

CouchDB, PHP & PHPillow

<http://joind.in/1445>

Kore Nordmann
<kore@php.net>
@koredn

May 15, 2010

- ▶ Kore Nordmann, <kore@php.net>
- ▶ Long time PHP developer
- ▶ Studies computer science in Dortmund, currently writing thesis
- ▶ Currently founding *Qafoo* (contact@qafoo.com / @qafoo)
- ▶ Active open source developer:
 - ▶ eZ Components (Graph, WebDav, Document), *Arbit*, PHPUnit, Torii, *PHPillow*, KaForkL, Image 3D, WCV, ...

Introduction

General

Structure

Views

Consistency

PHPillow

Applications

QA

- ▶ Structure

- ▶ Structure
- ▶ Consistency

- ▶ Structure
- ▶ Consistency
- ▶ API

- ▶ Structure
- ▶ Consistency
- ▶ API
- ▶ Applications

- ▶ Structure
- ▶ Consistency
- ▶ API
- ▶ Applications
 - ▶ I will show this by providing some views for a “blog”

Introduction

General

Structure

Views


Consistency

PHPillow

Applications

QA





- ▶ Apache top-level project



- ▶ Apache top-level project
- ▶ Build in Erlang / on Erlang/OTP

- ▶ RESTful HTTP access

- ▶ RESTful HTTP access
- ▶ HTTP is available on “all” platforms natively
 - ▶ No PHP extension required
 - ▶ Just use PHPs HTTP stream wrapper, pecl/http or curl

- ▶ RESTful HTTP access
- ▶ HTTP is available on “all” platforms natively
 - ▶ No PHP extension required
 - ▶ Just use PHPs HTTP stream wrapper, pecl/http or curl
- ▶ You can use all your known HTTP middleware
 - ▶ Reverse proxies for scaling reads (Varnish, Squid)
 - ▶ Simple custom proxy configuration for direct “AJAX” access

Introduction

General

Structure

Views

Consistency

PHPillow

Applications

QA

- ▶ Document based database

- ▶ Document based database
- ▶ No pre-defined structure (tables)

- ▶ Document based database
- ▶ No pre-defined structure (tables)
- ▶ Put in any JSON object you want

- ▶ Document based database
- ▶ No pre-defined structure (tables)
- ▶ Put in any JSON object you want
 - ▶ Even deep structures (arrays of objects)

- ▶ Document based database
- ▶ No pre-defined structure (tables)
- ▶ Put in any JSON object you want
 - ▶ Even deep structures (arrays of objects)
 - ▶ You may attach any number of files to documents

- ▶ We need to create a database first:

```
1 <?php
2
3 $fp = fopen(
4     'http://localhost:5984/phpday_blog', 'r', false,
5     stream_context_create( array(
6         'http' => array(
7             'method'      => 'PUT',
8             'content'     => null,
9             'ignore_errors' => true,
10            'header'      => 'Content-type:
11                               application/json',
12        ) ) ) );
13 var_dump( json_decode( stream_get_contents( $fp ), true
14             ) );
15 /* array(1) {
16     ["ok"] => bool(true)
17 } */
```

► Or do the same with curl:

```
1 $ curl -i -X PUT http://localhost:5984/phpday_blog
2
3 HTTP/1.1 412 Precondition Failed
4 Server: CouchDB/0.10.0 (Erlang OTP/R13B)
5 Date: Sun, 09 May 2010 10:11:09 GMT
6 Content-Type: text/plain; charset=utf-8
7 Content-Length: 95
8 Cache-Control: must-revalidate
9
10 {"error": "file_exists", "reason": "The database could not
    be created, the file already exists."}
```


► Store a document

```
1 <?php
2
3 $fp = fopen(
4     'http://localhost:5984/phpday_blog/hello_world', 'r', false,
5     stream_context_create( array(
6         'http' => array(
7             'method'     => 'PUT',
8             'content'    => json_encode( array(
9                 'type'   => 'blog_entry',
10                'title'  => 'Hello_world!',
11                'text'   => 'Some_text...',
12                'author' => 'kore',
13            ) ),
14            'ignore_errors' => true,
15            'header'       => 'Content-type:_application/json',
16        ) ) ) );
17
18 var_dump( json_decode( stream_get_contents( $fp ), true ) );
```

► Store a document

```
1 <?php
2
3 /* array(6) {
4     ["_id"] =>    string(11) "hello_world"
5     ["_rev"] =>  string(34) "1-0
6                 a73fbe05e737f0519ab2792d71d4911"
7     ["type"] =>  string(10) "blog_entry"
8     ["title"] => string(12) "Hello world!"
9     ["text"] =>  string(12) "Some ...text"
10    ["author"] => string(4) "kore"
11 } */
```

- ▶ Custom deterministic IDs can ensure uniqueness of documents
 - ▶ Just define the ID in the URL

- ▶ Custom deterministic IDs can ensure uniqueness of documents
 - ▶ Just define the ID in the URL
- ▶ CouchDB can also generate IDs for you
 - ▶ Which then are just unique hashes

► Retrieve a document

```
1 <?php
2
3 var_dump( json_decode( file_get_contents( 'http://
    localhost:5984/phpday-blog/hello_world' ) ) );
4
5 /* array(5) {
6     ["_id"]    => string(11) "hello_world"
7     ["_rev"]  => string(34) "1-0
        f563b09aa9a73b56a906919f013d492"
8     ["type"]  => string(10) "blog_entry"
9     ["text"]  => string(12) "Hello world!"
10    ["author"] => string(4) "kore"
11 } */
```

- ▶ Yeah, it is that easy.

- ▶ Yeah, it is that easy.
- ▶ Change document structure at any time

- ▶ Yeah, it is that easy.
- ▶ Change document structure at any time
- ▶ No need for (non-transaction-safe) Data Definition Language (DDL)

- ▶ Yeah, it is that easy.
- ▶ Change document structure at any time
- ▶ No need for (non-transaction-safe) Data Definition Language (DDL)
- ▶ Fits rapid development approaches with common customer requested changes to the data structure

- ▶ Yeah, it is that easy.
- ▶ Change document structure at any time
- ▶ No need for (non-transaction-safe) Data Definition Language (DDL)
- ▶ Fits rapid development approaches with common customer requested changes to the data structure
 - ▶ You need to handle this in your application properly, of course:
 - ▶ Incrementally update structure on modification
 - ▶ Liberal validation on read

Introduction

General

Structure

Views

Consistency

PHPillow

Applications

QA

- ▶ How to query such a mess?

- ▶ How to query such a mess?
 - ▶ Views are small scripts, run for all documents in a database

- ▶ How to query such a mess?
 - ▶ Views are small scripts, run for all documents in a database
 - ▶ Views are built iteratively, results stored in BTrees
 - ▶ Once built, they are *fast*

- ▶ How to query such a mess?
 - ▶ Views are small scripts, run for all documents in a database
 - ▶ Views are built iteratively, results stored in BTrees
 - ▶ Once built, they are *fast*
 - ▶ Mostly JavaScript, but also PHP, Ruby, Perl, Erlang, ...

- ▶ How to query such a mess?
 - ▶ Views are small scripts, run for all documents in a database
 - ▶ Views are built iteratively, results stored in BTrees
 - ▶ Once built, they are *fast*
 - ▶ Mostly JavaScript, but also PHP, Ruby, Perl, Erlang, ...
 - ▶ A view may emit any number of key-value pairs for each document

- ▶ How to query such a mess?
 - ▶ Views are small scripts, run for all documents in a database
 - ▶ Views are built iteratively, results stored in BTrees
 - ▶ Once built, they are *fast*
 - ▶ Mostly JavaScript, but also PHP, Ruby, Perl, Erlang, ...
 - ▶ A view may emit any number of key-value pairs for each document
 - ▶ Key and value may be any JSON structure

- ▶ Index all blog documents by their title

```
1 function( doc )
2 {
3     if ( doc.title && doc.text )
4     {
5         emit( doc.title , doc._id );
6     }
7 }
```

- ▶ Index all blog documents by their title

```
1 function( doc )
2 {
3     if ( doc.type == "blog_post" )
4     {
5         emit( doc.title , doc._id );
6     }
7 }
```

▶ Upload the view

```
1 <?php
2
3 $fp = fopen(
4     'http://localhost:5984/phpday_blog/_design/blog', '
5     r', false,
6     stream_context_create( array(
7         'http' => array(
8             'method'      => 'PUT',
9             'content'     => json_encode( array(
10                'language' => 'javascript',
11                'views'    => array(
12                    'list' => array(
13                        'map' => file_get_contents( '
14                            005_map.js' ),
15                    ), ), ),
16                'ignore_errors' => true,
17                'header'      => 'Content-type: _
18                    application/json',
19            ) ) ) );
```

▶ Fetching data

```
1 <?php
2
3 var_dump( json_decode( file_get_contents(
4     'http://localhost:5984/phpday_blog/_design/blog/
5     _view/list '
6 ) , true ) );
7
8 /* array(3) {
9     ["total_rows"] => int(1)
10    ["offset"]      => int(0)
11    ["rows"]        => array(1) {
12        [0] => array(3) {
13            ["id"]      => string(11) "hello_world"
14            ["key"]     => string(12) "Hello world!"
15            ["value"]  => string(11) "hello_world"
16        }
17    }
18 }
```

- ▶ Comparable to `SELECT * FROM table WHERE column = "value"`

- ▶ Comparable to `SELECT * FROM table WHERE column = "value"`
 - ▶ Allows range requests

- ▶ Comparable to `SELECT * FROM table WHERE column = "value"`
 - ▶ Allows range requests
 - ▶ Allows limit / offset

- ▶ Comparable to `SELECT * FROM table WHERE column = "value"`
 - ▶ Allows range requests
 - ▶ Allows limit / offset
 - ▶ Allows key across multiple fields

- ▶ Comparable to `SELECT * FROM table WHERE column = "value"`
 - ▶ Allows range requests
 - ▶ Allows limit / offset
 - ▶ Allows key across multiple fields
 - ▶ Even allow multiple indexes per document (tags)


- ▶ Comparable to `SELECT * FROM table WHERE column = "value"`
 - ▶ Allows range requests
 - ▶ Allows limit / offset
 - ▶ Allows key accross multiple fields
 - ▶ Even allow multiple indexes per document (tags)
 - ▶ Allow complex keys [category, title]

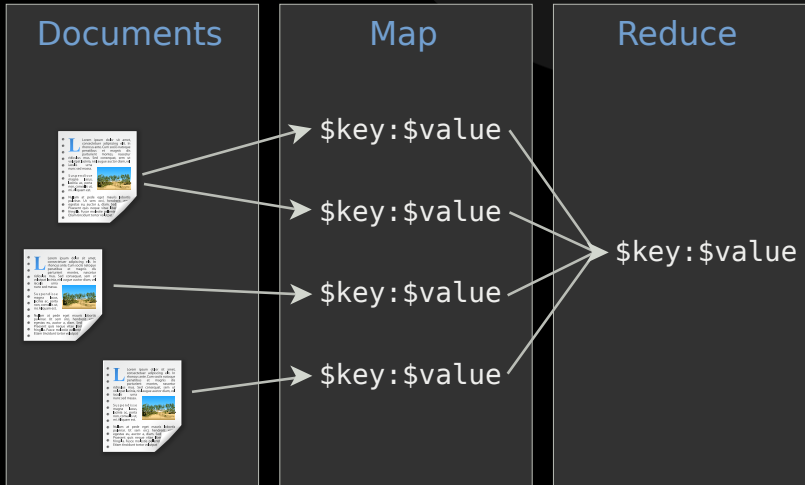
- ▶ “MapReduce is a software framework introduced by Google to support distributed computing on large data sets on clusters of computers.” [Wik09]
- ▶ Used by CouchDB to implement views

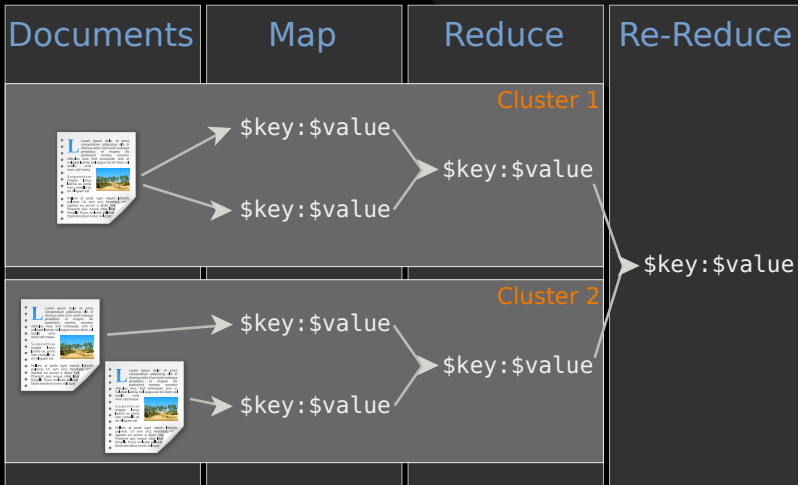
- ▶ “MapReduce is a software framework introduced by Google to support distributed computing on large data sets on clusters of computers.” [Wik09]
- ▶ Used by CouchDB to implement views
- ▶ Just a framework / pattern: You can implement “any” algorithm using map-reduce.

Documents









- ▶ Map and reduce functions are custom

- ▶ Map and reduce functions are custom
- ▶ Reduce is optional, plain view serves as a document index

- ▶ Map and reduce functions are custom
- ▶ Reduce is optional, plain view serves as a document index
- ▶ Reduce may be applied to subsets of the documents

- ▶ Map and reduce functions are custom
- ▶ Reduce is optional, plain view serves as a document index
- ▶ Reduce may be applied to subsets of the documents
- ▶ Reduce may be grouped

► The map function

```
1 function( doc )
2 {
3     if ( doc.type == "blog_post" )
4     {
5         date = new Date();
6         date.setTime( doc.edited * 1000 );
7         emit( [
8             date.getUTCFullYear(),
9             date.getUTCMonth() + 1,
10            date.getUTCDate(),
11            date.getUTCHours(),
12            date.getUTCMinutes(),
13            date.getUTCSeconds(),
14            ], 1 );
15        // You could also emit the whole doc as value
16    }
17 }
```

► The mapping result

```
1 [2008, 10, 11, 9, 11, 12] => 1
2 [2008, 10, 11, 9, 11, 12] => 1
3 [2008, 10, 11, 9, 11, 12] => 1
4 [2008, 10, 11, 9, 13, 8] => 1
5 [2008, 10, 11, 9, 13, 44] => 1
6 [2008, 10, 11, 9, 14, 2] => 1
7 [2008, 10, 12, 17, 46, 15] => 1
8 [2008, 10, 12, 17, 57, 52] => 1
9 [2008, 10, 12, 18, 0, 45] => 1
10 [2008, 10, 14, 8, 36, 29] => 1
11 [2008, 10, 14, 19, 33, 21] => 1
12 [2008, 10, 14, 19, 33, 35] => 1
```


- ▶ The simplest reduce function is just `count()`
 - ▶ Often used for statistics

```
1 function( keys , values , combine )
2 {
3     return sum( values );
4 }
```

▶ The reduce result

1 `null` \Rightarrow 42

► The grouped reduce result

```
1 [2008, 10, 11, 9, 11, 12] => 3
2 [2008, 10, 11, 9, 13, 8] => 1
3 [2008, 10, 11, 9, 13, 44] => 1
4 [2008, 10, 11, 9, 14, 2] => 1
5 [2008, 10, 12, 17, 46, 15] => 1
6 [2008, 10, 12, 17, 57, 52] => 1
7 [2008, 10, 12, 18, 0, 45] => 1
8 [2008, 10, 14, 8, 36, 29] => 1
9 [2008, 10, 14, 19, 33, 21] => 1
10 [2008, 10, 14, 19, 33, 35] => 1
```

- ▶ The filtered grouped reduce result
- ▶ `startkey=[2008,10,11]` and `endkey=[2008,10,12]`

```
1 [2008, 10, 11, 9, 11, 12] => 3
2 [2008, 10, 11, 9, 13, 8] => 1
3 [2008, 10, 11, 9, 13, 44] => 1
4 [2008, 10, 11, 9, 14, 2] => 1
```

- ▶ The grouped reduce result, with group level
- ▶ `group-level=3`

1	[2008, 10, 11]	⇒	6
2	[2008, 10, 12]	⇒	3
3	[2008, 10, 14]	⇒	3

Introduction

General

Structure

Views

Consistency

PHPillow

Applications

QA

- ▶ Multi-Version Concurrency Control

- ▶ Multi-Version Concurrency Control
- ▶ All documents in the database are versioned

- ▶ Multi-Version Concurrency Control
- ▶ All documents in the database are versioned
 - ▶ Don't use it for application level document versioning

- ▶ Multi-Version Concurrency Control
- ▶ All documents in the database are versioned
 - ▶ Don't use it for application level document versioning
- ▶ Updates and deletes need to specify the revision ID

- ▶ Multi-Version Concurrency Control
- ▶ All documents in the database are versioned
 - ▶ Don't use it for application level document versioning
- ▶ Updates and deletes need to specify the revision ID
- ▶ Changing outdated documents result in conflicts

- ▶ There is no ensured inter document consistency in CouchDB

ents:

- ▶ There is no ensured inter document consistency in CouchDB
- ▶ Different possibilities of relating documents:

- ▶ There is no ensured inter document consistency in CouchDB
- ▶ Different possibilities of relating documents:
 - ▶ List IDs of related documents in document (n:m)

- ▶ There is no ensured inter document consistency in CouchDB
- ▶ Different possibilities of relating documents:
 - ▶ List IDs of related documents in document (n:m)
 - ▶ ... both directions are feasible

- ▶ There is no ensured inter document consistency in CouchDB
- ▶ Different possibilities of relating documents:
 - ▶ List IDs of related documents in document (n:m)
 - ▶ ... both directions are feasible
 - ▶ Embed the whole related document (1:n)

- ▶ There is no ensured inter document consistency in CouchDB
- ▶ Different possibilities of relating documents:
 - ▶ List IDs of related documents in document (n:m)
 - ▶ ... both directions are feasible
 - ▶ Embed the whole related document (1:n)
- ▶ Solution depends on update-ratio

```
1  { "type": "blog_post",
2    "title": "Hello world",
3    "text": "...",
4    "comments": [
5      { "comment": "..."} ],
6  ],
7  "creator": "user-foo",
8  }
```


```
1 function( doc )
2 {
3     if ( doc.type == "blog" )
4     {
5         emit( [doc._id , 0], doc._id );
6     }
7
8     if ( doc.type == "blog_comment" )
9     {
10        emit( [doc.blog_post , doc._id], doc._id );
11    }
12 }
```

▶ JOIN query result


```
1 ["blog_post -1" , 0] => "blog_post -1"
2 ["blog_post -1" , "eaaa503adbec07"] => "eaaa503adbec07"
3 ["blog_post -1" , "e77ebd5ea0383c"] => "e77ebd5ea0383c"
4 ["blog_post -1" , "0502afb84a80a4"] => "0502afb84a80a4"
5 ["blog_post -10" , 0] => "blog_post -10"
6 ["blog_post -10" , "3872cd8a4d530f"] => "3872cd8a4d530f"
7 ["blog_post -10" , "af0dd3c4d184cf"] => "af0dd3c4d184cf"
```

- ▶ Can again be filtered...
- ▶ Using “?include_docs=true” also provides you all documents

- ▶ The CAP theorem, read more in “CouchDB: The Definitive Guide” [JCA09]



- ▶ CouchDB employs “Eventual Consistency” [Vog09]



- ▶ Important for seperated nodes (sharding)

Introduction

General

Structure

Views

Consistency

PHPillow

Applications

QA

- ▶ Object-oriented client for CouchDB
- ▶ PHP \geq 5.2 since last release (5.3 only before)
- ▶ $>96\%$ test coverage

- ▶ Object-oriented client for CouchDB
- ▶ PHP \geq 5.2 since last release (5.3 only before)
- ▶ $>96\%$ test coverage
- ▶ Still in alpha state

- ▶ Object-oriented client for CouchDB
- ▶ PHP \geq 5.2 since last release (5.3 only before)
- ▶ $>96\%$ test coverage
- ▶ Still in alpha state
 - ▶ Since CouchDB just got “beta” recently, and no new release was required.

- ▶ Lightweight layer

- ▶ Lightweight layer
- ▶ Features
 - ▶ Simple document validation constraints

- ▶ Lightweight layer
- ▶ Features
 - ▶ Simple document validation constraints
 - ▶ Automatic synchronization of views

- ▶ Lightweight layer
- ▶ Features
 - ▶ Simple document validation constraints
 - ▶ Automatic synchronization of views
 - ▶ Automatic versioning of documents

- ▶ Lightweight layer
- ▶ Features
 - ▶ Simple document validation constraints
 - ▶ Automatic synchronization of views
 - ▶ Automatic versioning of documents
 - ▶ couchdb-python compatible tool for dump and import

- ▶ Lightweight layer
- ▶ Features
 - ▶ Simple document validation constraints
 - ▶ Automatic synchronization of views
 - ▶ Automatic versioning of documents
 - ▶ couchdb-python compatible tool for dump and import
- ▶ Different connection handlers
 - ▶ PHP HTTP stream wrapper
 - ▶ Custom HTTP protocol implementation

- ▶ Lightweight layer
- ▶ Features
 - ▶ Simple document validation constraints
 - ▶ Automatic synchronization of views
 - ▶ Automatic versioning of documents
 - ▶ couchdb-python compatible tool for dump and import
- ▶ Different connection handlers
 - ▶ PHP HTTP stream wrapper
 - ▶ Custom HTTP protocol implementation
 - ▶ Which is faster, most likely because of Connection: Keep-Alive

- ▶ Query a view

```
1 $users = myUserView::all( array(  
2     'key' => 'Kore_Nordmann',  
3 ) );
```

- ▶ PHPillow validates and converts allowed view parameters

- ▶ Query a view

```
1 $users = myUserView::all( array(  
2     'key' => 'Kore_Nordmann',  
3 ) );
```

- ▶ PHPillow validates and converts allowed view parameters
- ▶ What happens:
 - ▶ View function will be uploaded to the database

- ▶ Query a view

```
1 $users = myUserView::all( array(  
2     'key' => 'Kore_Nordmann',  
3 ) );
```

- ▶ PHPillow validates and converts allowed view parameters
- ▶ What happens:
 - ▶ View function will be uploaded to the database
 - ▶ CouchDB will index all documents in the database using the view function, if the view is new

- ▶ Query a view

```
1 $users = myUserView::all( array(  
2     'key' => 'Kore_Nordmann',  
3 ) );
```

- ▶ PHPillow validates and converts allowed view parameters
- ▶ What happens:
 - ▶ View function will be uploaded to the database
 - ▶ CouchDB will index all documents in the database using the view function, if the view is new
 - ▶ View results will be returned as an array

Introduction

General

Structure

Views

Consistency

PHPillow

Applications

QA

- ▶ CouchDB allows you to attach files to documents

- ▶ CouchDB allows you to attach files to documents
- ▶ Files are replicated
 - ▶ Even incrementally since 0.11

- ▶ CouchDB allows you to attach files to documents
- ▶ Files are replicated
 - ▶ Even incrementally since 0.11
- ▶ You can serve full Web-Applications from a CouchDB
 - ▶ See CouchApp
- ▶ Deploy using PUSH-replication

- ▶ Mirror database into userspace

- ▶ Mirror database into userspace
- ▶ Offline usage and synchronization of Browser applications

- ▶ Mirror database into userspace
- ▶ Offline usage and synchronization of Browser applications
- ▶ Mozilla develops a JavaScript implementation of the CouchDB API [Moz09]

- ▶ Ubuntu One uses CouchDB
- ▶ Synchronize contacts & date between nodes, or to a server

- ▶ Ubuntu One uses CouchDB
- ▶ Synchronize contacts & date between nodes, or to a server
- ▶ Yes, all Ubuntu Karmic users already have a CouchDB running

- ▶ Arbit uses CouchDB for issue tracking, wiki, FAQ and more
- ▶ Other applications: `http://wiki.apache.org/couchdb/CouchDB_in_the_wild`

- ▶ CouchDB is fast (enough)
- ▶ Document oriented approach allows new application development approaches
- ▶ CouchDB scales really well, horizontally and vertically
- ▶ CouchDB fits web applications really well
 - ▶ RDBMS are still better for single-cluster scalable applications with strong integrity requirements.

Introduction

General

Structure

Views

Consistency

PHPillow

Applications

QA

- ▶ Apache CouchDB: <http://couchdb.org/>
- ▶ Free CouchDB book: <http://books.couchdb.org/relax/>
- ▶ PHPillow: <http://arbitracker.org/phpillow.html>

- ▶ Open questions?
- ▶ Further remarks?
- ▶ Contact
 - ▶ Mail: <kore@php.net>
 - ▶ Web: <http://kore-nordmann.de/> (Slides will be available here soonish)
 - ▶ Twitter: <http://twitter.com/koredn>
 - ▶ Comment: <http://joind.in/1445>

- ▶ Index all documents by all their words

```
1 function( doc ) {
2   if ( doc.type == "wiki" ) {
3     // Simple word indexing, does not respect overall
4     // occurrences of words,
5     // stopwords, different word separation characters,
6     // or word variations.
7     var text = doc.title.replace( /\s:.,!?-]+/g, "_" )
8     +
9     doc.text.replace( /\s:.,!?-]+/g, "_" )
10    ;
11    var words = text.split( "_" );
12    for ( var i = 0; i < words.length; ++i ) {
13      value = {};
14      value[doc._id] = 1;
15      emit( words[i].toLowerCase(), value );
16    }
17  }
18 }
```

► Index all documents by all their words

```
1  ...
2  "a"      => {wiki-8: 1}
3  "a"      => {wiki-8: 1}
4  "a"      => {wiki-8: 1}
5  "a"      => {wiki-8: 1}
6  "a"      => {wiki-81: 1}
7  "a"      => {wiki-83: 1}
8  "a"      => {wiki-83: 1}
9  "able"   => {wiki-39: 1}
10 "able"   => {wiki-56: 1}
11 "able"   => {wiki-73: 1}
12 "able"   => {wiki-80: 1}
13 "about"  => {wiki-24: 1}
14 "about"  => {wiki-43: 1}
15 "about"  => {wiki-85: 1}
16  ...
```

► Reduce by word count

```
1 function( keys , values ) {
2     var count = {};
3     for ( var i in values ) {
4         for ( var id in values[i] ) {
5             if ( count[id] ) {
6                 count[id] = values[i][id] + count[id];
7             } else {
8                 count[id] = values[i][id];
9             }
10        }
11    }
12    return count;
13 }
```

► Index all documents by all their words

```
1  ...
2  "a"      => {
3          wiki-68: 6,
4          wiki-66: 6,
5          wiki-22: 4,
6          wiki-63: 3,
7          wiki-60: 2,
8          wiki-35: 2,
9          wiki-34: 1,
10         wiki-31: 1,
11         ...
12     }
13  "able"   => {wiki-86: 1, wiki-80: 1, wiki-73: 1, wiki
14             -56: 1, wiki-39: 1}
15  "about"  => {wiki-85: 1, wiki-43: 1, wiki-24: 1}
```

- [JCA09] Noah Slater J. Chris Anderson, Jan Lehnardt, *Couchdb: The definitive guide*, O'Reilly Media, Inc., 2009.
- [Moz09] Mozilla, *Browsercouch documentation*, November 2009.
- [Vog09] Werner Vogels, *Eventually consistent - revisited*, http://www.allthingsdistributed.com/2008/12/eventually_consistent.html, December 2009.
- [Wik09] Wikipedia, *Mapreduce — wikipedia, the free encyclopedia*, 2009, [Online; accessed 27-August-2009].