Strings

```
SET key "value" ~ 4 billions of keys

GET key

=> "value"

DEL key
```

Expiration

```
EXPIRE key 5
GET key
=> "value"
TTL key
=> 1
GET key
=> (nil)
```

Cache

http://antirez.com/post/redis-as-LRU-cache.html

Sessions

https://github.com/mattmatt/redis-session-store

Atomic Increments

```
GET key
=> nil
INCR key
=> 1
INCR key
=> 2
GET key
=> 2
```

```
Counters (downloads, hits, votes, ...)
$ curl http://example.com/downloads/file1.mpg
INCR downloads:total
INCR downloads:total:today
INCR downloads:total:2011-05-10
INCR downloads:/downloads/file1.mpg:total
INCR downloads:/downloads/file1.mpg:today
INCR downloads:/downloads/file1.mpg:2011-05-10
```

```
Counters (downloads, hits, votes, ...)
# Total downloads for server, all time
GET downloads:total
# Total downloads for server, today
GET downloads:total:today
# Total downloads for file
GET downloads:/downloads/file1.mpg:total
# Total downloads for file today
GET downloads:/downloads/file1.mpg:today
```

Counters (downloads, hits, votes, ...)

```
# Expire at 2011-05-10 23:59:59
```

EXPIREAT downloads:total:today 1305064799

All this runs at super-sonic speed, with minimal overhead and resource consumption.

See implementation for Rubygems.org: https://gist.github.com/296921

However, you'll hit denormalization bloat once you start adding metrics (eg. downloads per country, per category, ...)

```
Variations: Rate limiting

$ curl http://api.example.com/list.json

INCR api:<TOKEN>:hits
=> 1
```

```
if INCR('api:abc123:hits') > LIMIT
  return 420 Enhance Your Calm
end
```

```
# Every hour...
DEL api:<TOKEN>:hits
```

Lists

```
LPUSH key 1
                     RPOP key
                     => "1"
=> 1
                     LRANGE key 0 -1
LPUSH key 2
                     => "3"
=> 2
                     => "2"
LPUSH key 3
                     LLEN key
=> 3
LRANGE key 0 -1
                     => 2
=> "3"
                     LTRIM key 0 1
=> "2"
                     => OK
=> "1"
```

```
Indexes (list of comments, ...)
LPUSH article:comments:<ID> "New comment"
Timelines (of all sorts: messages, logs, ...)
LPUSH user:<ID>:inbox "Message from Alice"
LPUSH user:<ID>:inbox "Message from Bob"
# Limit the messages to 100
LTRIM user:<ID>:inbox 0 99
# Get last 10 messages
LRANGE user:<ID>:inbox 0 9
# Get next 10 messages
LRANGE user:<ID>:inbox 10 19
```

Queues

```
# Publisher
RPUSH queue "task-1"
RPUSH queue "task-2"

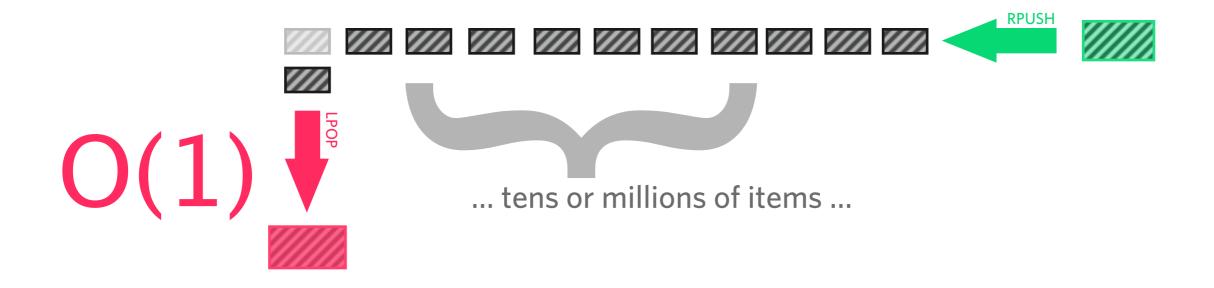
# Worker (blocks and waits for tasks)
BLPOP queue 0
```

Queues

```
# publisher.sh
for i in {1..10}; do
    redis-cli RPUSH "queue" "task-$i"
done

# worker.sh
while true; do
    redis-cli BLPOP "queue" 0
done
```

Resque: Background Processing from Github

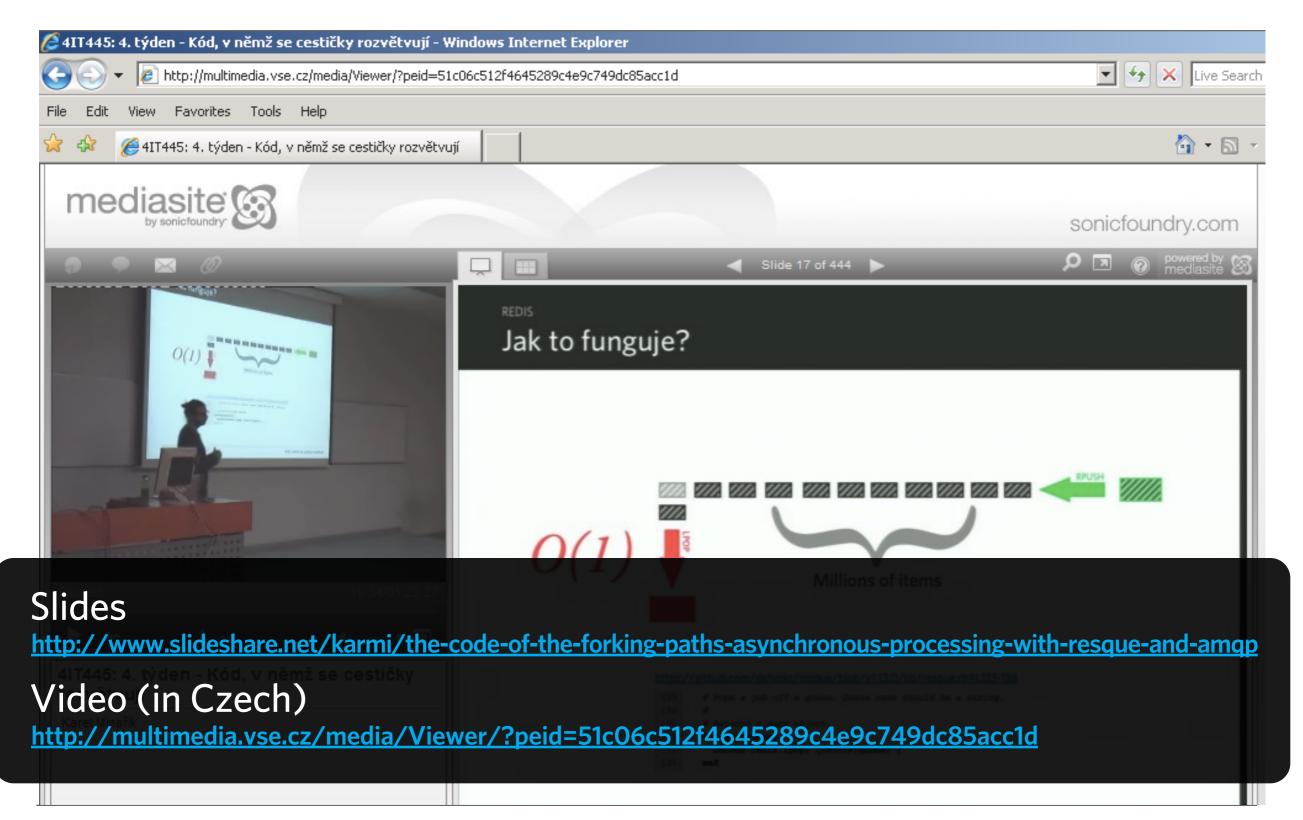


https://github.com/defunkt/resque/blob/v1.13.0/lib/resque.rb#L133-138

```
# Pops a job off a queue. Queue name should be a string.
#
#
# Returns a Ruby object.

def pop(queue)
    decode redis.lpop("queue:#{queue}")
end
```

The Code of the Forking Paths



Sets

```
SISMEMBER key 1
SADD key 1
                     => "1"
=> 1
SADD key 2
                     SISMEMBER key 5
                     => "0"
=> 2
SADD key 3
                     SRANDMEMBER key
                     => "<RAND>"
=> 3
SMEMBERS key
                     SREM key 3
=> "3"
                     => 1
=> "1"
=> "2"
```

Set Operations

SADD A 1

SADD A 2

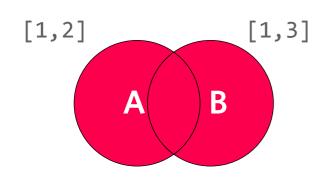
SADD B 1

SADD B 3

SMEMBERS A

SMEMBERS B

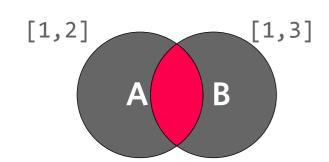
Set Operations



Union

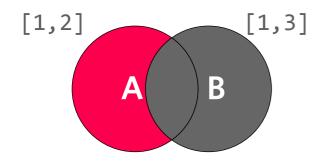
SUNION A B





Intersection

SINTER A B



Difference

SDIFF A B

=> 2

Ad serving

```
SADD ads:cars "Check out Toyota!"

SADD ads:cars "Check out Ford!"

...

SADD ads:movies "Check out Winter's Bone!"

...
```

SRANDMEMBER ads:cars

SRANDMEMBER ads:movies

Note: Time complexity is O(1). "Check out ODER BY RAND()!"

```
Relations (Friends/followers)
```

```
SADD users:A:follows B
```

SADD users:B:follows C

SADD users:B:follows D

SADD users:C:follows A

SADD users:C:follows D

Relations (Friends/followers)

Joint network of A and B

SUNION users:A:follows users:B:follows

- 1) "C"
- 2) "D"
- 3) "B'

```
Relations (Friends/followers)
# Common for A and B
SINTER users:A:follows users:B:follows
# Common for B and C
SINTER users:B:follows users:C:follows
1) "D"
# Unique to B compared to C
SDIFF users:B:follows users:C:follows
```

Relations (Friends/followers)

Relations (Friends/followers)

```
# Mutual relationships
SINTER friends followers
# Who does not follow me back?
SDIFF friends followers
1) "A"
# Who am I not following back?
SDIFF followers friends
```

Relations (Article tags/categories)

SADD tags:ruby article-1

SADD tags:java article-2

SADD tags:web article-1

SADD tags:web article-2

Relations (Article tags/categories) # ruby OR java SUNION tags:ruby tags:java 1) "article-2" 2) "article-1" # ruby AND java SINTER tags:ruby tags:java # web AND NOT ruby SDIFF tags:web tags:ruby 1) "article-2"

Friends Online

```
# My friends
SADD friends A
SADD friends B
SADD friends C
# Friend A performs an action on the site...
SADD online: fresh A
SUNIONSTORE online online: fresh online: stale
# Who's online now?
SINTER friends online
\Gamma A I
```

Friends Online

```
# Every minute, rename the "fresh" to "stale" ...
RENAME online: fresh online: stale
# ... and update the "online" set
SUNIONSTORE online online: fresh online: stale
# Friend B performs an action on the site...
SADD online: fresh B
SUNIONSTORE online online: fresh online: stale
# Who's online now?
SINTER friends online
[A,B]
```

Friends Online

```
# Time passes ...
# Rename the "fresh" to "stale", every minute ...
RENAME online: fresh online: stale
# ... and update the "online" set
SUNIONSTORE online online: fresh online: stale
# Who's online now?
SINTER friends online
IBI
```

Sorted Sets

ZADD key 100 A

ZADD key 10 C

ZADD key 80 B

ZRANGE key 0 -1

- 1) "C"
- 2) "B"
- 3) "A"

ZREVRANGE key 0 -1

- 1) "A"
- 2) "B"
- 3) "C"

ZINCRBY key 10 C

"20"

Sorted Sets

ZREVRANGE key 0 -1

WITHSCORES

- 1) "A"
- 2) "100"
- 3) "B"
- 4) "80"
- 5) "C"
- *6)* "20"

ZREVRANGEBYSCORE

key 100 50

- 1) "A"
- 2) "B"

4) "15"

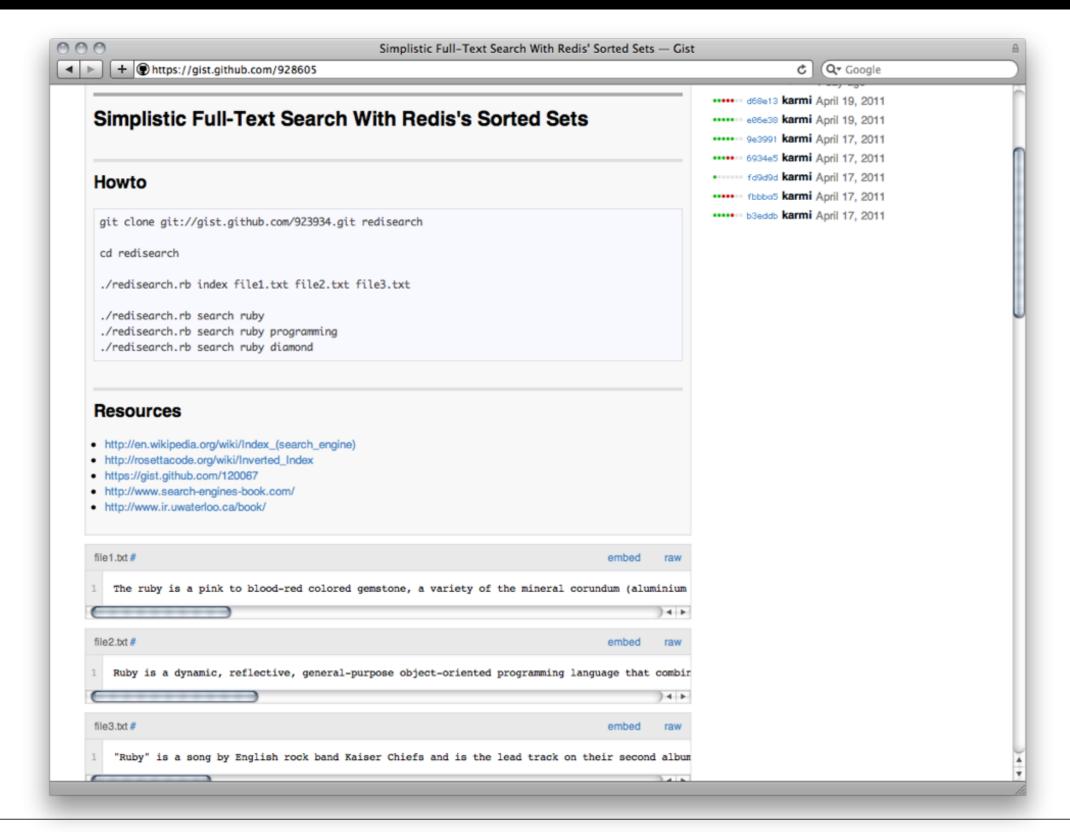
Leaderboards

```
# User A got 10 points
ZINCRBY scores 10 A
# User B got 15 points
ZINCRBY scores 15 B
# User A got another 10 points
ZINCRBY scores 10 A
# Display scores
ZREVRANGE scores 0 -1 WITHSCORES
1) "A"
2) "20"
```

Inverted Index

```
# Index document A
ZINCRBY index: foo 1 document-A
ZINCRBY index: foo 1 document-A
ZINCRBY index:bar 1 document-A
# Index document B
ZINCRBY index: foo 1 document-B
ZINCRBY index:baz 1 document-B
# Search for token foo, sort by occurences
ZREVRANGE index: foo 0 -1 WITHSCORES
1) "document-A"
3) "document-B"
```

Inverted Index



```
Counters (downloads, hits, votes, ...)
```

Total downloads for server, all time GET downloads:total

Mm'kay. But I need the "leaderboard" of downloads!

GET downloads:/downloads/file1.mpg:total

Total downloads for file today
INCR downloads:/downloads/file1.mpg:today

Counters (downloads, hits, votes, ...)

```
INCR downloads:total
INCR downloads:total:today
                                        Easy!
# > Individual files
  GET /downloads/file1.mpg
ZINCRBY downloads:files:total 1 /downloads/file1.mpg
ZINCRBY downloads:files:2011-05-10 1 /downloads/file1.mpg
    GET /downloads/file2.mpg
ZINCRBY downloads:files:total 1 /downloads/file2.mpg
ZINCRBY downloads:files:2011-05-10 1 /downloads/file2.mpg
```

```
Counters (downloads, hits, votes, ...)
# 10 most downloaded, all the time
ZREVRANGE downloads:files:total 0 10 WITHSCORES
# 10 most downloaded on 2011-05-10
ZREVRANGE downloads:files:2011-05-10 0 10 WITHSCORES
# 10 most downloaded between 2011-05-10 and <OTHER DATE>
ZUNIONSTORE downloads:timeline 2 \
            downloads:files:2011-05-10 \
            <OTHER DATE>
```

ZREVRANGE downloads:timeline 0 10 WITHSCORES

Hashes

```
HMSET users:1 username j name John
HMSET users:2 username m name Mary
HGETALL users:1
1) "username"
2) "j"
HKEYS users:1
1) "username"
2) "name"
HSET users:1 score 100
HGET users:1 score
1) "100"
```

```
Structured data (Articles, users, ...)
HMSET articles:1 title
                               "Redis is cool!" \
                    content "I recently ..." \
                    published "2011-05-10"
HGETALL articles:1
1) "title"
2) "Redis is cool!"
3) "content"
4) "I recently ..."
5) "published"
6) "2011-05-10"
HSET articles:1 title "Redis is very cool!"
HGET articles:1 title
"Redis is very cool!"
```

User Preferences (No login)

```
# Save preferences from <FORM>
HMSET prefs:<COOKIE HASH> background #ccc color #333
```

Keep it for one year
EXPIRE prefs:<COOKIE HASH> 31556926

Retrieve preferences
HGETALL prefs:<COOKIE HASH>

Publish/Subscribe

```
SUBSCRIBE log.error "ERROR"

PSUBSCRIBE log.*

PUBLISH log.error "ERROR"

-> "ERROR"

-> "ERROR"
```

PUBLISH

log.info "INFO"

=> "INFO"

Scripting

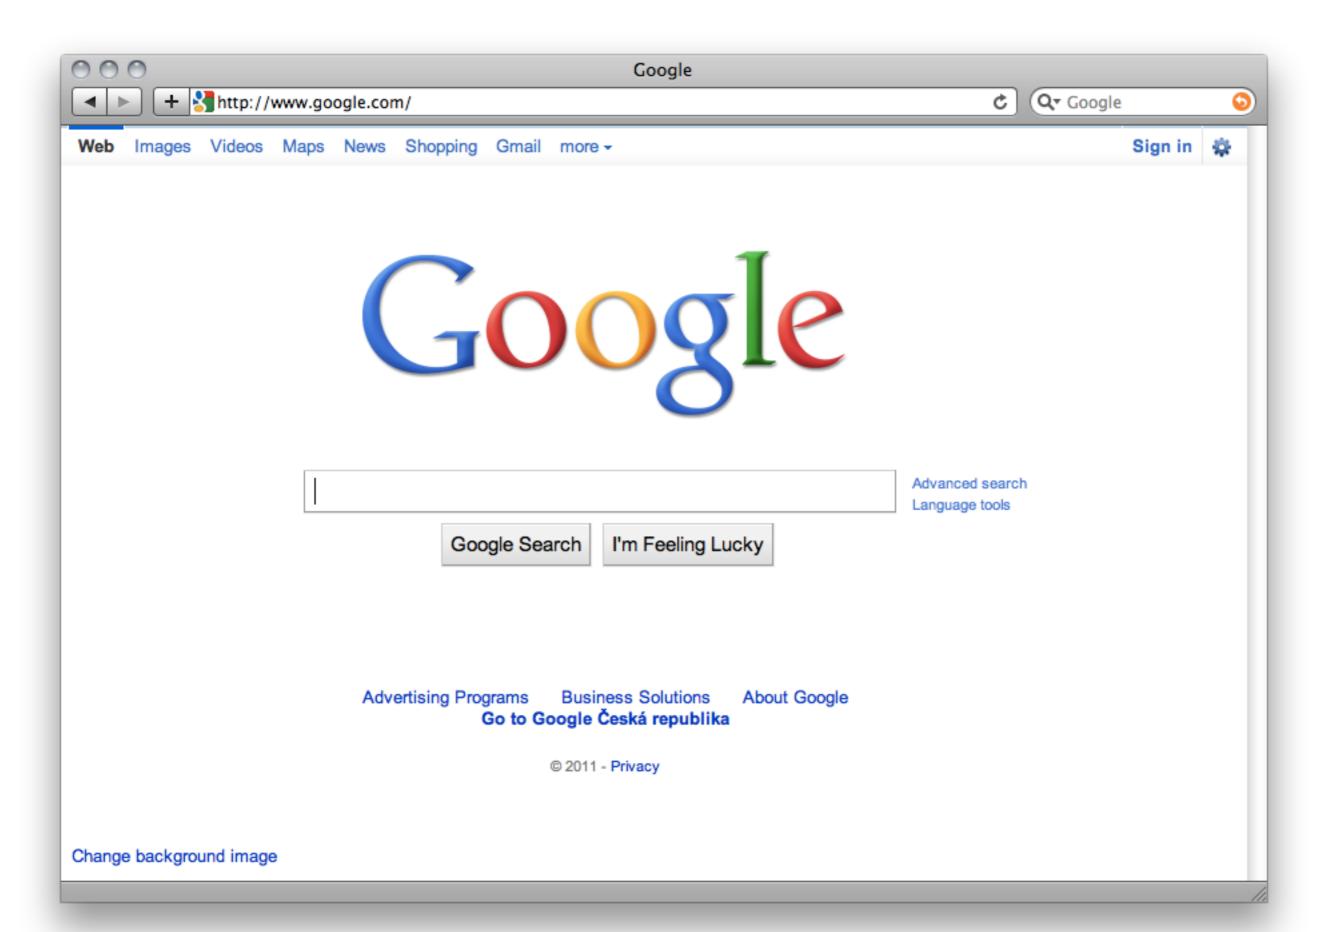
```
SET myscript "return 'HELLO'"
EVAL "return loadstring(redis.call('get', KEYS[1]))()" 1 myscript
```

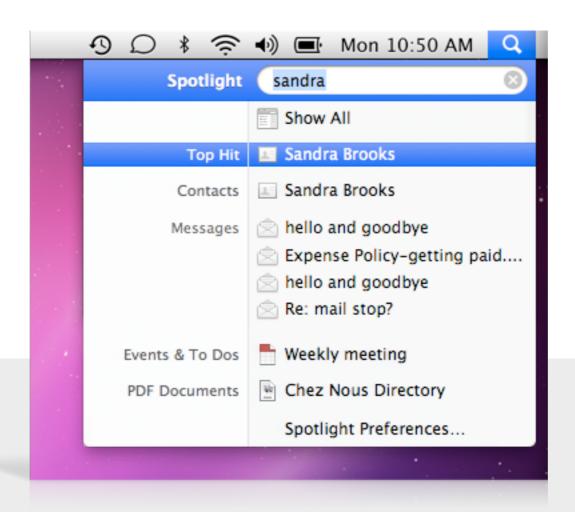
=> "HELLO"



elasticsearch.

Search is the primary interface for getting information today.





http://www.apple.com/macosx/what-is-macosx/spotlight.html



http://www.apple.com/iphone/features/search.html

How does search work?

A collection of documents



file_1.txt

The ruby is a pink to blood-red colored gemstone ...



file_2.txt

Ruby is a dynamic, reflective, general-purpose object-oriented programming language ...



file_3.txt

"Ruby" is a song by English rock band Kaiser Chiefs ...

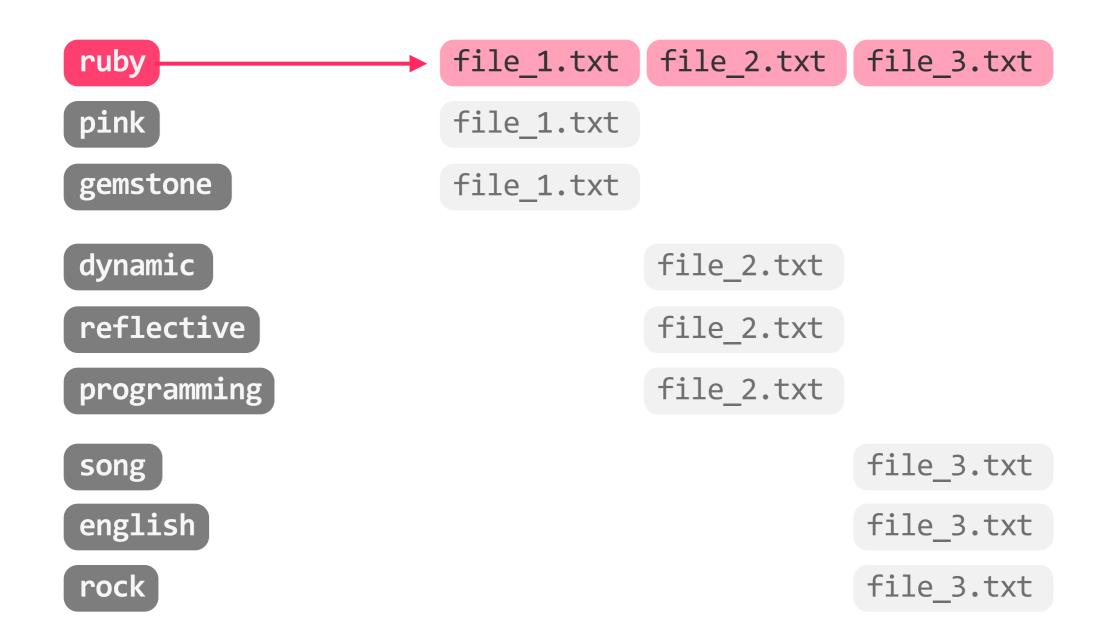
How do you search documents?

```
File.read('file_1.txt').include?('ruby')
File.read('file_2.txt').include?('ruby')
```

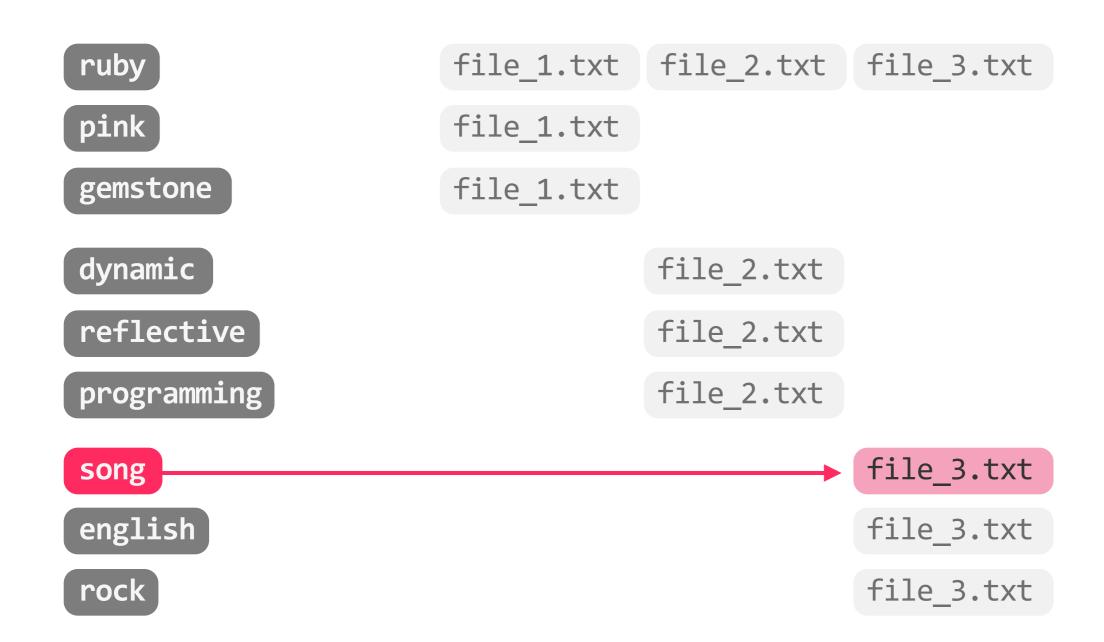
The inverted index

TOKENS	POSTINGS		
ruby	file_1.txt	file_2.txt	file_3.txt
pink	file_1.txt		
gemstone	file_1.txt		
dynamic		file_2.txt	
reflective		file_2.txt	
programming		file_2.txt	
song			file_3.txt
english			file_3.txt
rock			file 3.txt

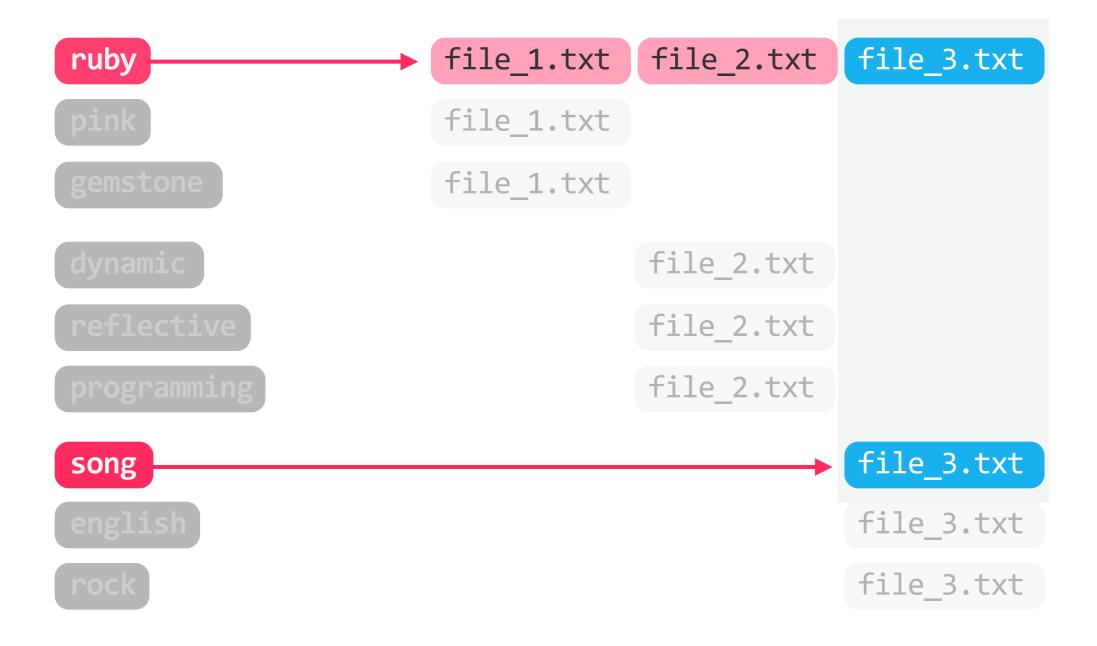
MySearchLib.search "ruby"



MySearchLib.search "song"



MySearchLib.search "ruby AND song"



```
def index document, content
    tokens = analyze content
    store document, tokens
    puts "Indexed document #{document} with tokens:", tokens.inspect, "\n"
 end
 def analyze content
    # >>> Split content by words into "tokens"
   content.split(/\W/).
   # >>> Downcase every word
   map { | word | word.downcase }.
   # >>> Reject stop words, digits and whitespace
    reject { |word| STOPWORDS.include?(word) || word =~ /^\d+/ || word == '' }
 end
 def store document id, tokens
    tokens.each do | token |
      # >>> Save the "posting"
       (INDEX[token] | = []) << document id ).uniq!
   end
 end
 def search token
    puts "Results for token '#{token}':"
   # >>> Print documents stored in index for this token
    INDEX[token].each { |document| " * #{document}" }
 end
  INDEX = \{\}
 STOPWORDS = %w a an and are as at but by for if in is it no not of on or that the then there
 extend self
end
```

Indexing documents

```
SimpleSearch.index "file1", "Ruby is a language. Java is also a language.
SimpleSearch.index "file2", "Ruby is a song."
SimpleSearch.index "file3", "Ruby is a stone."
SimpleSearch.index "file4", "Java is a language."
Indexed document file1 with tokens:
["ruby", "language", "java", "also", "language"]
Indexed document file2 with tokens:
["ruby", "song"] ←

    Words downcased,

                                                        stopwords removed.
Indexed document file3 with tokens:
["ruby", "stone"]
Indexed document file4 with tokens:
["java", "language"]
```

The index

```
puts "What's in our index?"
p SimpleSearch::INDEX
 "ruby" => ["file1", "file2", "file3"],
 "language" => ["file1", "file4"],
 "java" => ["file1", "file4"],
 "also" => ["file1"],
 "stone" => ["file3"],
 "song" => ["file2"]
```

Search the index

SimpleSearch.search "ruby"

```
Results for token 'ruby':
```

- * file1
- * file2
- * file3

TOKENS	POSTINGS

ruby 3

pink 1

gemstone

dynamic

reflective

programming

song

english

rock

file_1.txt

file_1.txt

file_2.txt

file_2.txt

file_2.txt

file_3.txt

file_3.txt

file_3.txt

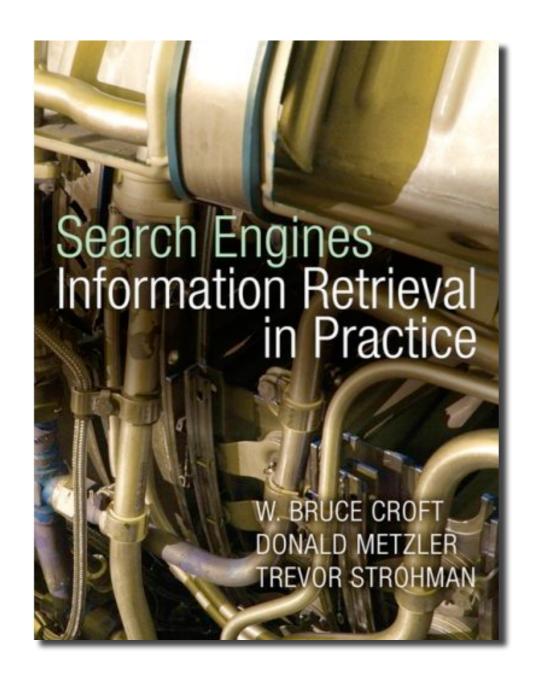
It is very practical to know how search works.

For instance, now you know that the analysis step is very important.

It's more important than the "search" step.

```
def index document, content
 tokens = analyze content
  store document, tokens
 puts "Indexed document #{document} with tokens:", tokens.inspect, "\n"
end
def analyze content
 # >>> Split content by words into "tokens"
  content.split(/\W/).
 # >>> Downcase every word
  map { | word | word.downcase }.
 # >>> Reject stop words, digits and whitespace
  reject { |word| STOPWORDS.include?(word) || word =~ /^{d+}/|| word == '' }
end
def store document id, tokens
 tokens.each do | token |
    # >>> Save the "posting"
      (INDEX[token] ||= []) << document_id ).uniq!</pre>
 end
end
def search token
  puts "Results for token '#{token}':"
 # >>> Print documents stored in index for this token
  INDEX[token].each { |document| " * #{document}" }
end
INDEX = \{\}
STOPWORDS = %w a an and are as at but by for if in is it no not of on or that the then there
extend self
                                                         A naïve Ruby implementation
```

The Search Engine Textbook



Search Engines Information Retrieval in Practice

Bruce Croft, Donald Metzler and Trevor Strohma Addison Wesley, 2009

\$ curl -X GET "http://localhost:9200/_search?q=<YOUR QUERY>"

Terms	apple apple iphone
Phrases	"apple iphone"
Proximity	"apple safari"~5
Fuzzy	apple~0.8
Wildcards	app* *pp*
Boosting	apple^10 safari
Range	[2011/05/01 TO 2011/05/31] [java TO json]
Boolean	apple AND NOT iphone +apple -iphone (apple OR iphone) AND NOT review
Fields	title:iphone^15 OR body:iphone published_on:[2011/05/01 TO "2011/05/27 10:00:00"]

http://lucene.apache.org/java/3_1_0/queryparsersyntax.html

```
curl -X GET localhost:9200/articles/_search -d '{
  "query" : {
    "filtered" : {
     "query" : {
       "bool" : {
          "must" : {
            "match" : {
             "author.first_name" : {
              "query" : "claire",
               "fuzziness" : 0.1
             }
           }
          },
          "must" : {
           "multi_match" : {
             "query" : "elasticsearch",
             "fields" : ["title^10", "body"]
     },
      "filter": {
       "and" : [
         { "terms" : { "tags" : ["search"] } },
         { "range" : { "published_on": {"from": "2013"} } },
         { "term" : { "featured" : true } }
 }
}'
```

```
curl -X GET loca
                    "query" : {
 "query"
                       "filtered" : {
   "filtered"
     "query" :
                          "query" : {
       "bool"
                             "bool" : {
         "must
           "matu
            "author.first_name" : {
              "query" : "claire",
              "fuzziness" : 0.1
             }
          }
         },
         "must" : {
           "multi_match" : {
            "query" : "elasticsearch",
            "fields" : ["title^10", "body"]
     },
     "filter": {
       "and" : [
         { "terms" : { "tags" : ["search"] } },
         { "range" : { "published_on": {"from": "2013"} } },
         { "term" : { "featured" : true } }
 }
}'
```

```
curl -X GET localhost:9200/articles/_search -d '{
  "query" : {
   "filtered" : {
     "query" : {
       "bool" : {
                           "must" : {
                               "match" : {
        "must" : {
          "match" : {
                                  "author.first_name" : {
            "author.first
              "query" : "
                                     "query" : "claire",
              "fuzziness"
                                     "fuzziness" : 0.1
            }
          }
        },
        "must" : {
          "multi_match" :
           "query" : "e
            "fields" : ["tic
     },
     "filter": {
       "and" : [
        { "terms" : { "tags" : ["search"] } },
        { "range" : { "published_on": {"from": "2013"} } },
        { "term" : { "featured" : true } }
 }
}'
```

```
curl -X GET localhost:9200/articles/_search -d '{
  "query" : {
   "filtered" : {
     "query" : {
      "bool" : {
        "must" : {
          "match" : {
            "author.first_name" : {
             "query" : "claire",
              "fuzziness" : 0.1
            }
                             },
          }
        },
        "must" : {
                            "must" : {
          "multi match" :
                               "multi_match" : {
           "query" : "@
            "fields" : ['
                                   "query" : "elasticsearch",
                                  "fields" : ["title^10", "body"]
     },
     "filter": {
       "and" : [
        { "terms" : { "tag
        { "range" : { "publisheu_on : \ ITOM : ZUIS } } ,
        { "term" : { "featured" : true } }
 }
}'
```

```
curl -X GET localhost:9200/articles/_search -d '{
 "query" : {
   "filtered" : {
    "query" : {
      "bool" : {
        "must" : {
         "match" : {
          "author.first_name" : {
            "query" : "claire",
           "fuzziness" : 0.1
           }
         }
        },
        "must" : {
         "multi_match" : {
          "query" : "elasticsearch",
           "fields" : ["title^10", "body"]
                          J 1
     },
                          "filter": {
     "filter": {
                            "and" : [
      "and" : [
       { "terms" : { "
                               { "terms" : { "tags" : ["search"] } },
       { "range" : { "
                               { "range" : { "published_on": {"from": "2013"} } },
        { "term" : { "
                               { "term" : { "featured" : true } }
}'
```

Full-text Search

"Find all articles with 'search' in their title or body, give matches in titles higher score"

Structured Search

"Find all articles from year 2013 tagged 'search'"

Custom Scoring

See custom_score and custom_filters_score queries

The _analyze API

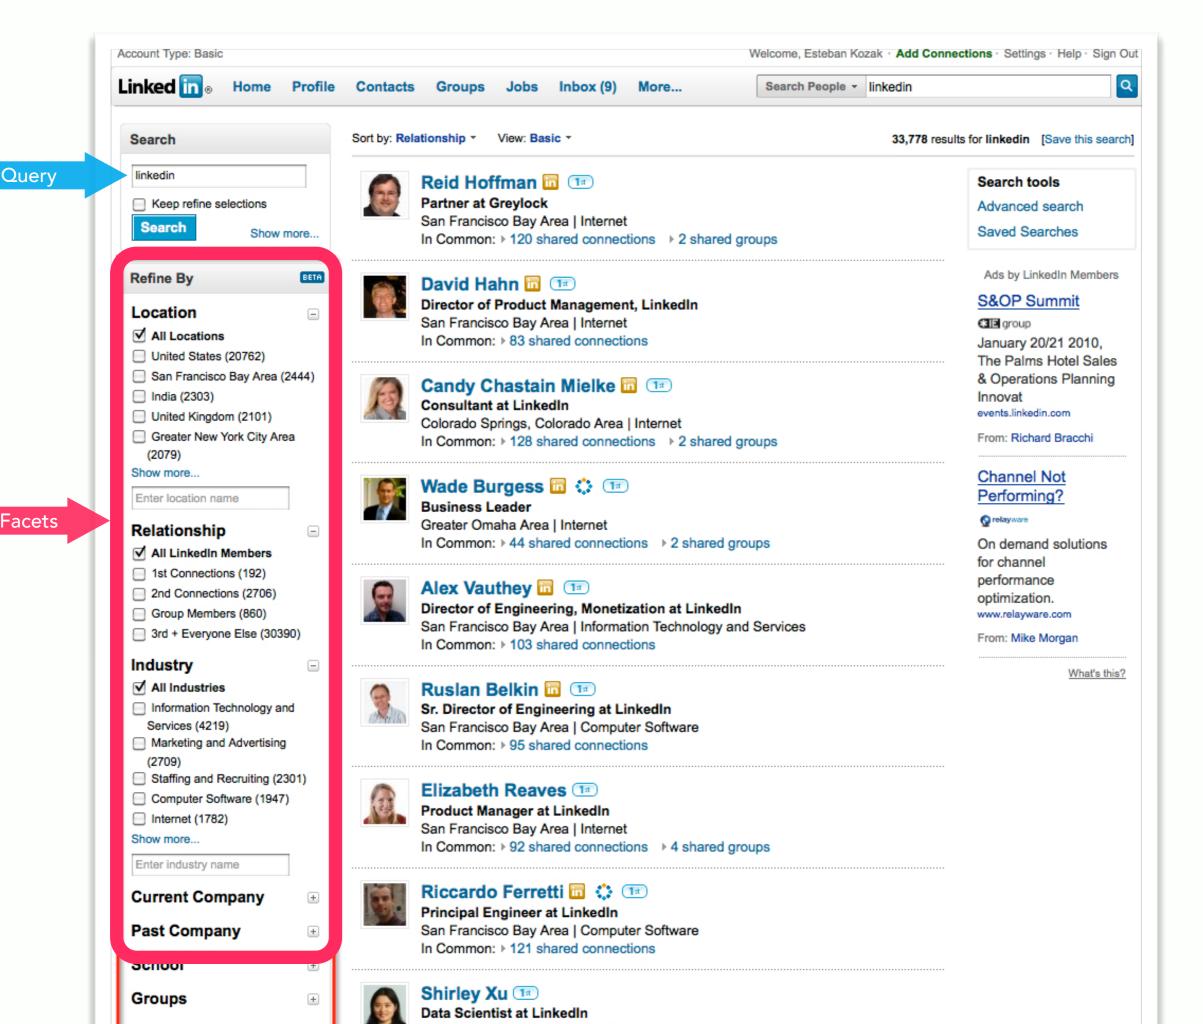
_analyze?text=...&tokenizer=X&filters=A,B,C

```
[žluťoučký:0->9:<ALPHANUM>]\n
\n2: \n[kůň:10-
>13:<ALPHANUM>]\n\n3:
\n[skákal:14->20:<ALPHANUM>]
\n\n4: \n[přes:21-
>25:<ALPHANUM>]\n\n5:
\n[potok:26->31:<ALPHANUM>]
```

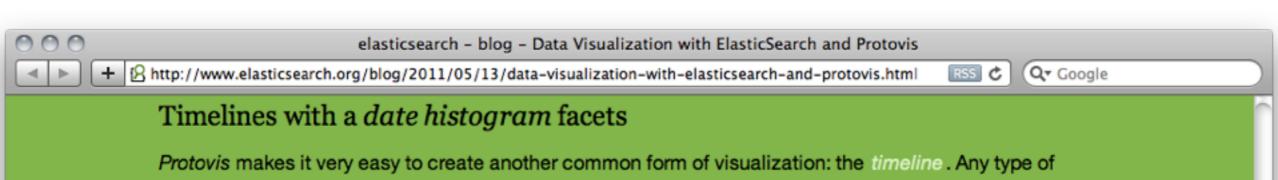
_analyze?pretty&format=text&text=Žluťoučký+kůň+skákal+přes+potok

_analyze?pretty&format=text&text=Žluťoučký+kůň+skákal+přes +potok&analyzer=czech

```
[žluťoučk:0->9:<ALPHANUM>]\n
\n2: \n[koň:10-
>13:<ALPHANUM>]\n\n3:
\n[skákal:14->20:<ALPHANUM>]
\n\n5: \n[potok:26-
>31:<ALPHANUM>]\n
```

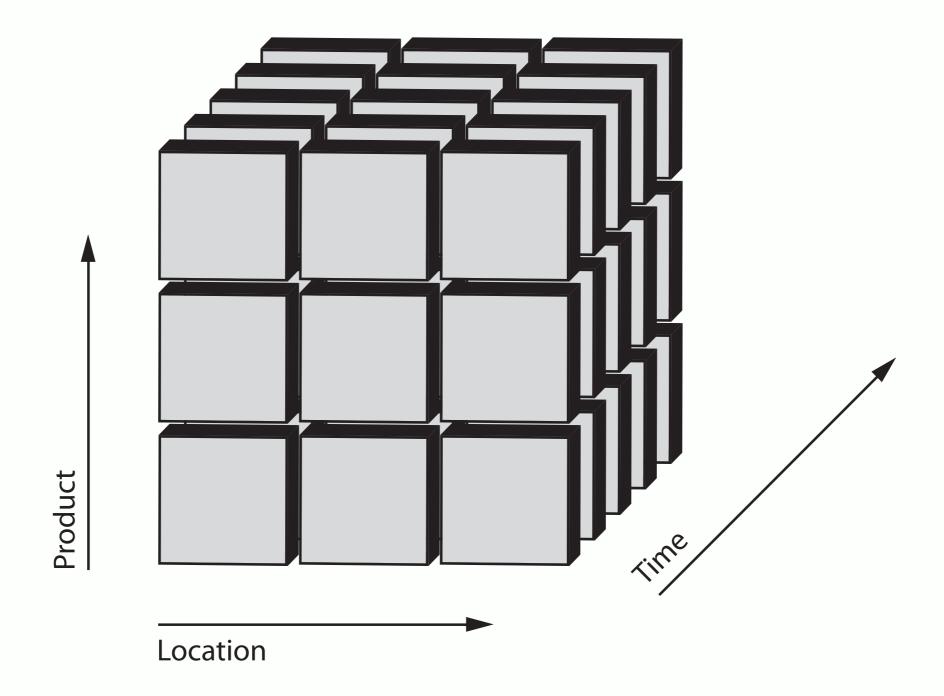


Using facets as analytics



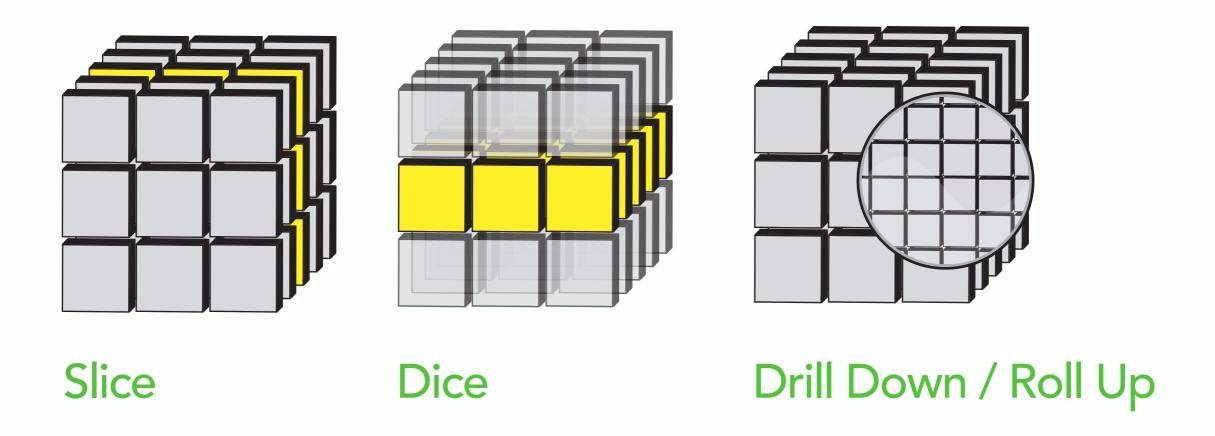
data, tied to a certain date, such as an article being published, an event taking place, a purchase being completed can be visualized on a timeline.

The end result should look like this (again, you may want to check out the working version):



OLAP Cube

Dimensions, measures, aggregations



Show me sales numbers for all products across all locations in year 2013

Show me products sales numbers in location X over all years

Category Count With the terms Aggregation

Simplest "map/reduce" / GROUP_BY aggregation: document count per tag

```
curl -X POST 'localhost:9200/articles/_search?search_type=count&pretty' -d '{
  "aggregations": {
    "categories": {
      "terms" : {
                                               "facets" : {
        "field" : "category"
                                                   "tag-cloug" : {
                                                     "buckets" : [ {
                                                      "key": "ruby",
                                                      "doc_count" : 3
                                                    }, {
                                                      "key" : "java",
                                                      "doc_count" : 2
                                                    },
```

Inner Aggregation within Category

Aggregating number of clicks per category

```
curl -X GET 'localhost:9200/articles/_search/?search_type=count&pretty' -d
"aggregations": {
```

```
"aggregations" : {
    "categories" : {
        "buckets" : [ {
            "key" : "ruby",
            "doc_count" : 3,
            "clicks" : {
                 "value" : 420
            }
        },
        ... ]
    }
}
```

Limit the Aggregation Scope With a Fulltext Query

Realtime filtering with queries and filters

```
curl -X GET 'localhost:9200/articles/_search/?search_type=count&pretty' -d
   "query" : {
        "match" : {
            "body" : "Twitter"
        }
    },
    "aggregations": {
        ....
    }
}'
```

Aggregations

Terms

Significant Terms

Stats

Percentiles

Cardinality

Top Hits

Scripted Metric

Filter

Range

Histogram

http://www.elasticsearch.org/guide/en/elasticsearch/guide/current/aggregations.html

Thanks!

