

# CouchDB — Databáze pro web

---

*Karel Minařík*

# Karel Minařík

- Web designer a vývojář na volné noze od roku 2000
- V minulosti Flash vývojář, art director, informační architekt v internetové agentuře, „více na LinkedIn
- Ruby, Rails, Git, CouchDB *propagandista* v České Republice

→ [karmi.cz](http://karmi.cz)



**CouchDB**

relax



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

<http://wiki.apache.org/couchdb>



Apache CouchDB is **RESTful** and **JSON API**.  
It is **Distributed**, **Schema-Free**, **Document Oriented**, and  **tolerant**.  
Database access.

<http://wiki.apache.org/couchdb>





**NO SQL!!!!**

Replication!  
Decentralization!  
Schema FREE!!!  
JSON, not XML!  
Clusters!  
HTTP!!!!

FAIL





Project Voldemort



# Důvody vzniku a rozšíření „netradičních“ databází

NoSQL není ani „protestní hnutí“ „přechodná móda“.

Důvody pro vznik a rozšíření nových databází jsou **skutečné**.

A většinou pocházejí ze skutečné „bolesti“.

# Denormalizace databáze v Digg.com

```
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>

# Redis: „asymptotická složitost“ součástí balení

```
redis> rpush mylist 1
redis> rpush mylist 2
redis> lpop mylist
"1"
...
redis> llen mylist
(integer) 1 000 000
redis> lpop mylist
"2"
```



Redis **redis**  
A persistent key-value database with built-in net interface written in ANSI-C for

Project Home Downloads Wiki Issues Source

Search  Current pages for  Search

>> ☆ LpopCommand

**LPOP key**


**RPOP key**

**Time complexity:  $O(1)$**

Atomically return and remove the first (LPOP) or last (RPOP) element of the list. return "a" and the list will become "b","c".

```
$ redis-benchmark
...
===== LPOP =====
10010 requests completed in 0.47 seconds
...
96.22% <= 3 milliseconds
```

# Příklad: Asynchronní fronty



<http://github.com/defunkt/resque/blob/master/lib/resque.rb#L133-138>


```
133 # Pops a job off a queue. Queue name should be a string.
134 #
135 # Returns a Ruby object.
136 def pop(queue)
137   decode redis.lpop("queue:#{queue}")
138 end
```

1

## Příběh CouchDB



# Damien Katz: CouchDB and Me



**Damien Katz**

(RubyFringe 2008)

Sell my house,  
move my family,  
and live off  
savings? **WHY?**

# Damien Katz: CouchDB and Me

Na začátku bylo C++, XML a vlastní dotazovací jazyk.

Zkrátka všechno to, co za ještě nikoho nevyhodili.

Pak přišel Erlang, HTTP, JSON a Map/Reduce.

# 2

## Bezeschémové dokumenty

# “Relační data”



[Home](#) [Profile](#) [Find People](#) [Settings](#) [Help](#) [Sign out](#)

OH: “Svět je relační!!!”

17 minutes ago via Tweetie for Mac  
Retweeted by 10000 people

[↩ Reply](#) [↻ Retweet](#)

To ale neznamená, že svět odpovídá  
třetí normální formě.

# Učebnicový příklad

Navrhněte **databázi zákazníků**.

Lidé mají jména, e-maily, telefonní čísla, ...



**Kolik telefonních čísel?**

# Učebnicový příklad

Relační databáze, 1. semestr

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

Customers	
<b>id</b> INTEGER	A N P
first_name	VARCHAR
last_name	VARCHAR
phone	VARCHAR

Tak. A co s více telefonními čísly?

# Učebnicový příklad

„Řešení“, pokus 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

Customers	
<b>id</b> INTEGER	A N P
first_name	VARCHAR
last_name	VARCHAR
phone	VARCHAR

“Stejně budeme s tou databází pracovat jen z naší aplikace.”



# Učebnicový příklad

## „Řešení“, pokus 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		

Customers	
id	INTEGER (A N P)
first_name	VARCHAR
last_name	VARCHAR
phone_1	VARCHAR
phone_2	VARCHAR
phone_3	VARCHAR

“Tohle je přece o moc lepší *architektura!*”

Dobře. Pak zkuste odpovědět na následující otázky:

- Jak vyhledáte zákazníky podle telefonního čísla?
- Kteří zákazníci mají stejné telefonní číslo?
- Kolik telefonních čísel má zákazník?

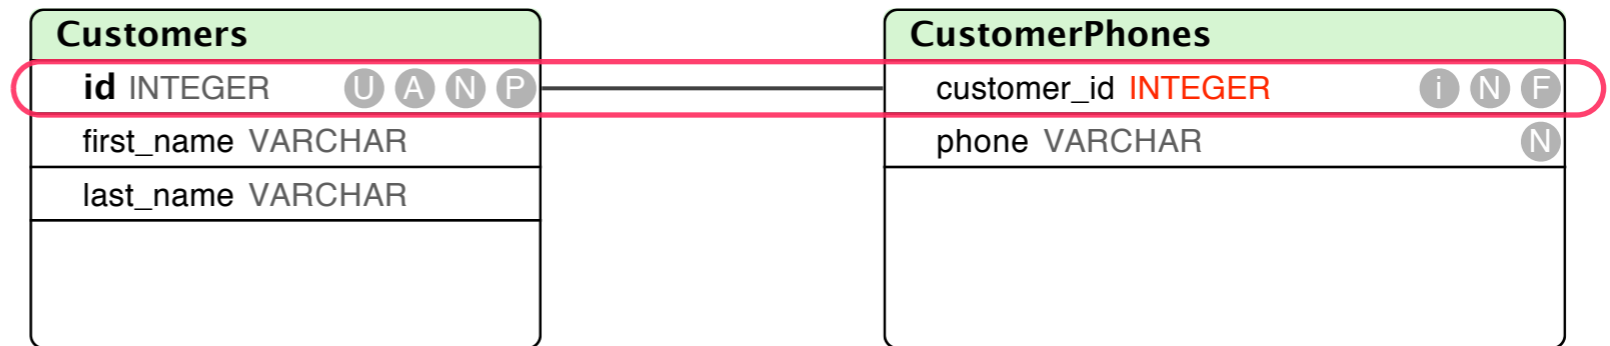
A nyní prosím přidejte možnost ukládat **čtyři** telefonní čísla. Díky.

# Učebnicový příklad

## Správné řešení

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



# Učebnicový příklad

```
mysql> SELECT * FROM Customers LEFT JOIN CustomerPhones
        ON Customers.id = CustomerPhones.customer_id;
```

```
+-----+-----+-----+-----+-----+
| id | first_name | last_name | customer_id | phone |
+-----+-----+-----+-----+-----+
|  1 | John      | Smith    | 1           | 123   |
|  1 | John      | Smith    | 1           | 456   |
+-----+-----+-----+-----+-----+
```

# Učebnicový příklad

```
mysql> SELECT * FROM Customers WHERE id = 1;
```

```
+-----+-----+-----+
| id | first_name | last_name |
+-----+-----+-----+
|  1 | John      | Smith    |
+-----+-----+-----+
```

```
mysql> SELECT phone FROM CustomerPhones WHERE customer_id = 1;
```

```
+-----+
| phone |
+-----+
| 123   |
| 456   |
+-----+
```

# Strukturované dokumenty

Ale hrome!, potřebujeme spíše něco jako:

```
{  
  "id"          : 1,  
  
  "first_name" : "Clara",  
  "last_name"  : "Rice",  
  
  "phones"     : ["0000 777 888 999", "0000 123 456 789", "0000 314 181 116"]  
}
```

“V pohodě. Prostě iteruj přes řádky a sestav si objekt tak, jak ho potřebuješ. Tak to prostě je.”

“Když to bude zlobit, dáme tam nějakou *cache*.”

# Efemérní data

Ne všechno ale musí být uděláno „správně“. Nebo musí?

```
class User < ActiveRecord::Base
  serialize :preferences
end
```

# „Konzistence“

Může „správné řešení“ také selhat?

Může. A zatraceně snadno.

PŘÍKLAD

Když navrhujete fakturační aplikaci, ukládáte zákazníka k faktuře „správně“, přes cizí klíče.

Pak se změní adresa zákazníka.

Změnila se také adresa na faktuře?

# Dokumenty v reálném světě





# Dokumenty v reálném světě



```
{
  "_id"      : "clara-rice",
  "_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"
}
```


# Dokumenty v reálném světě


Field	Value
<code>_id</code>	<code>"lottie-armstrong"</code>
<code>_rev</code>	<code>"2-fcb71b26096957b3ff3ffd2970f3c933"</code>
<code>_attachments</code>	<code>portrait.jpg</code> 72.5 KB, image/jpeg
<code>addresses</code>	<code>home</code> <code>number</code> <code>"753"</code> <code>city</code> <code>"Murphyville"</code> <code>street</code> <code>"Jaqueline Greens"</code> <code>country</code> <code>"United Kingdom"</code> <code>work</code> <code>number</code> <code>"366/99"</code> <code>city</code> <code>"Myrnaland"</code> <code>street</code> <code>"Ruben Stravenue"</code> <code>country</code> <code>"New Zealand"</code>
<code>birthday</code>	<code>"1952/10/24"</code>
<code>first_name</code>	<code>"Lottie"</code>
<code>groups</code>	<code>0</code> <code>"work"</code>
<code>last_name</code>	<code>"Armstrong"</code>
<code>occupation</code>	<code>"programmer"</code>
<code>phones</code>	<code>cell</code> <code>"1-171-069-1373 x74099"</code> <code>work</code> <code>"1-247-705-7343 x45353"</code>

Showing revision 2 of 2

← Previous Version

Address Book



 **Lottie Armstrong**

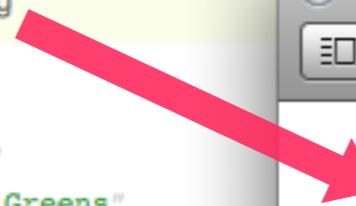
`work` 1-247-705-7343 x45353  
`mobile` 1-171-069-1373 x74099

`work` 366/99 Ruben Stravenue  
Myrnaland

`home` 753 Jaqueline Greens  
Murphyville

Note:

+ Edit 210 cards








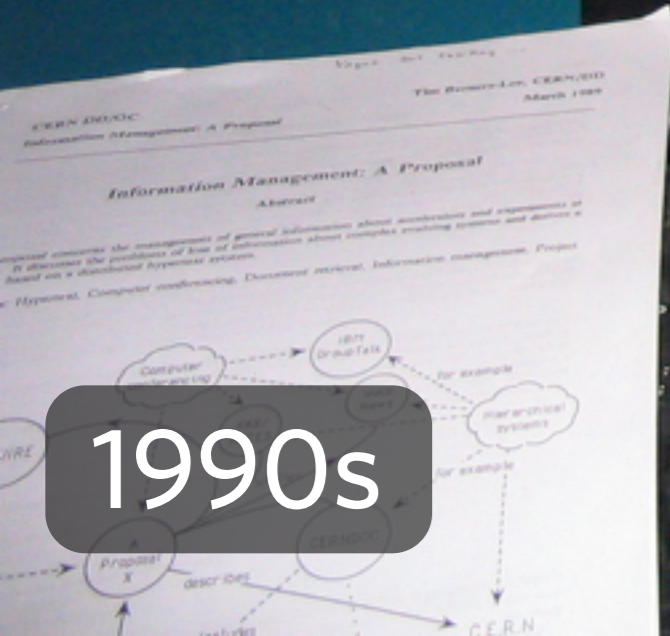

Prokrustovo lože





## RESTful HTTP







# Vytvořeno „z webu“

*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)

# Vytvořeno „z webu“

*HTTP is the lingua franca of our age; if you speak HTTP, it opens up all sorts of doors.*

*There's something almost subversive about CouchDB; it's completely language-, platform-, and OS-agnostic.*

# HTTP rozhraní

```
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"}
```



# Jednoduchá implementace API

```
require 'rubygems'  
require 'ostruct'
```

```
require 'restclient' ← 1 HTTP knihovna  
require 'json' ← 2 JSON knihovna
```

```
class Article < OpenStruct  
  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
```

# Jednoduchá implementace API

```
Article.create :_id => 'my-first-post',
              :title => 'CouchDB is easy',
              :body => 'So relax!',
              :tags => ['couchdb', 'databases'] rescue RestClient::Conflict

article = Article.find('my-first-post')

puts "Got an article:"
p article

puts "\n-----"
puts "Title: %s" % article.title + " (class: #{article.title.class})"
puts "Tags: %s" % article.tags.inspect + " (class: #{article.tags.class})"
puts "-----\n\n"

puts "Deleting article..."
article.destroy
```

# HTTP „od sklepa až na půdu“

```
$ curl -X POST http://localhost:5984/_replicate \  
-d '{"source": "database",  
    "target": "http://example.org/database"}'
```

# Skutečné využití HTTP

```
$ curl -i -X GET $HOST/my-database/abc123
```

```
HTTP/1.1 200 OK
```

```
Server: CouchDB/1.0.1 (Erlang OTP/R14B)
```

```
Etag: "4-f04f2435e031054d6b5298c5841ae052"
```

```
Date: Thu, 23 Sep 2010 12:56:37 GMT
```

```
Content-Type: text/plain;charset=utf-8
```

```
Content-Length: 73
```

```
Cache-Control: must-revalidate
```

```
{"_id":"abc123","_rev":"4-f04f2435e031054d6b5298c5841ae052","foo":"bar"}
```

```
$ cat /etc/squid3/squid.conf
```

```
cache_peer 192.168.100.2 parent 5984 0 no-query originserver name=master  
acl master_acl method GET  
cache_peer_access master allow master_acl
```



**SQUID**

The word "SQUID" is written in a bold, blue, bubbly font with a yellow outline. To the left of the text, there are several blue and white bubbles of varying sizes, suggesting an underwater theme.

# Notifikace o změnách (\_changes)

The image shows a Ruby script on the left and the Apache CouchDB Futon interface on the right. The script is listening for changes in a CouchDB database named 'addressbook'. The output shows three updates to a document with ID 'jarvis-lowe'.

```

karmi@aluir:addressbook$ rake changes DATABASE=addressbook
(in /Users/karmi/Playground/CouchDB/Prednasky/Webexpo2010/code/addressbook)
Listening for changes in 'addressbook' since 0...

[09:51:02] Document ID jarvis-lowe received updates:
- {}
+ {}

[09:51:12] Document ID jarvis-lowe received updates:
- {:occupation=>"supermodel"}
+ {:occupation=>"programmer"}

[09:52:10] Document ID jarvis-lowe received updates:
- {:phones=>{:home=>"1-385-854-9763"}}
+ {:phones=>{:mobile=>"1-777-888-999", :home=>"1-385-854-9763"}}
  
```

The Apache CouchDB Futon interface shows the document 'jarvis-lowe' with the following fields and values:

Field	Value
_id	"jarvis-lowe"
_rev	"3-c56bcc63ceb524f3a0db3052c49cb6cf"
addresses	home
birthday	"1971/11/21"
first_name	"Jarvis"
groups	0 "family" 1 "friends" 2 "work"
last_name	"Lowe"
occupation	"programmer"
phones	home "1-385-854-9763" mobile "1-777-888-999"

```

190 end
191
192 # ---> Run the _changes listener
193 begin
194   CouchDB::Changes.new(ENV['DATABASE']).listen
195 rescue ArgumentError => e
196   puts "[!] #{e.message}"
197   exit(1)
198 end
  
```

# 4

## Spolehlivost a paralelizace





CouchDB nemá vypínač.

CouchDB nemá příkaz *repair*.

```
$ kill -9 <PID>
```



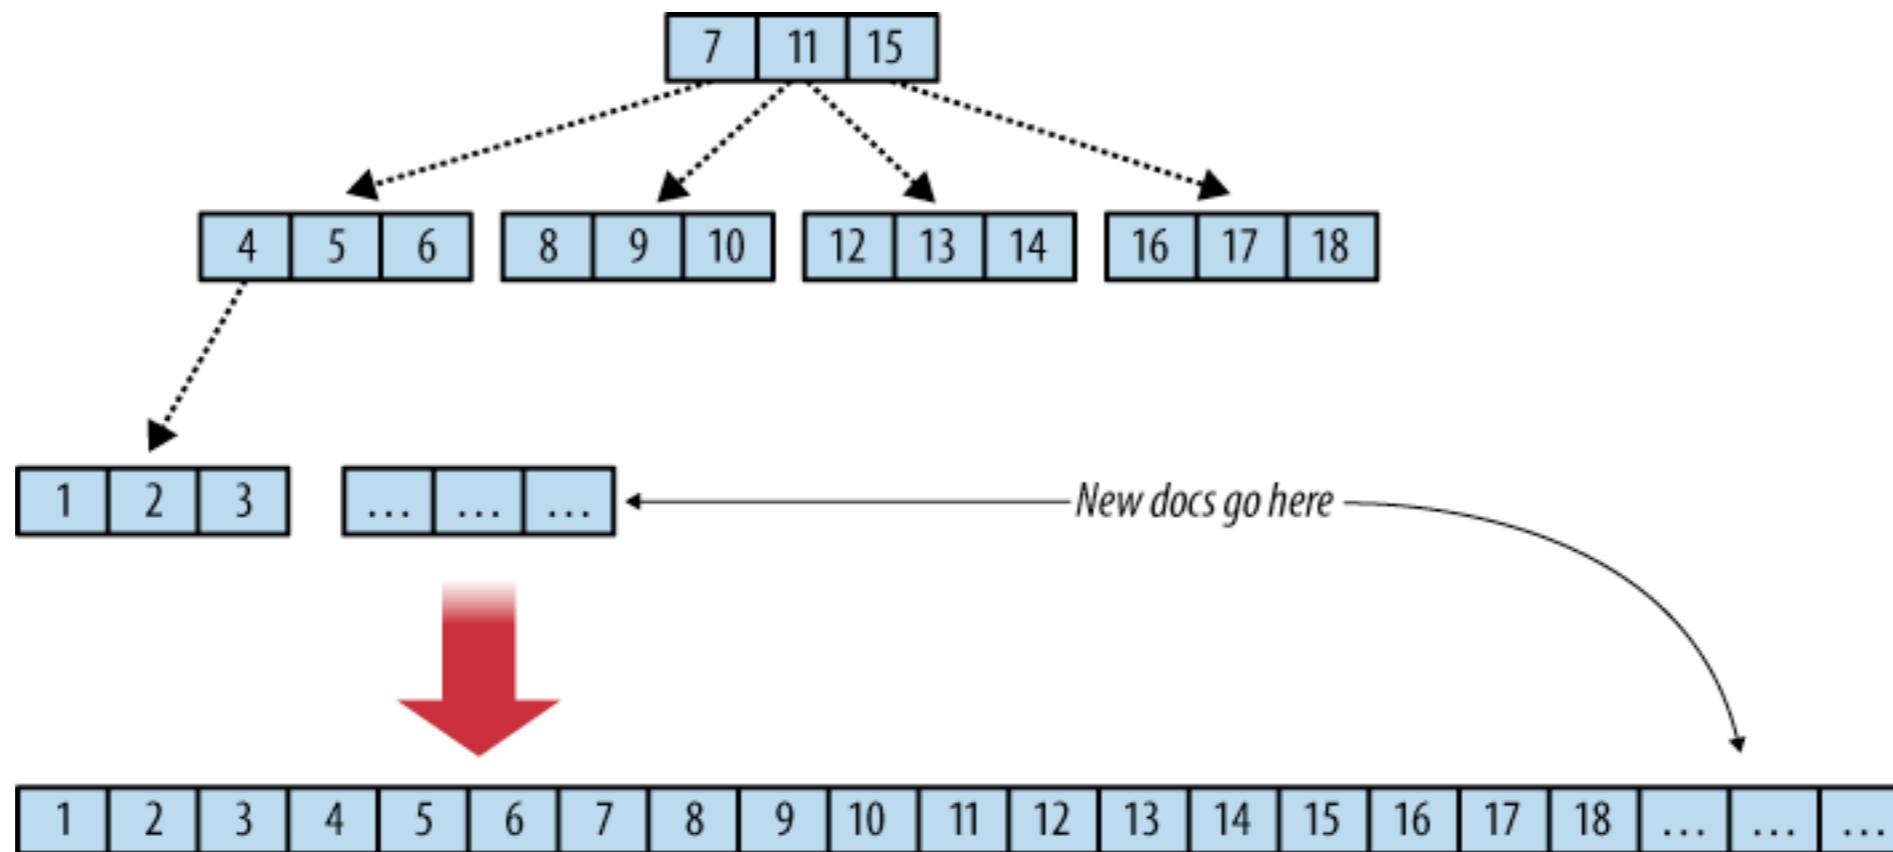
# Erlang



# Erlang

*Erlang's main strength is support for concurrency. It has a small but powerful set of primitives to create processes and communicate among them.  
(...) a **benchmark with 20 million processes** has been successfully performed.*

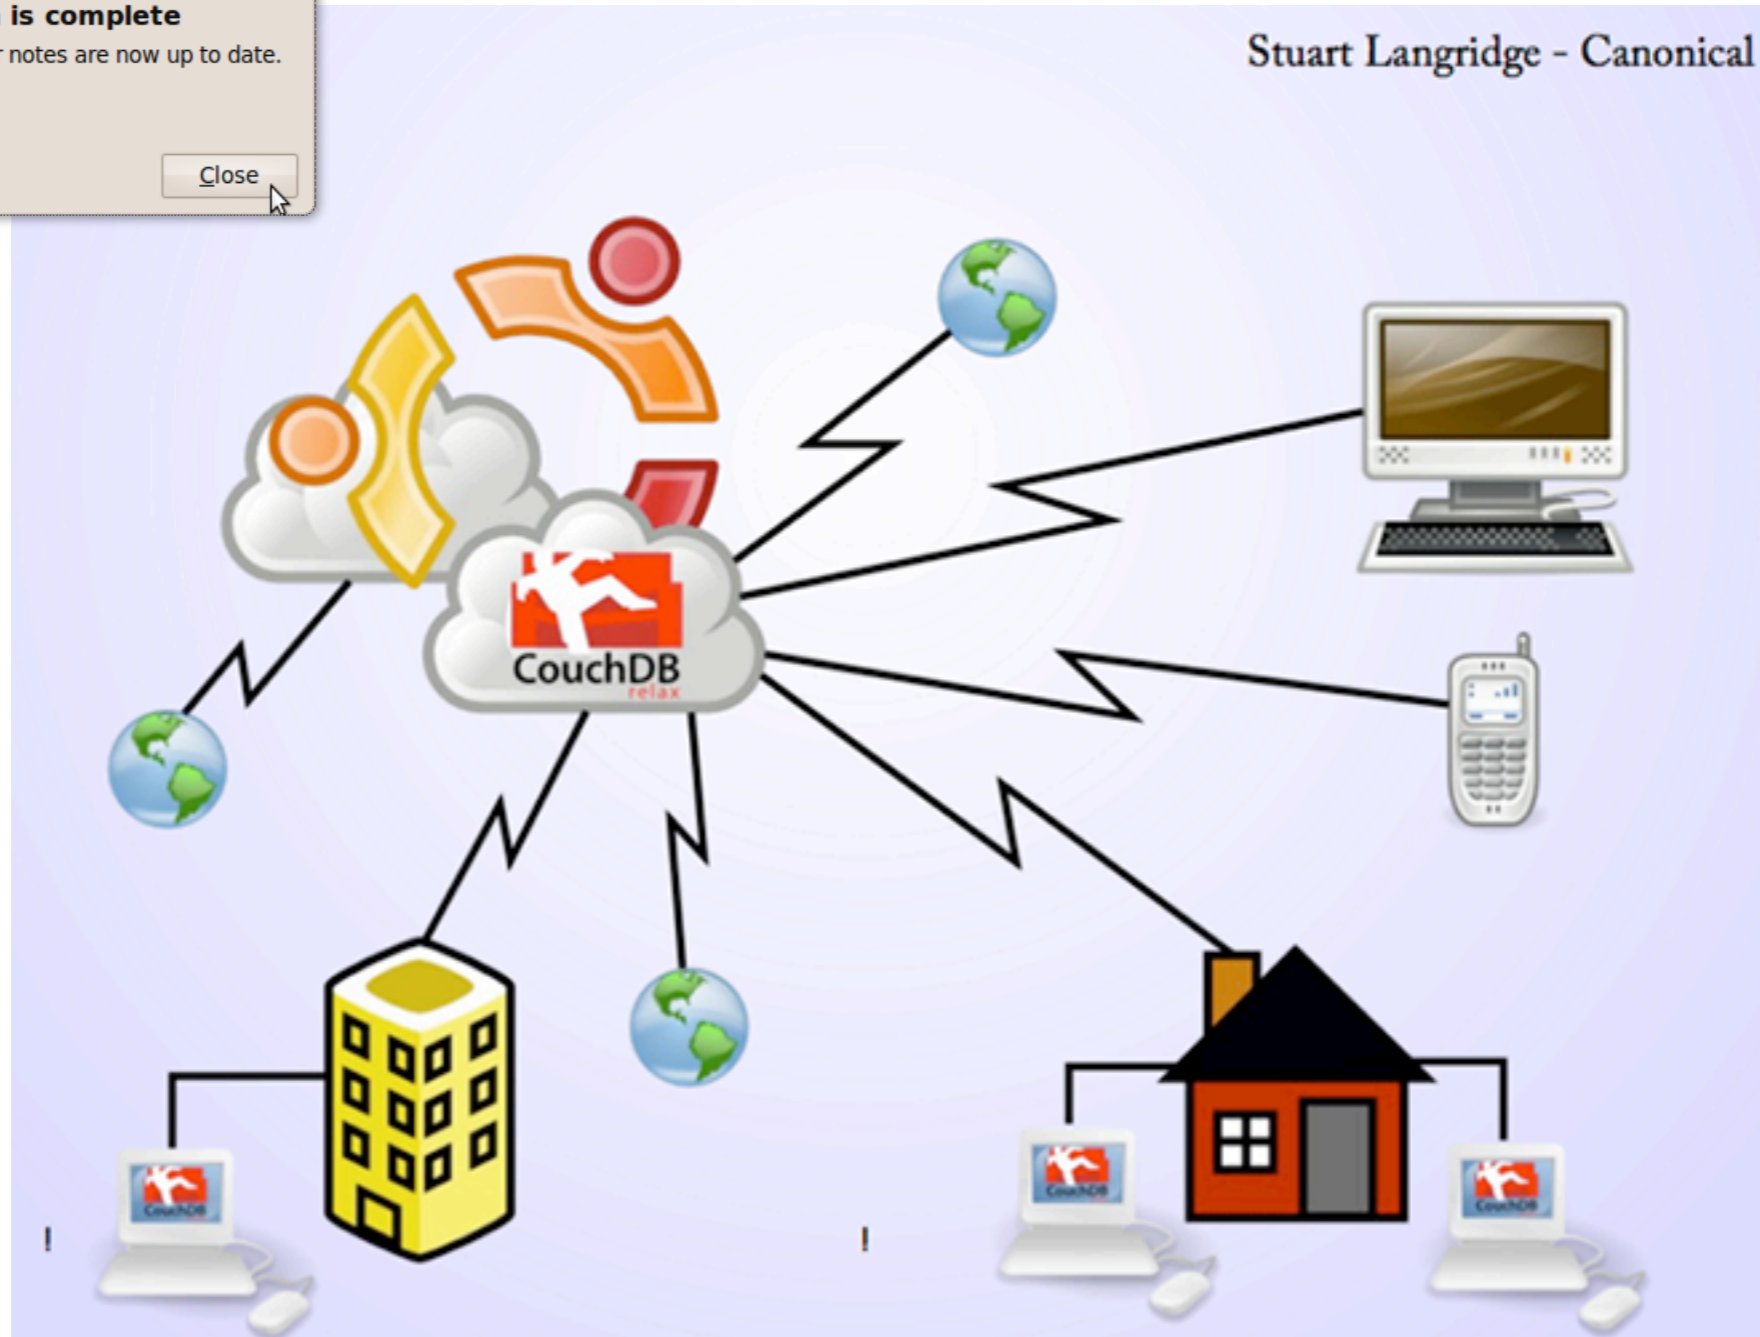
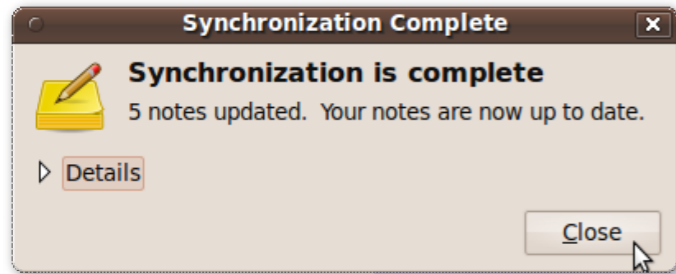
# „Append-Only B-Tree“





## Decentralizace

# Ubuntu One



# Replikace

The screenshot shows the Apache CouchDB Futon web interface for configuring a replicator. The browser window title is "Apache CouchDB - Futon: Replicator" and the address bar shows "http://localhost:5984/\_utils/replicator.html".

The main content area is titled "Replicator" and contains the following configuration options:

- Replicate changes from:**
  - Local database:
  - Remote database:
- to:**
  - Local database:
  - Remote database:

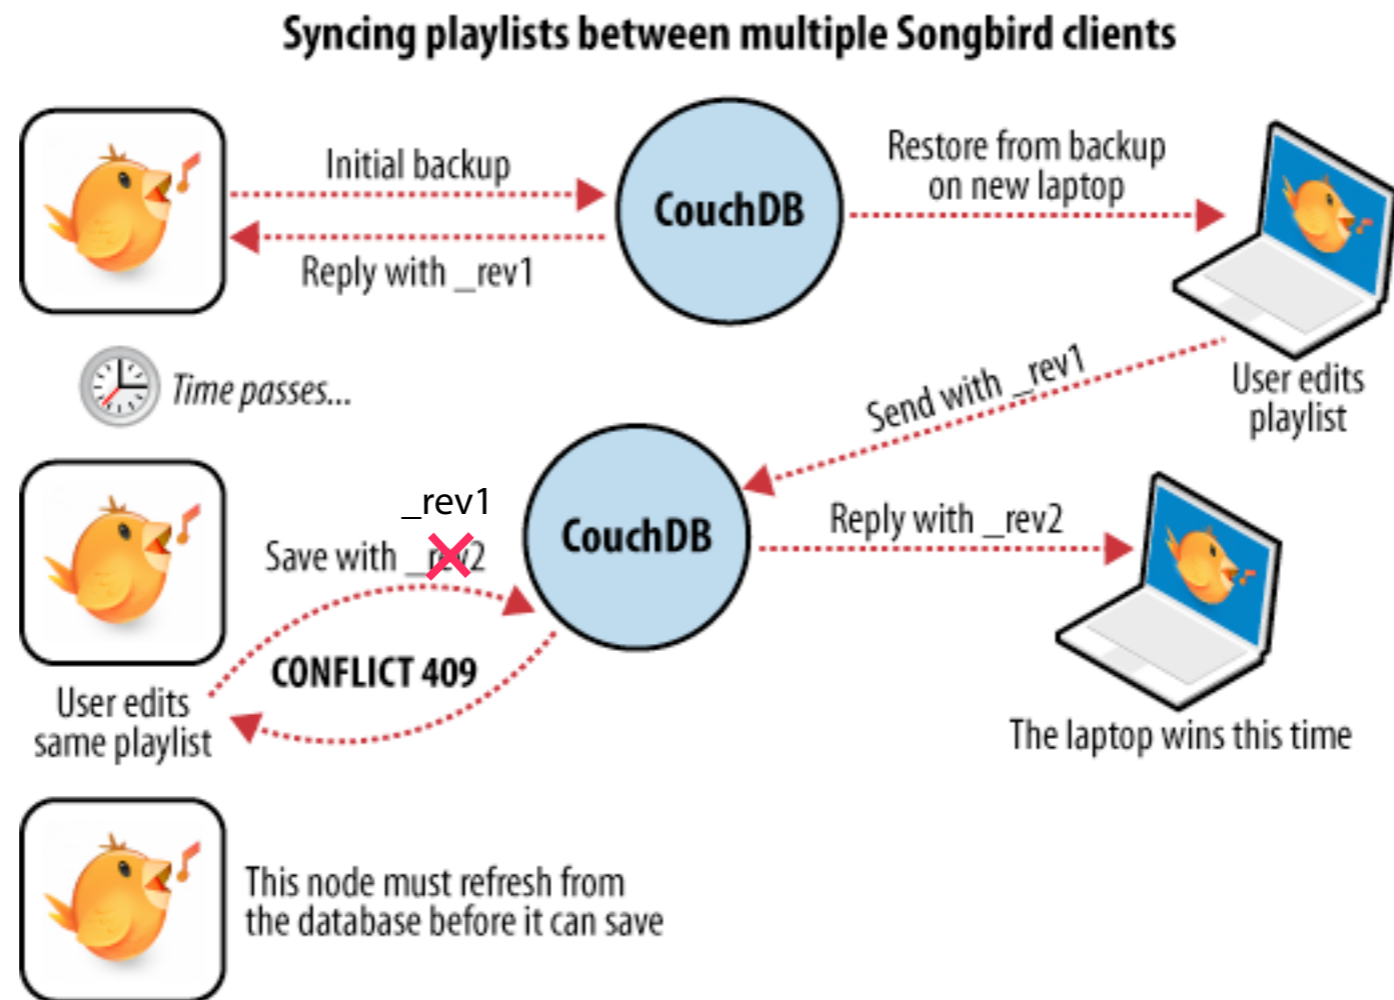
Below the configuration, there is a "Continuous" checkbox which is checked and circled in red, and a "Replicate" button. Below that is an "Event" section with the text "No replication".

On the right side, there is a sidebar with the CouchDB logo and the text "CouchDB relax". Below the logo, there is a "Tools" section with links for "Overview", "Configuration", and "Replicator". A message in the sidebar reads: "Welcome to Admin Party! Everyone is admin. Fix this". At the bottom of the sidebar, it says "Futon on Apache CouchDB 1.0.1".

The browser's address bar at the bottom shows "http://localhost:5984/\_utils/document.html?addressbook/isom-fritsch".



# Řešení konfliktů



# Jednoduchý multi-master cluster s Nginx

```
$ cat /opt/local/etc/couchdb/serverA.ini
```

```
...
```

```
port = 6001
```

```
$ cat /opt/local/etc/couchdb/serverB.ini
```

```
...
```

```
port = 6002
```

```
$ cat /opt/nginx/conf/nginx.conf
```

```
http {
    # ...
    upstream couchdb_cluster {
        server 127.0.0.1:6001;
        server 127.0.0.1:6002;
    }

    server {
        listen 80;
        server_name couchdb_cluster.local;

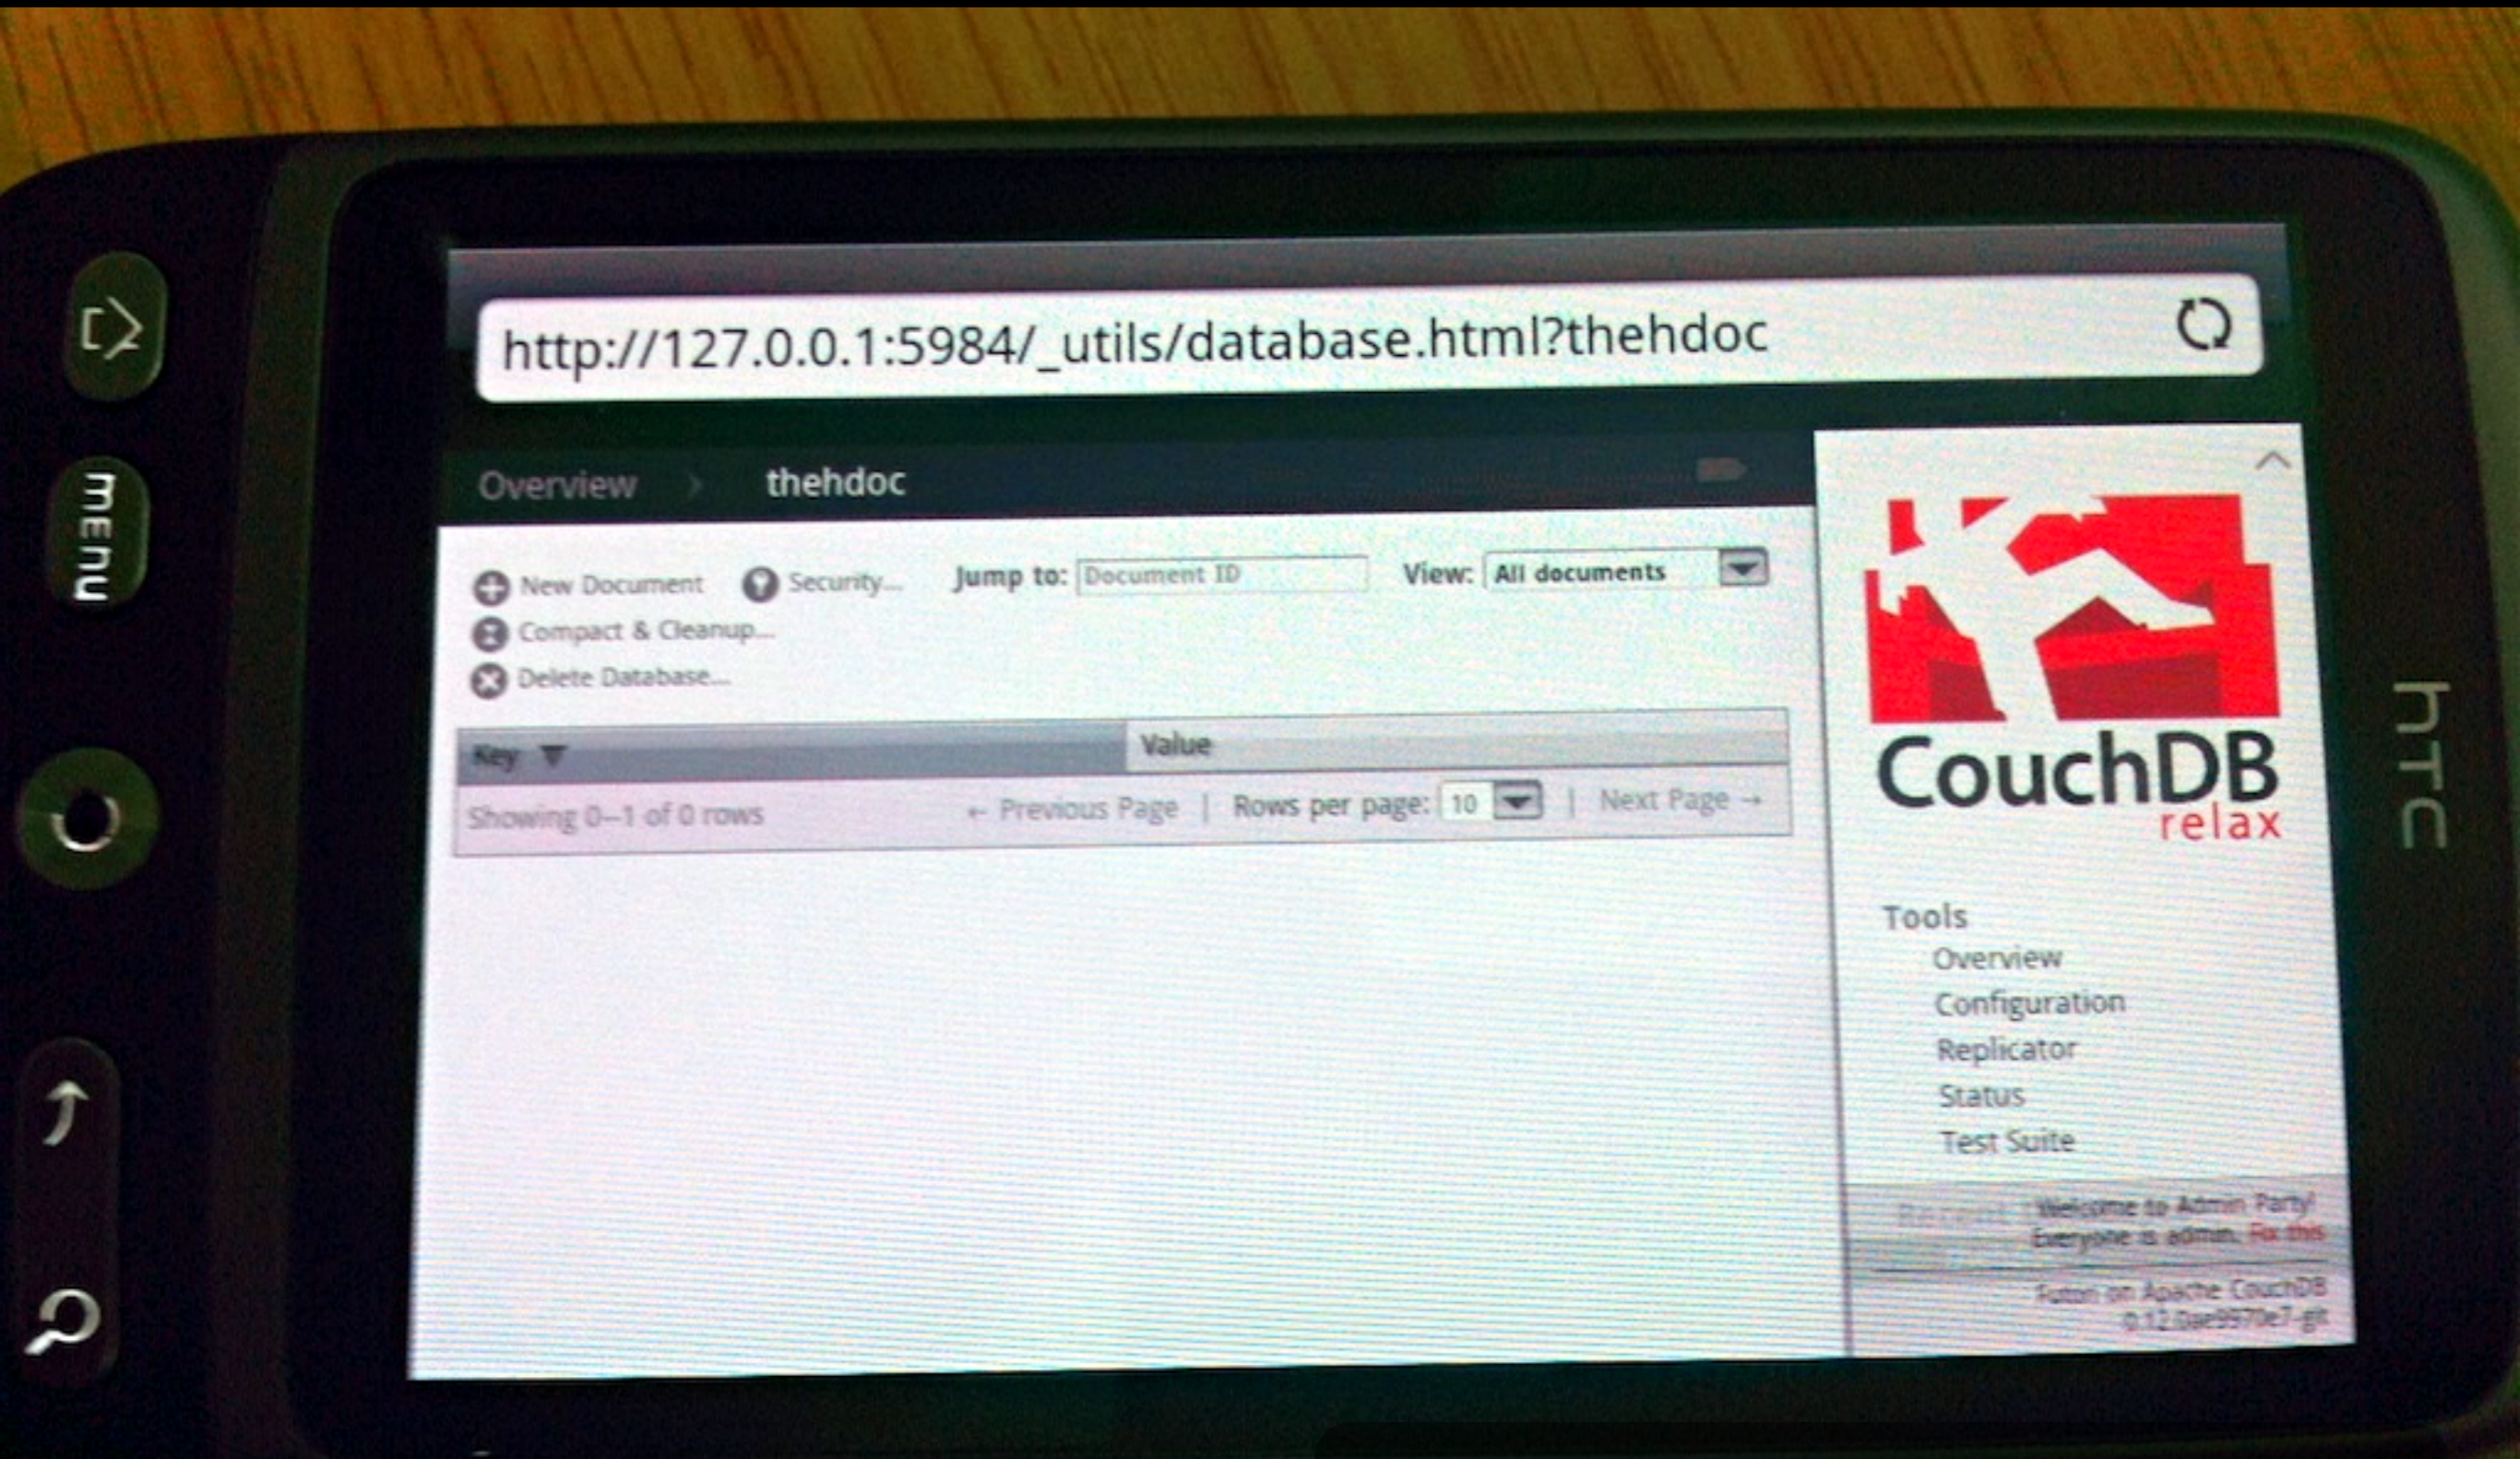
        location / {
            proxy_pass http://couchdb_cluster;
            break;
        }
    }
    # ...
}
```

```
$ curl -X POST http://localhost:6001/_replicate \
-d '{"source":"my_database", "target":"http://localhost:6002/my_database", "continuous":true}'
```

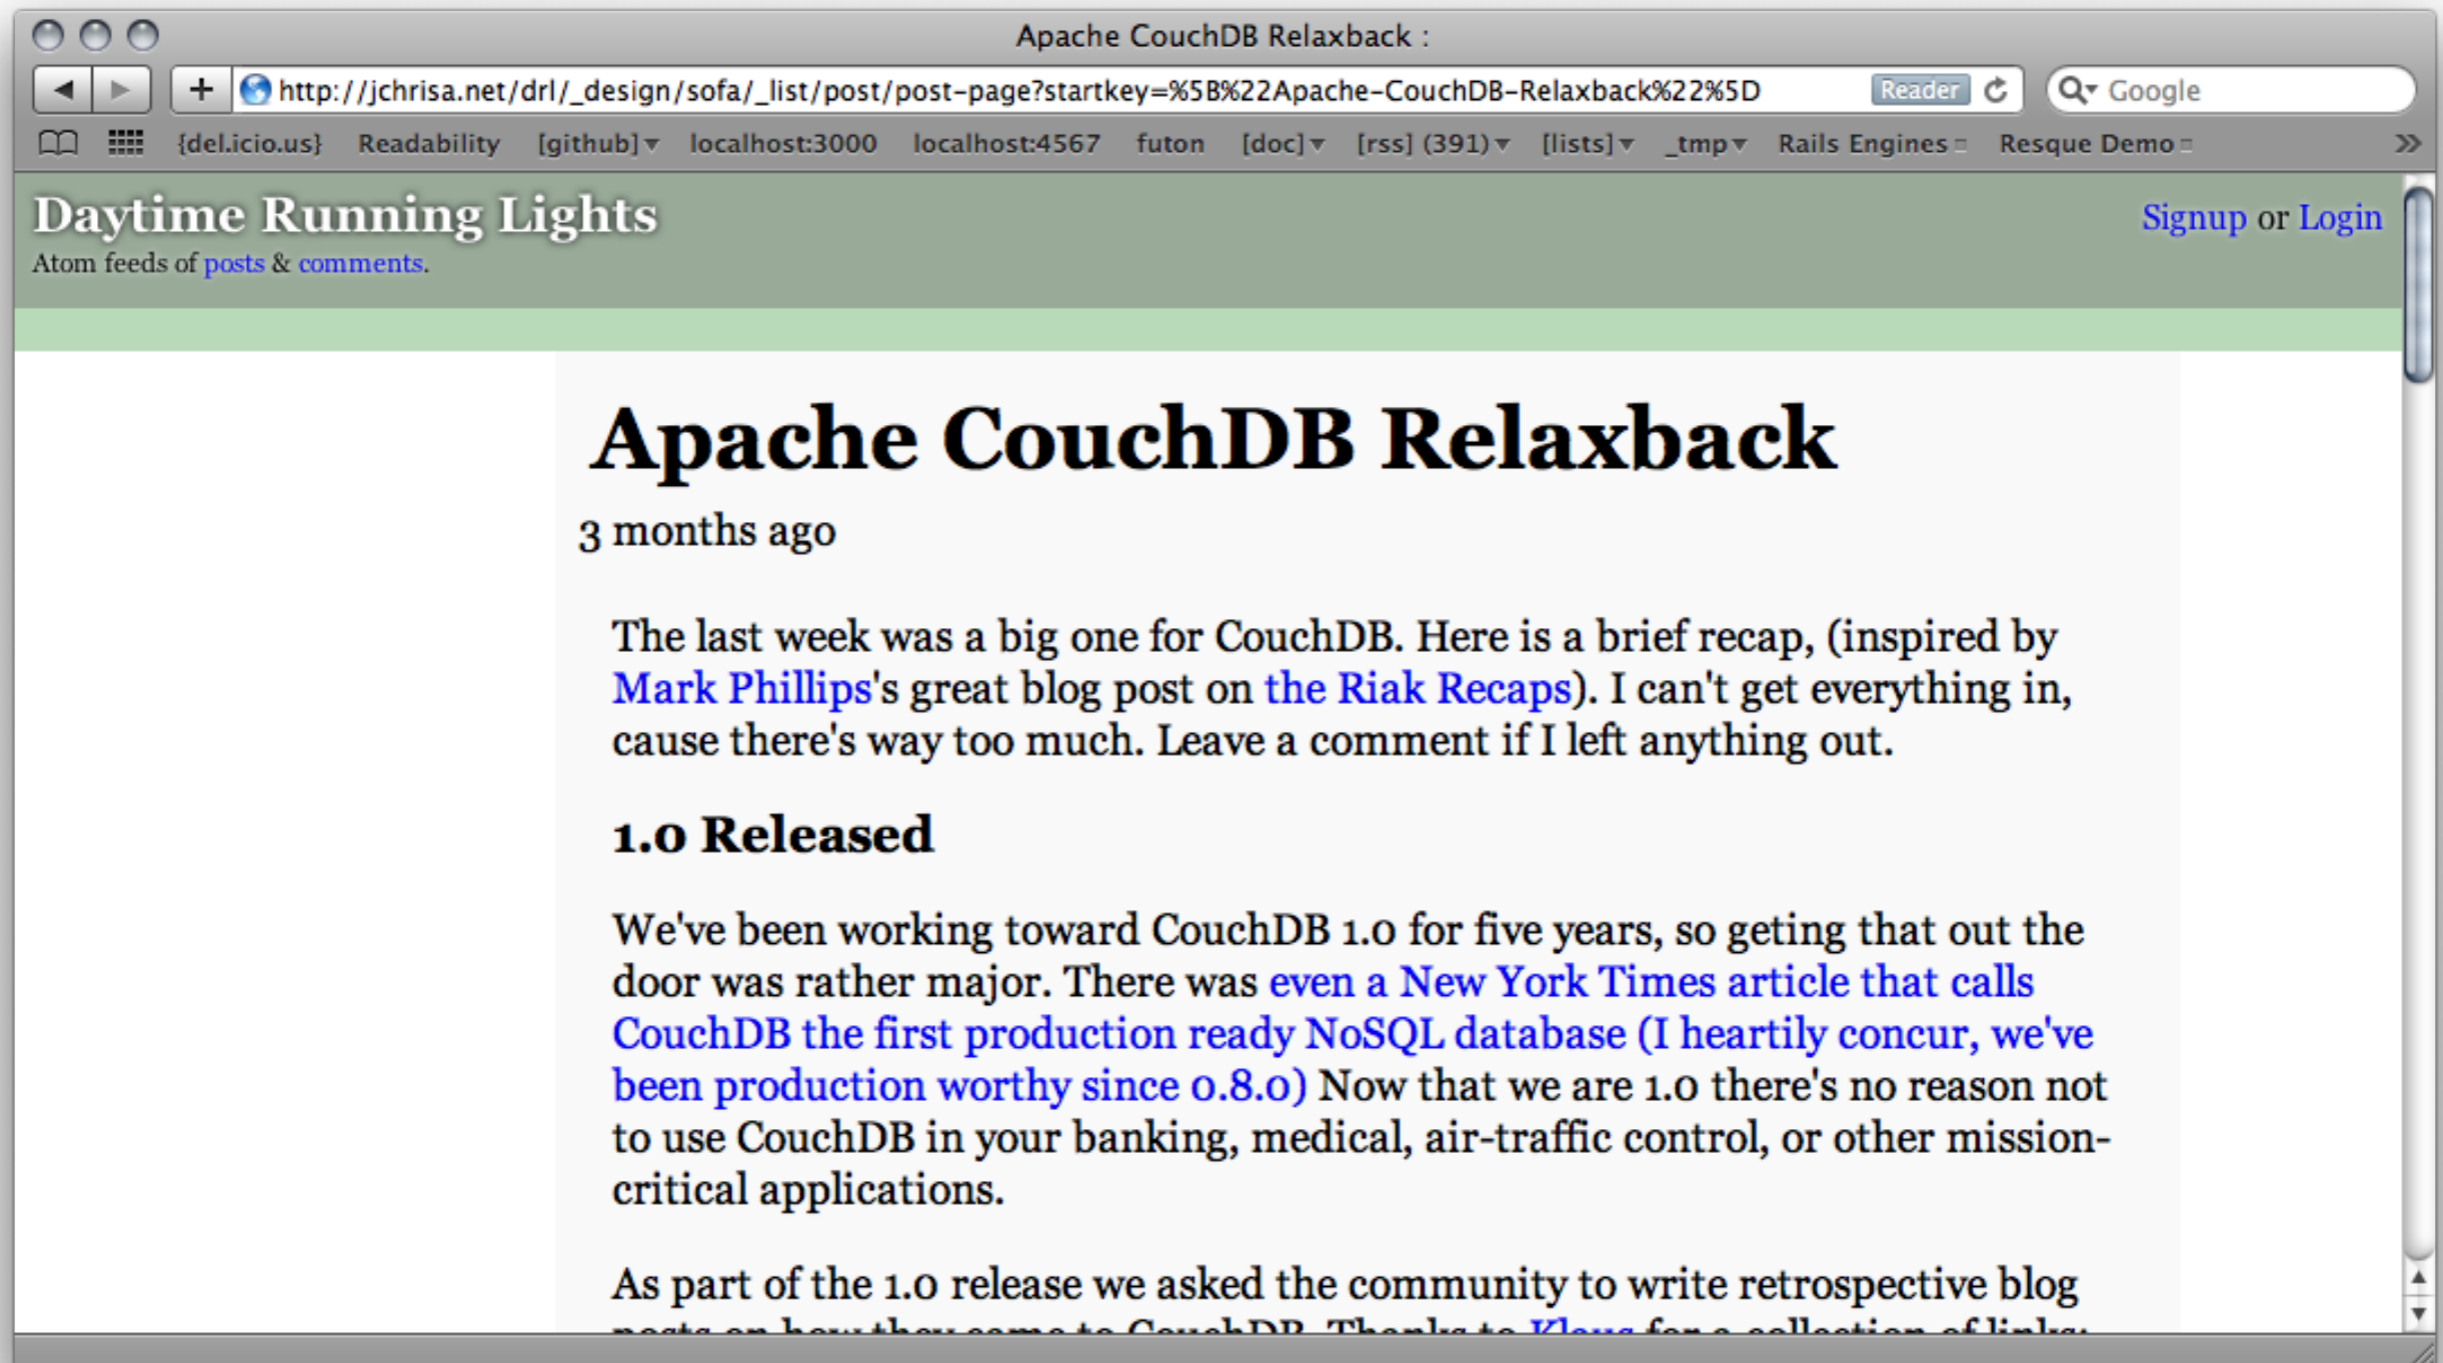
```
$ curl -X POST http://localhost:6002/_replicate \
-d '{"source":"my_database", "target":"http://localhost:6001/my_database", "continuous":true}'
```



# Škálování směrem dolů







# CouchApps



Kabul War Diary

http://127.0.0.1:5984/afgwardiary/\_design/afgwardiary/index.html

Reader Google

{del.icio.us} Readability [github] localhost:3000 localhost:4567 futon [doc] [rss] (391) [lists] \_tmp Rails Engines Resque Demo

## Kabul War Diary (unofficial)

Couchapp by benoitc

### Look the Timeline

Look the content in a timeline with each reports displayed on the map.

Sunday, July 26 5pm EST.

WikiLeaks today released over 75,000 secret US military reports covering the war in Afghanistan.

The Afghan War Diary an extraordinary secret compendium of over 91,000 reports covering the war in Afghanistan from 2004 to 2010. The reports describe the majority of lethal military actions involving the United States military. They include the number of persons internally stated to be killed, wounded, or detained during each action, together with the precise geographical location of each event, and the military units involved and major weapon systems used.

The Afghan War Diary is the most significant archive about the reality of war to have ever been released during the course of a war. The deaths of tens of thousands is normally only a statistic but the archive reveals the locations and the key events behind each most of these deaths. We hope its release will lead to a comprehensive understanding of the war in Afghanistan and provide the raw ingredients necessary to change its course.

**Timeline**

**Browse by Type**

**Browse by Category**



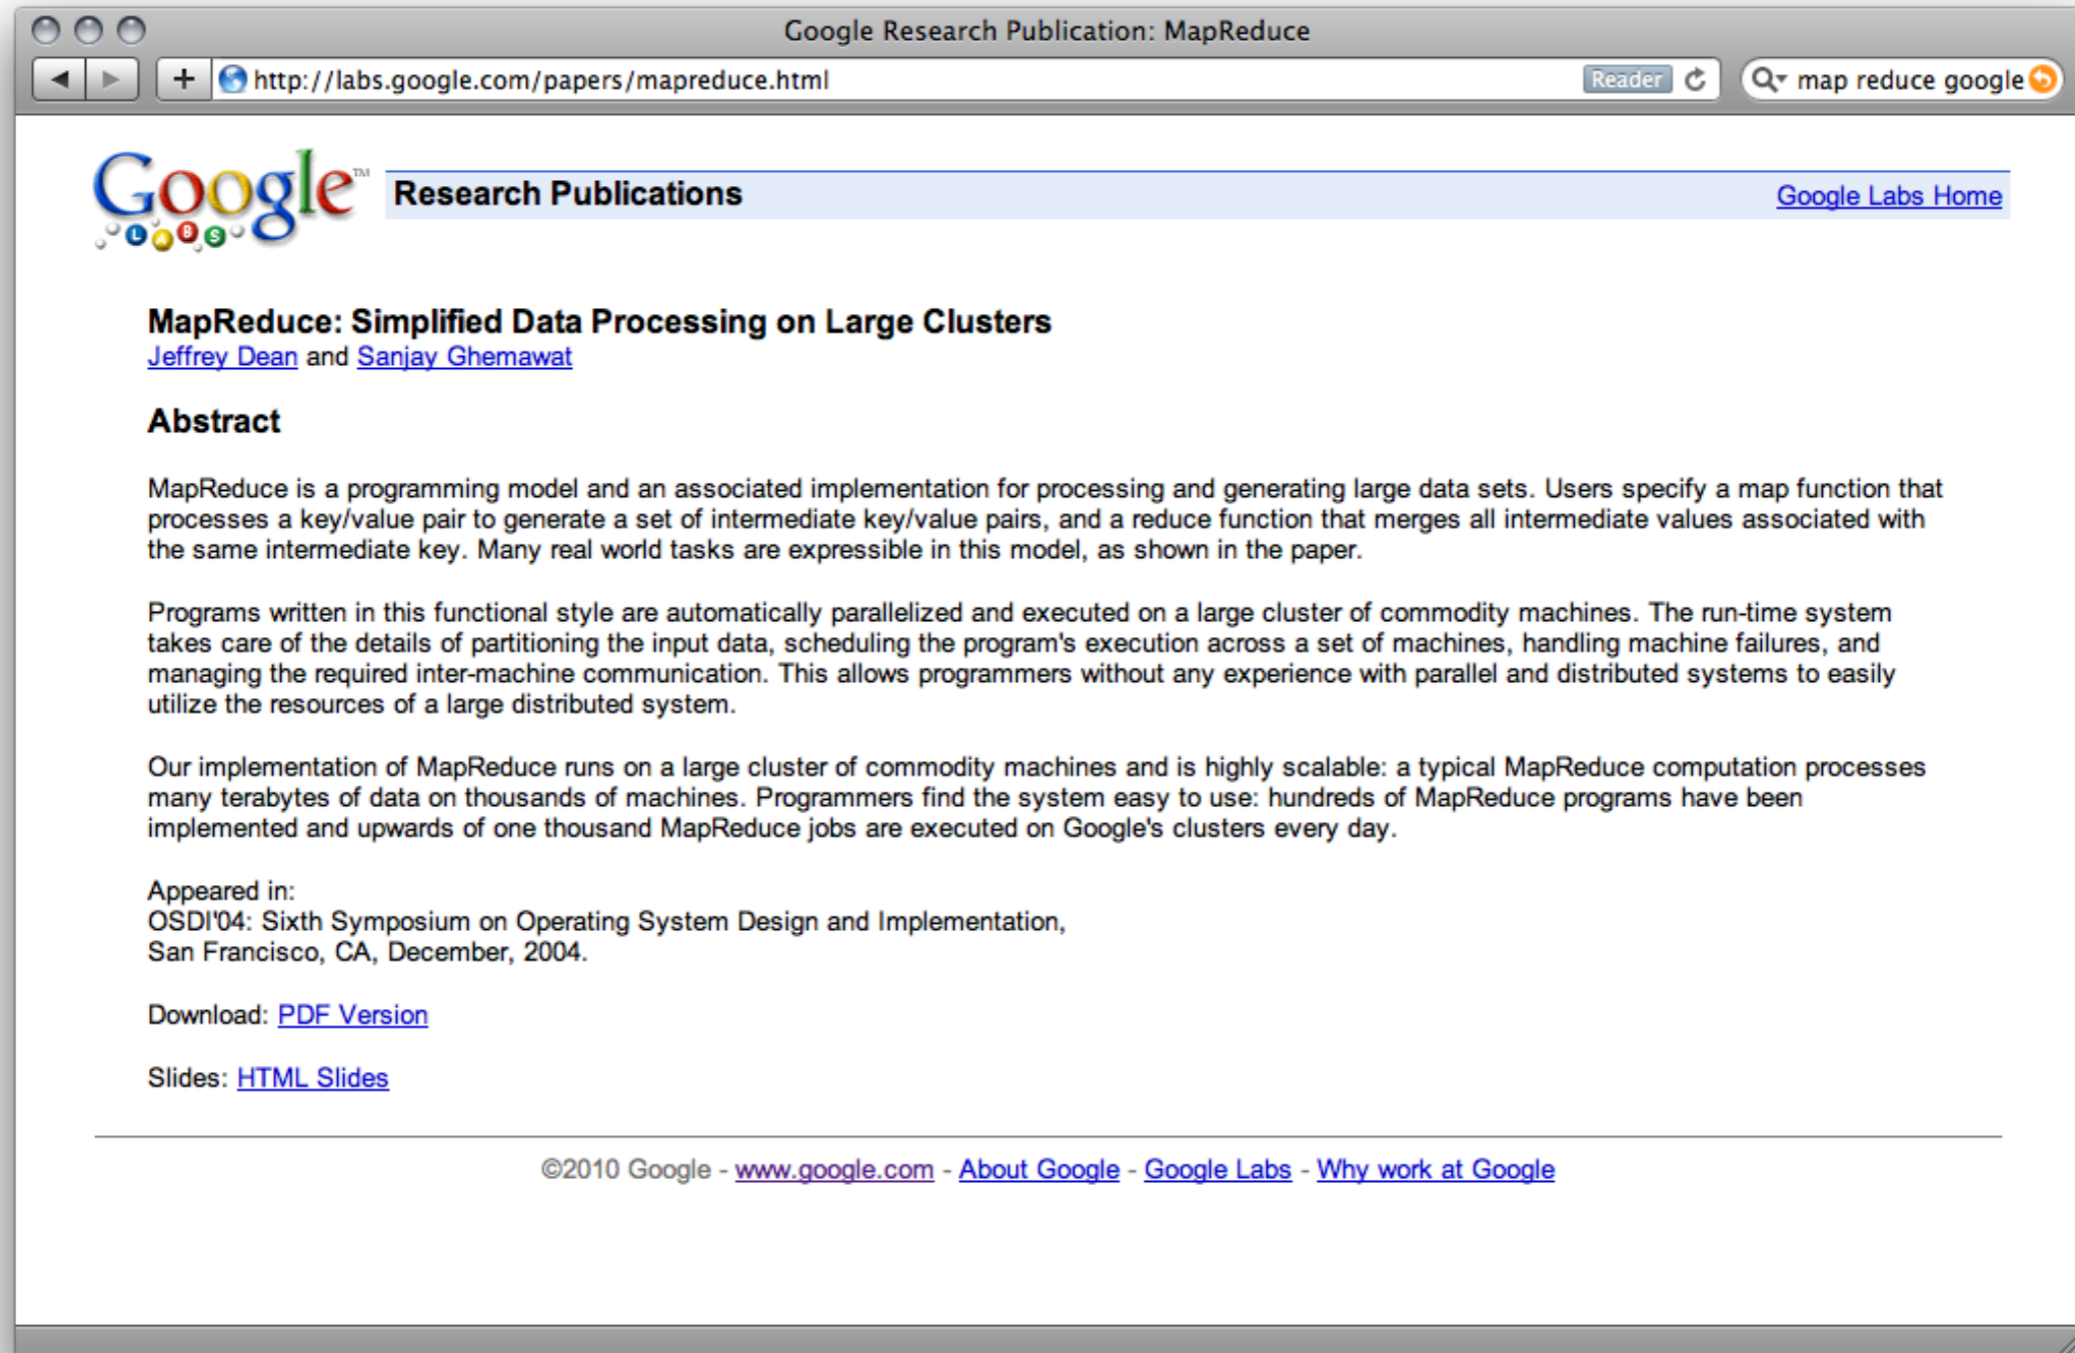
# CouchApps

The screenshot shows a web browser window displaying a GitHub repository page for 'benoitc/afgwardiary'. The browser's address bar shows the URL 'git http://github.com/benoitc/afgwardiary'. The page header includes the GitHub logo and the user 'karmi' with navigation links like 'Dashboard', 'Inbox', 'Account Settings', and 'Log Out'. Below the repository name, there are buttons for 'Watch', 'Fork', and 'Download Source', along with statistics for 12 watchers and 1 fork. The 'Source' tab is active, showing a commit message: 'couchapp to render afgwardiary data from wikileaks'. Below the message, there are links for 'Read more' and a URL 'http://pollen.nymphormation.org/afgwar/\_design/afgwardiary/index.html'. A section for cloning the repository is visible, with 'Git Read-Only' selected and the URL 'http://github.com/benoitc/afgwardiary.git'. A commit summary for 'link the timeline' by 'benoitc' (author) on August 04, 2010, is shown with commit hash '0be2094f00e8fcbcd2d4', tree '42a416b8bf81863cdd8b', and parent 'cfc26685d5edf1b8d8e7'. At the bottom, a commit history table is displayed.

name	age	message	history
<code>.couchappignore</code>	July 27, 2010	initial release [benoitc]	

# 5

## Dotazování s Map/Reduce



The screenshot shows a web browser window with the title "Google Research Publication: MapReduce". The address bar contains the URL "http://labs.google.com/papers/mapreduce.html". The search bar shows "map reduce google". The page content includes the Google Labs logo, the title "MapReduce: Simplified Data Processing on Large Clusters" by Jeffrey Dean and Sanjay Ghemawat, an abstract, and a footer with copyright information.

Google Research Publications [Google Labs Home](#)

## MapReduce: Simplified Data Processing on Large Clusters

[Jeffrey Dean](#) and [Sanjay Ghemawat](#)

### Abstract

MapReduce is a programming model and an associated implementation for processing and generating large data sets. Users specify a map function that processes a key/value pair to generate a set of intermediate key/value pairs, and a reduce function that merges all intermediate values associated with the same intermediate key. Many real world tasks are expressible in this model, as shown in the paper.

Programs written in this functional style are automatically parallelized and executed on a large cluster of commodity machines. The run-time system takes care of the details of partitioning the input data, scheduling the program's execution across a set of machines, handling machine failures, and managing the required inter-machine communication. This allows programmers without any experience with parallel and distributed systems to easily utilize the resources of a large distributed system.

Our implementation of MapReduce runs on a large cluster of commodity machines and is highly scalable: a typical MapReduce computation processes many terabytes of data on thousands of machines. Programmers find the system easy to use: hundreds of MapReduce programs have been implemented and upwards of one thousand MapReduce jobs are executed on Google's clusters every day.

Appeared in:  
OSDI'04: Sixth Symposium on Operating System Design and Implementation,  
San Francisco, CA, December, 2004.

Download: [PDF Version](#)

Slides: [HTML Slides](#)

---

©2010 Google - [www.google.com](#) - [About Google](#) - [Google Labs](#) - [Why work at Google](#)

# Koncept

```
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
```

# Nejjednodušší *view*

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



# Nejjednodušší view

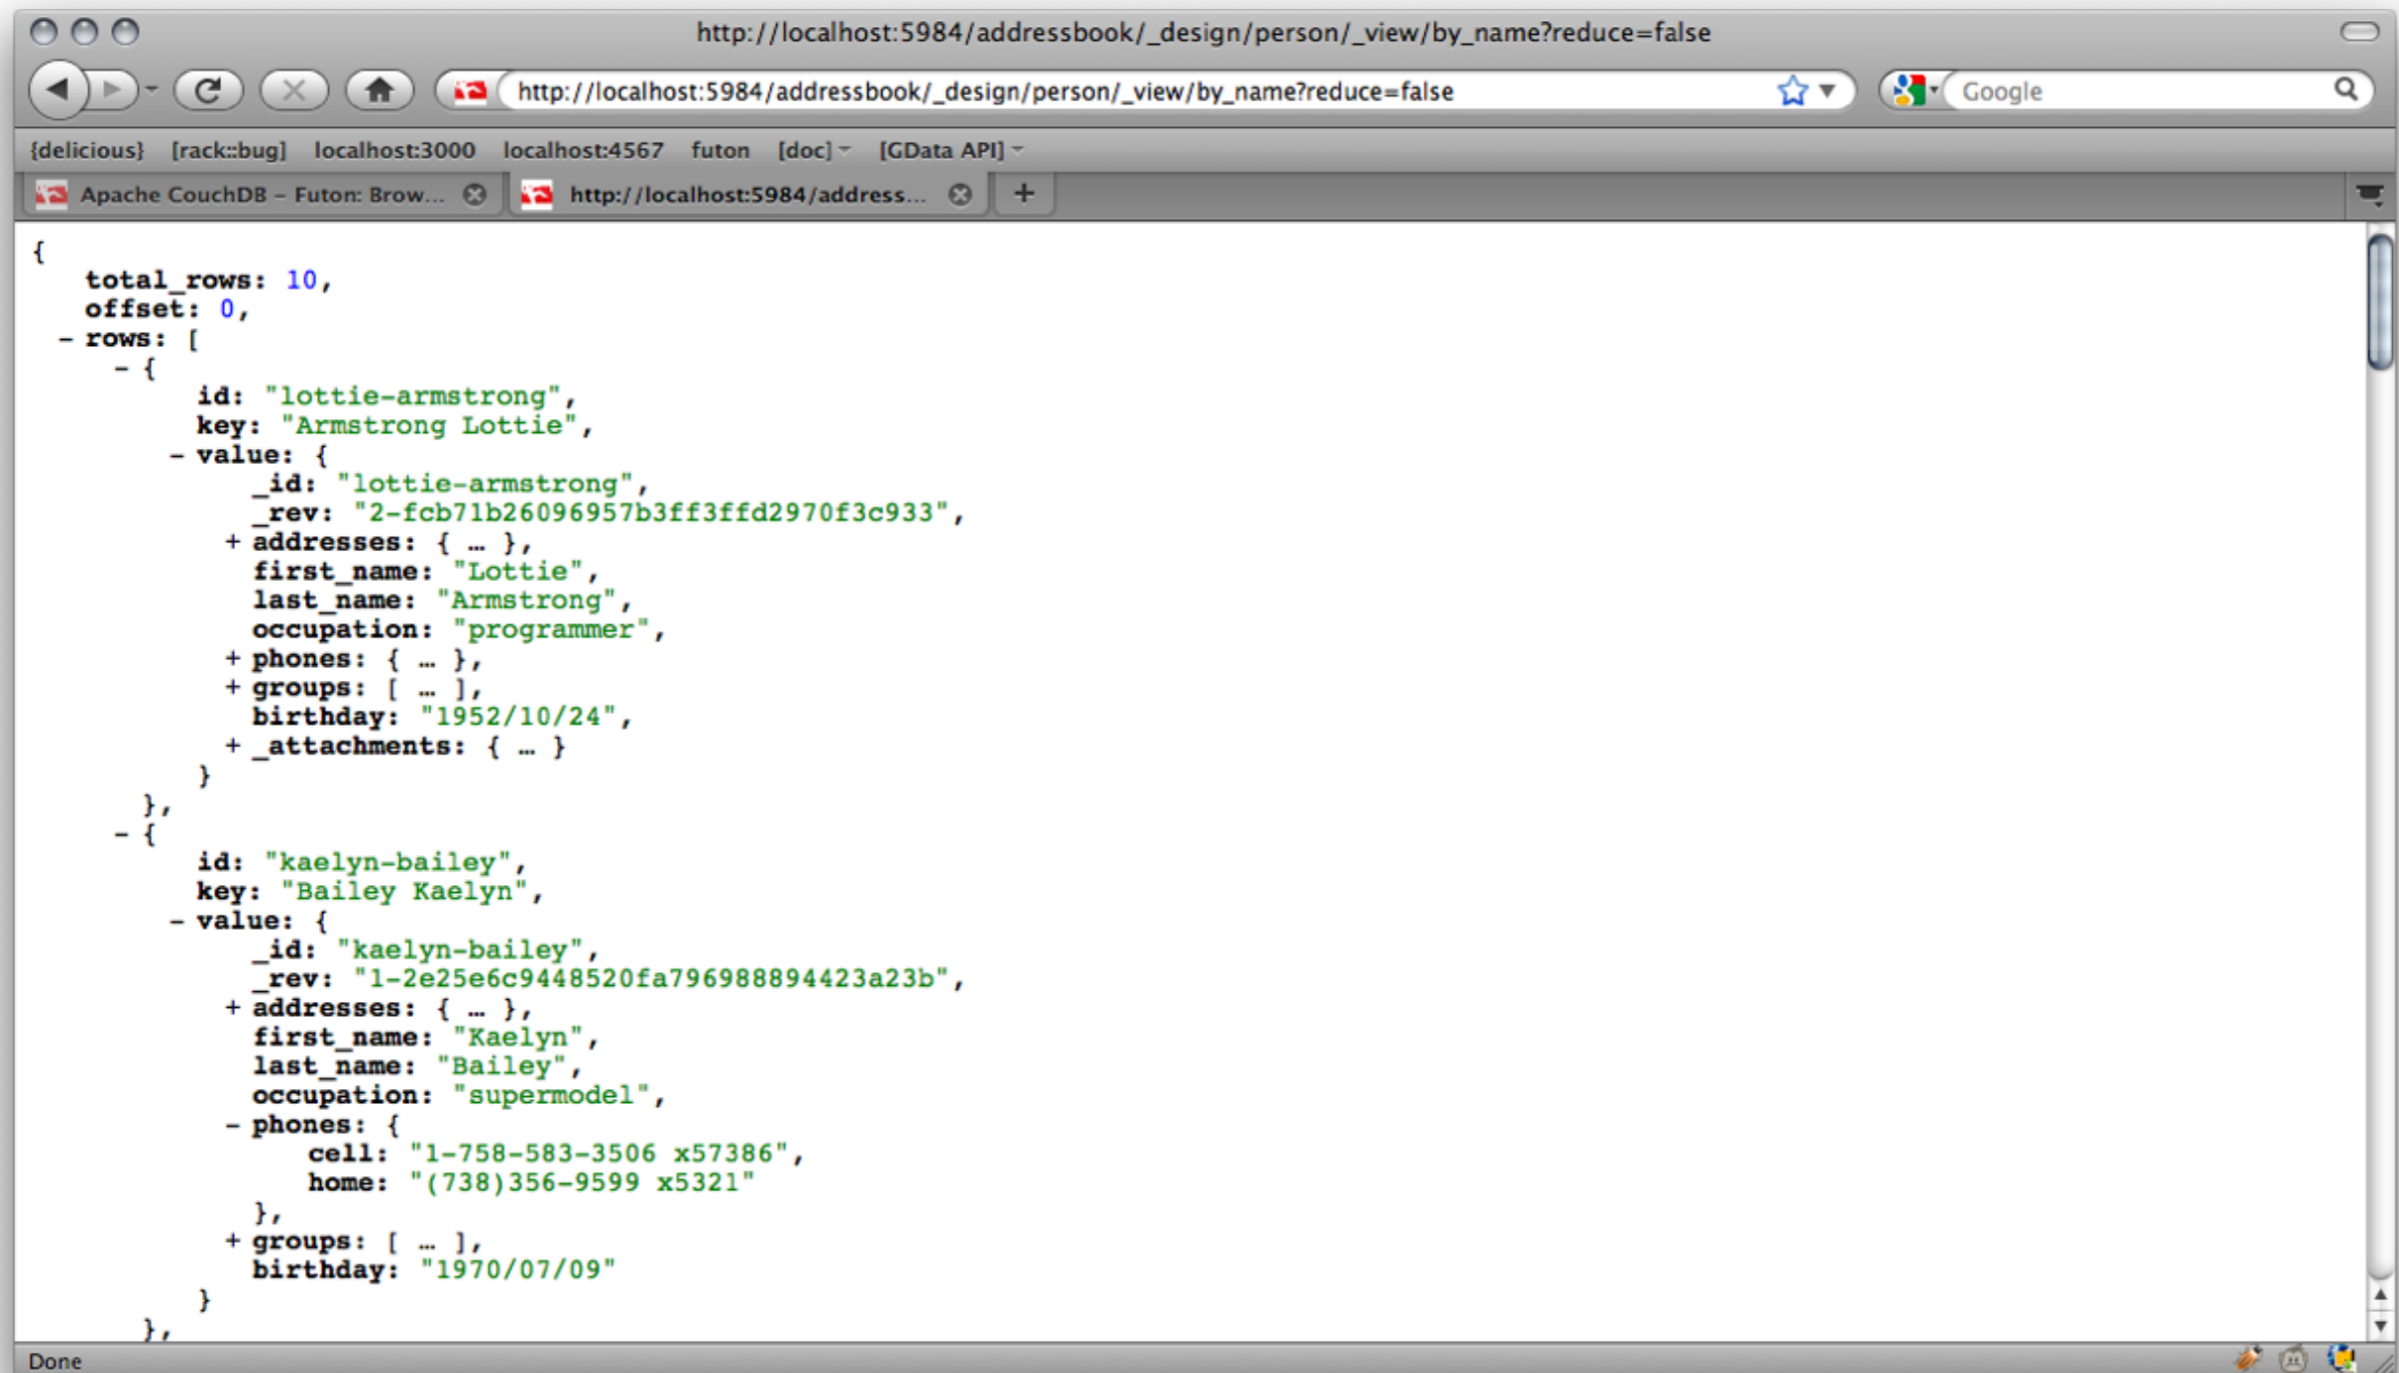
```
function(doc) {  
  if (doc.last_name && doc.first_name) {  
    emit( doc.last_name + ' ' + doc.first_name, doc )  
  }  
}
```

**INPUT**

**KEY**                      **VALUE**

**OUTPUT**

# Výsledek: dokumenty seřazené podle příjmení



```
{
  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",
        - phones: {
          cell: "1-758-583-3506 x57386",
          home: "(738)356-9599 x5321"
        },
        + groups: [ ... ],
        birthday: "1970/07/09"
      }
    }
  ],
}
```

# Ještě jednodušší *view*

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

# Výsledek?

Key ▼	Grouping: exact ▼	Value	<input type="checkbox"/> Reduce
"supermodel" ID: kaelyn-bailey		1	
"supermodel" ID: jeramie-feest		1	
"supermodel" ID: edd-mosciski		1	
"supermodel" ID: desmond-stokes		1	
"supermodel" ID: bertram-bogan		1	
"supermodel" ID: andreanne-jakubowski		1	
"programmer" ID: lottie-armstrong		1	
"designer" ID: khalid-thiel		1	
"designer" ID: jo-lubowitz		1	
"designer" ID: dorothy-von		1	

Showing 1-10 of 10 rows      ← Previous Page | Rows per page: 50 ▼ | Next Page →

# Jednoduchý *reduce*

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



# Výsledek!

Key ▼	Grouping: exact ▼	Value	<input checked="" type="checkbox"/> Reduce
"supermodel"		6	
"programmer"		1	
"designer"		3	

Showing 1-3 of unknown rows

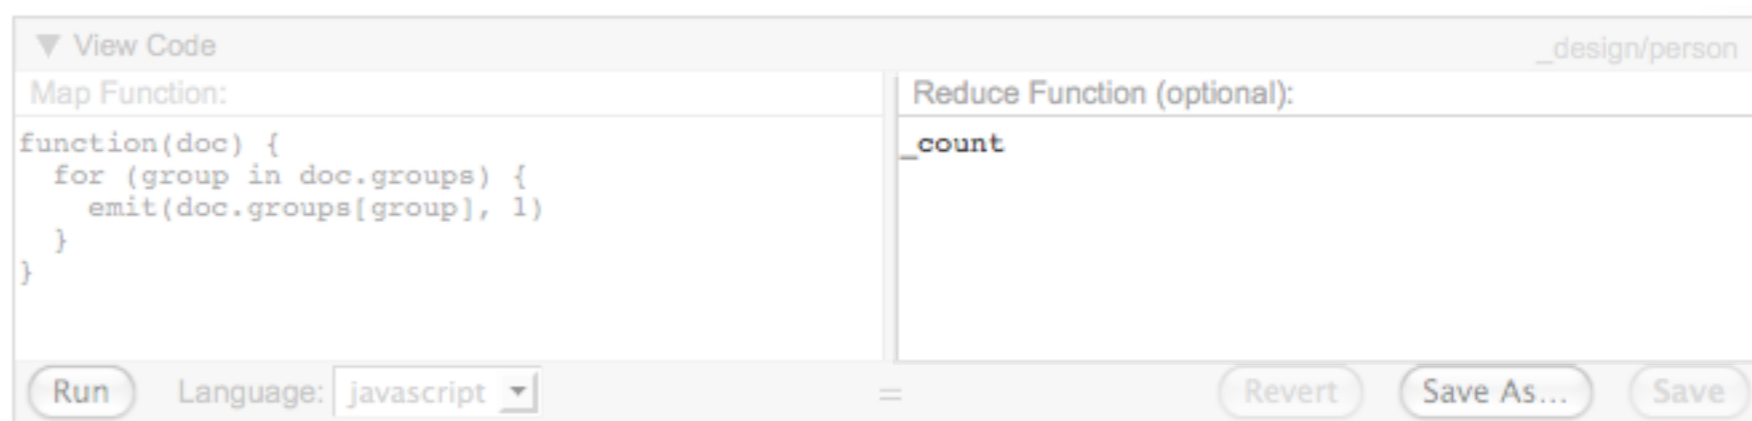
← Previous Page | Rows per page: 50 ▼ | Next Page →

# Zabudované Erlang funkce

```
$ couchdb
```

```
Apache CouchDB has started. Time to relax.
```

**\_count**  
**\_sum**  
**\_stats**



The screenshot shows the CouchDB interface for editing a view. The title bar indicates the view is named `_design/person`. The interface is split into two main sections: "Map Function" and "Reduce Function (optional)".

**Map Function:**

```
function(doc) {  
  for (group in doc.groups) {  
    emit(doc.groups[group], 1)  
  }  
}
```

**Reduce Function (optional):**

```
_count
```

At the bottom of the interface, there are several controls: a "Run" button, a "Language:" dropdown menu set to "javascript", an equals sign "=", a "Revert" button, a "Save As..." button, and a "Save" button.

# Map/Reduce pro počítání „tagů a podobných věcí“

```
function(doc) {  
  for (group in doc.groups) {  
    emit(doc.groups[group], 1)  
  }  
}
```

`_count`

# Výsledek *map* fáze

Key ▾	Grouping: <input type="text" value="exact"/> ▾	Value	<input type="checkbox"/> Reduce
"work" ID: lottie-armstrong		1	
"work" ID: khalid-thiel		1	
"work" ID: kaelyn-bailey		1	
"work" ID: jeramie-feest		1	
"work" ID: desmond-stokes		1	
"friends" ID: khalid-thiel		1	
"friends" ID: kaelyn-bailey		1	
"friends" ID: jo-lubowitz		1	
"friends" ID: dorothy-von		1	
"friends" ID: andreanne-jakubowski		1	
"family" ID: kaelyn-bailey		1	
"family" ID: jo-lubowitz		1	
"family" ID: edd-mosciski		1	
"family" ID: desmond-stokes		1	
"family" ID: bertram-bogan		1	

Showing 1-15 of 15 rows      ← Previous Page | Rows per page:  ▾ | Next Page →

# Výsledek *reduce* fáze

Key ▼	Grouping: exact ▼	Value	<input checked="" type="checkbox"/> Reduce
"work"		5	
"friends"		5	
"family"		5	

Showing 1-3 of unknown rows      ← Previous Page | Rows per page: 50 ▼ | Next Page →



# Seskupování klíčů (*group levels*)

```
function(doc) {  
  var date = new Date(doc.birthday)  
  emit([date.getFullYear(), date.getMonth()+1, date.getDate()], 1 )  
}
```

COMPOSITE KEY (ARRAY)

\_count

# Výchozí seskupení („přesné“)

Key ▼	Grouping: exact ▼	Value	<input checked="" type="checkbox"/> Reduce
[1990, 5, 9]		1	
[1990, 5, 3]		1	
[1988, 11, 20]		1	
[1986, 11, 9]		1	
[1983, 5, 4]		1	
[1983, 2, 27]		1	
[1979, 11, 21]		1	
[1978, 5, 29]		1	
[1975, 11, 24]		1	
[1971, 5, 5]		1	

Showing 1-10 of unknown rows      ← Previous Page | Rows per page: 50 ▼ | Next Page →

# Seskupení podle prvních dvou položek klíče

Key ▼	Grouping: level 2 ▼	Value	<input checked="" type="checkbox"/> Reduce
[1990, 5]		2	
[1988, 11]		1	
[1986, 11]		1	
[1983, 5]		1	
[1983, 2]		1	
[1979, 11]		1	
[1978, 5]		1	
[1975, 11]		1	
[1971, 5]		1	

Showing 1-9 of unknown rows

← Previous Page | Rows per page: 50 ▼ | Next Page →

# Seskupení podle první položky klíče

Key ▾	Grouping: level 1 ▾	Value	<input checked="" type="checkbox"/> Reduce
[1990]		2	
[1988]		1	
[1986]		1	
[1983]		2	
[1979]		1	
[1978]		1	
[1975]		1	
[1971]		1	

Showing 1-8 of unknown rows      ← Previous Page | Rows per page: 50 ▾ | Next Page →

# Parametry pro dotazování indexu

key

startkey

startkey\_docid

endkey

endkey\_docid

limit

stale

descending

skip

group

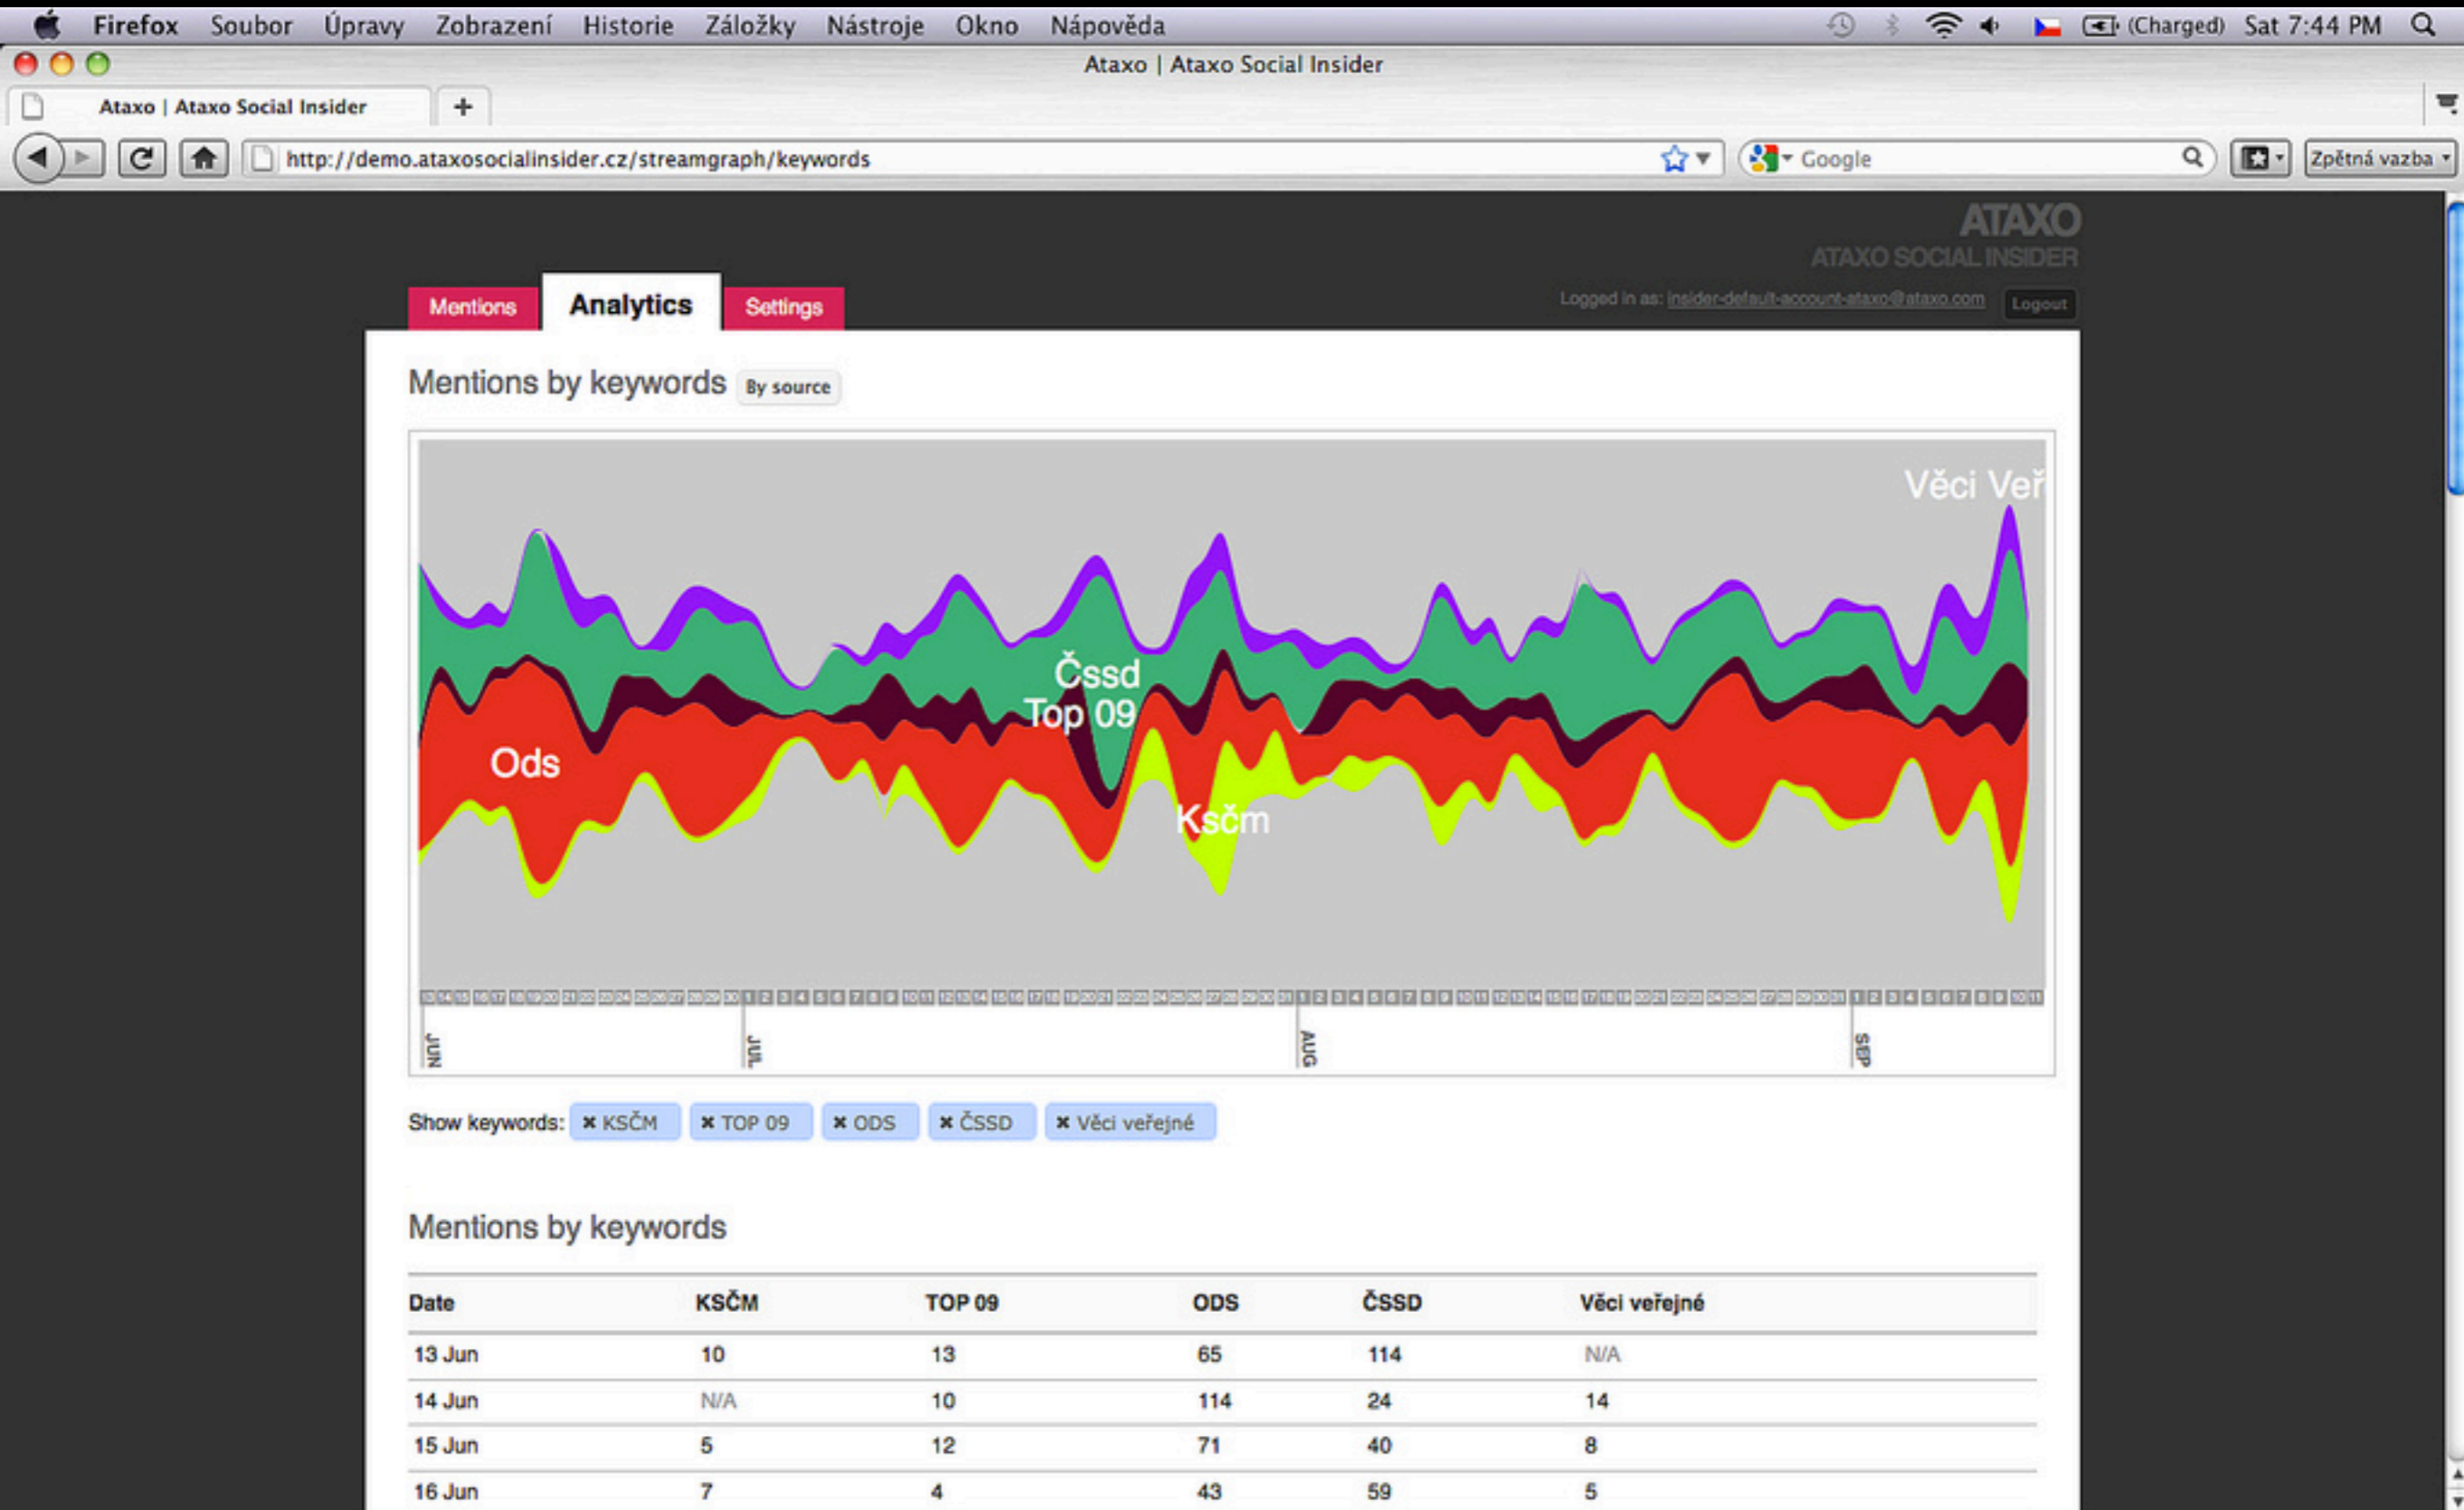
group\_level

reduce

include\_docs



# Komplexní Map/Reduce operace



# Komplexní Map/Reduce operace

```
SELECT
  COUNT(*) AS count,
  DATE_FORMAT(published_at, "%Y/%m/%d") AS date,
  keywords.value AS keyword
FROM feed_entries
  INNER JOIN feeds ON feed_entries.feed_id = feeds.id
  INNER JOIN keywords ON feeds.keyword_id = keywords.id
WHERE DATE_SUB(CURDATE(), INTERVAL 90 DAY) <= feed_entries.published_at
GROUP BY date, keyword
ORDER BY date, keyword ASC;
```

# Komplexní Map/Reduce operace

Jenže. Já nechci „řádky“. Já chci data zhruba v tomto formátu:

```
Streamgraph.load_data({
  max : 170,

  keywords : ['ruby', 'python', 'erlang', 'javascript', 'haskell'],

  values : [
    { date: '2010/01/01', ruby: 50, python: 20, erlang: 5, javascript: 30, haskell: 50 },
    { date: '2010/02/01', ruby: 20, python: 20, erlang: 2, javascript: 40, haskell: 43 },
    { date: '2010/03/01', ruby: 70, python: 20, erlang: 10, javascript: 80, haskell: 15 },
    { date: '2010/04/01', ruby: 20, python: 40, erlang: 8, javascript: 30, haskell: 12 },
    { date: '2010/05/01', ruby: 150, python: 30, erlang: 12, javascript: 40, haskell: 18 },
    { date: '2010/06/01', ruby: 30, python: 10, erlang: 14, javascript: 17, haskell: 14 }
  ]
});
```

# Fáze *map*

```
function(doc) {
  // Fix potentially junk date formats in docs
  var fix_date = function(junk) {
    var formatted = junk.toString().replace(/-/g, "/").replace("T", " ").substring(0,19);
    return new Date(formatted);
  };
  // Format integers to have at least two digits.
  var f = function(n) { return n < 10 ? '0' + n : n; }
  // This is a format that collates in order and tends to work with
  // JavaScript's new Date(string) date parsing capabilities, unlike rfc3339.
  Date.prototype.toJSON = function() {
    return this.getUTCFullYear()   + '/' +
           f(this.getUTCMonth() + 1) + '/' +
           f(this.getUTCDate())      + ' ' +
           f(this.getUTCHours())     + ':' +
           f(this.getUTCMinutes())   + ':' +
           f(this.getUTCSeconds())   + ' +0000';
  };
  if (doc['couchrest-type'] == 'Mention') {
    for ( keyword in doc.keywords ) {
      var key = fix_date(doc.published_at).toJSON().substring(0,10);
      var value = {};
      value[ doc.keywords[keyword] ] = 1;
      emit( key, value);
    }
  }
}
```

# Fáze *reduce*

```
function(keys, values, rereduce) {
  if (rereduce) {
    var result = {}
    for ( item in values ) {
      for ( prop in values[item] ) {
        if ( result[prop] ) { result[prop] += values[item][prop] }
        else { result[prop] = values[item][prop] }
      }
    }
    return result;
  }
  else {
    // Prepare the data for the re-reduce
    var result = {}
    for ( value in values ) {
      var item = values[value];
      for ( prop in item ) {
        if ( result[prop] ) { result[prop] += item[prop] }
        else { result[prop] = item[prop] }
      }
    }
    return result;
  }
}
```



# Výsledek

```
$ curl http://localhost:5984/customer_database/_design/Mention/_view/by_date_and_keyword?group=true
```

```
{
  "rows": [
    {
      "key": "2010/09/22",
      "value": { "ruby": 8, "python": 19 }
    },
    {
      "key": "2010/09/23",
      "value": { "ruby": 24, "python": 12 }
    },
    {
      "key": "2010/09/24",
      "value": { "ruby": 7, "python": 8 }
    }
  ]
}
```

# I ♥ JS. Nebo... ne?



# Komplexní dotazy


No... a co když potřebuji něco jako:

**Vypiš všechny programátory z Plzně.**

Smůla?

# CouchDB-Lucene

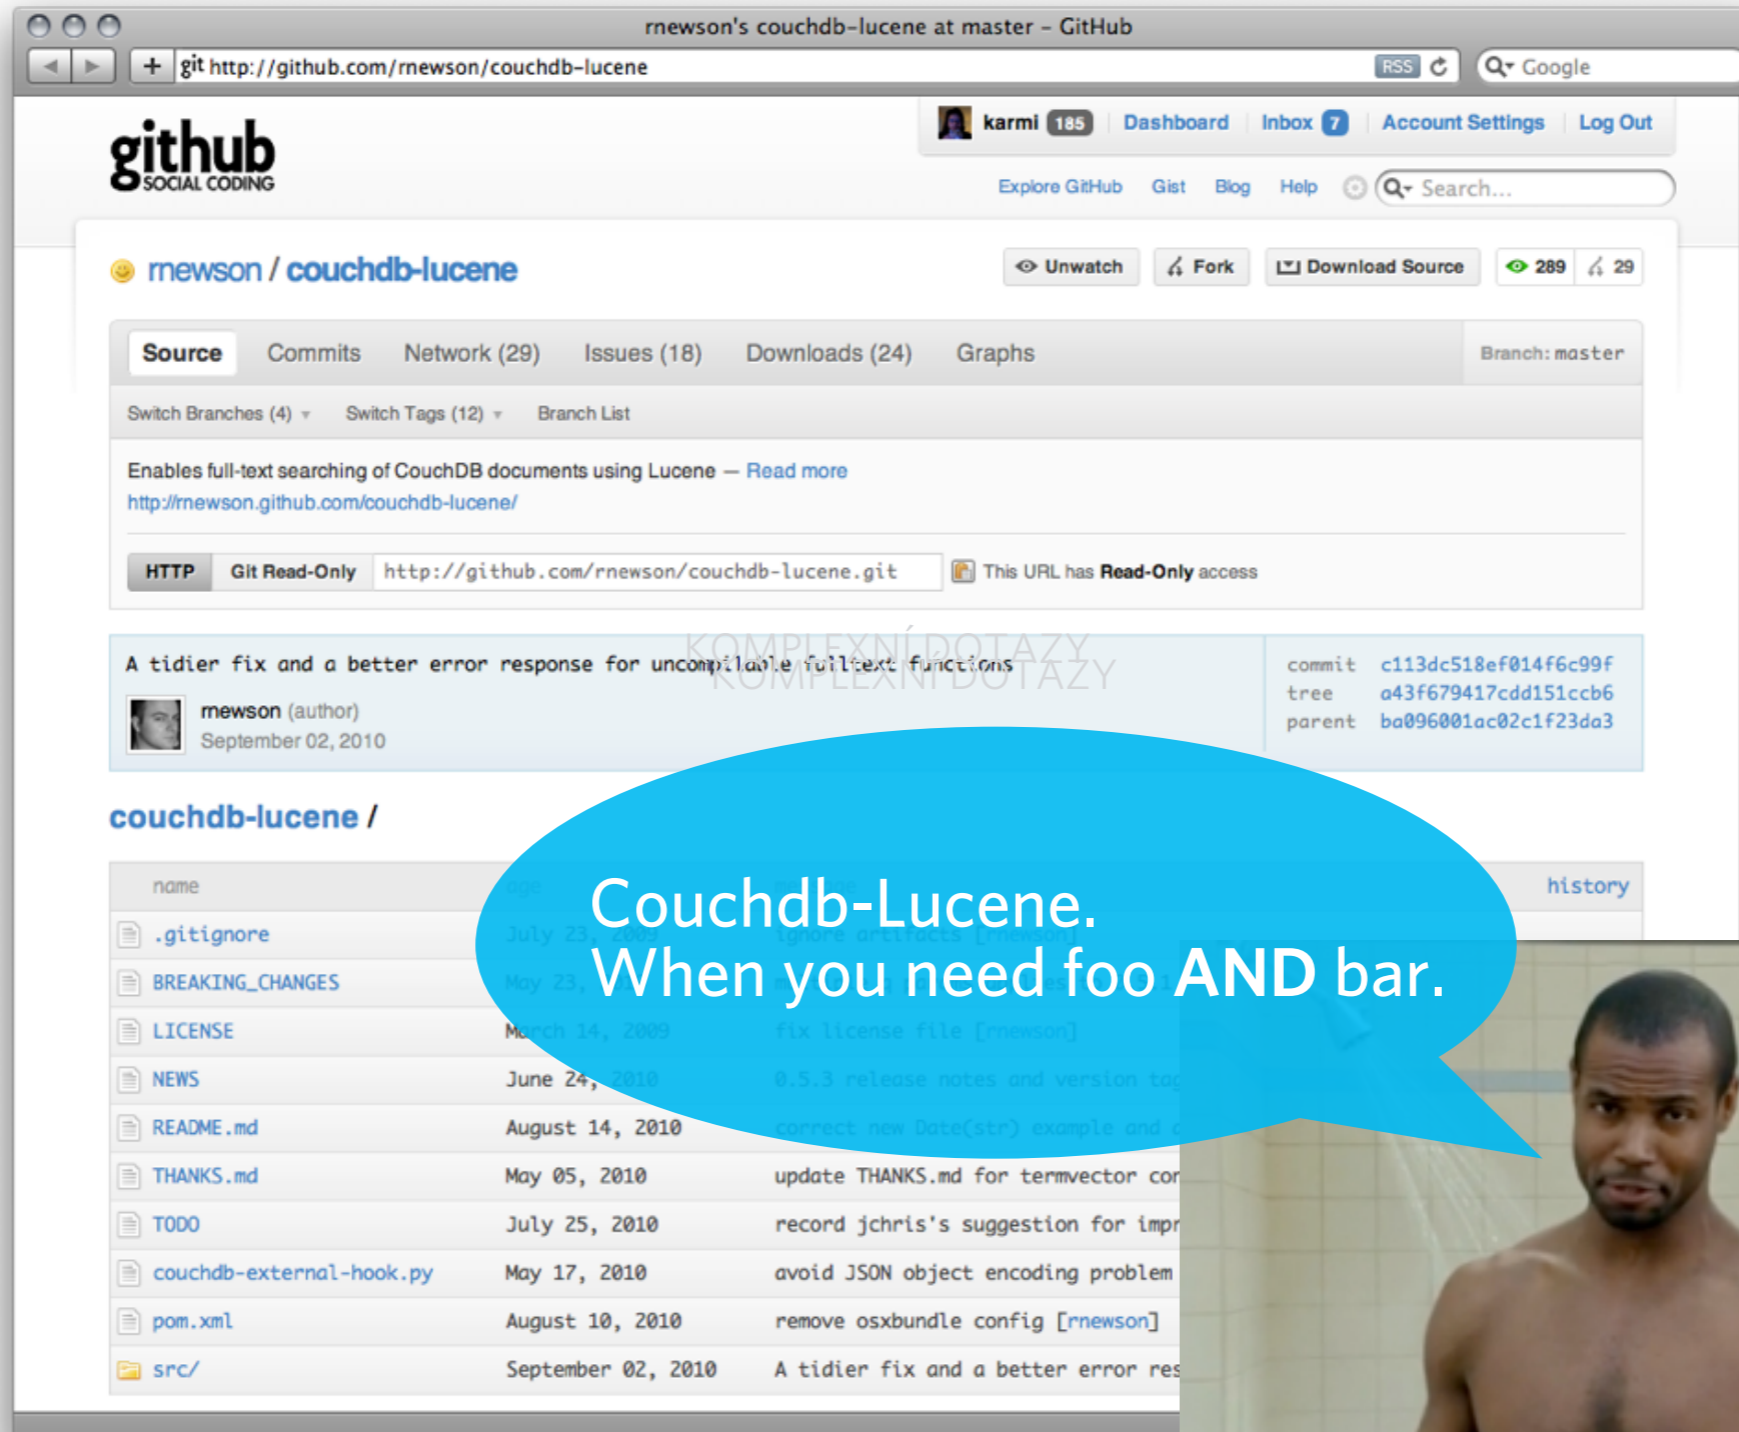
Robert Newson



Tenhle chlapík to ví.

**Vypiš všechny programátory z Plzně.**

# foo AND bar



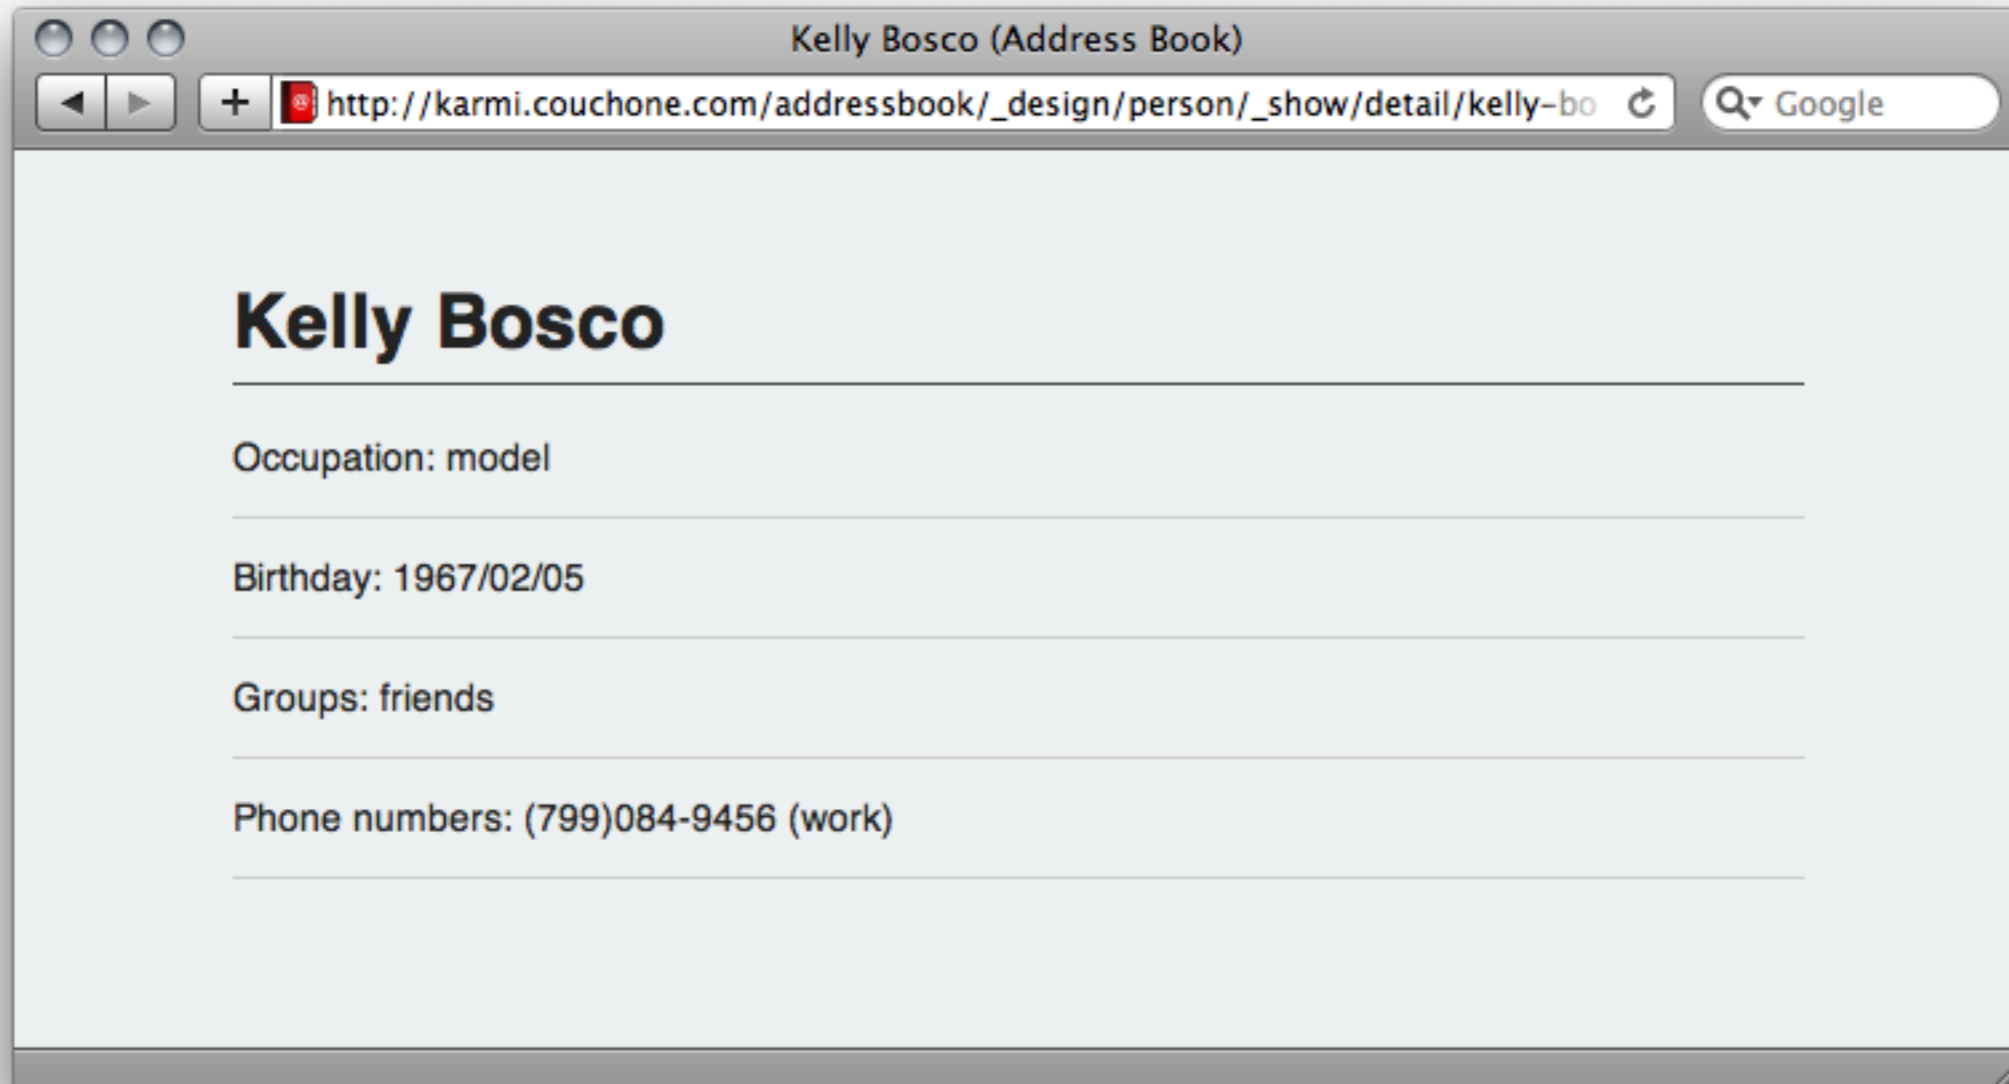


# Definice indexu

```
function(doc) {  
  
    var result = new Document();  
  
    if (doc.occupation) {  
        result.add(doc.occupation, {"field": "occupation"})  
    }  
  
    if (doc.addresses) {  
        for (address in doc.addresses) {  
            result.add(doc.addresses[address].city, {"field": "city"})  
        }  
    }  
  
    return result;  
  
}
```

[http://localhost:5984/addressbook/\\_fti/\\_design/person/search?q=occupation:programmer AND city:Plzeň](http://localhost:5984/addressbook/_fti/_design/person/search?q=occupation:programmer AND city:Plzeň)

# Ukázková aplikace



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

# Resources

- <http://guide.couchdb.org>
- <https://nosqleast.com/2009/#speaker/miller>
- <http://www.couchone.com/case-studies>
- <http://www.couchone.com/migrating-to-couchdb>
- <http://wiki.apache.org/couchdb/>
- <http://blog.couchone.com/>
- <http://stackoverflow.com/tags/couchdb/>



Otázky, prosím!

