
Distributed CouchApps – Embracing Eventual Consistency

International PHP Conference – Spring Edition

Kore Nordmann (@koredn, kore@qafoo.com)

05.06.2012

Part I

Introduction

Outline

Couch Apps

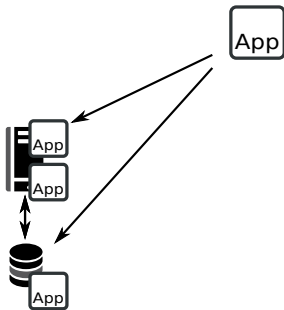
The idea



The idea



The idea

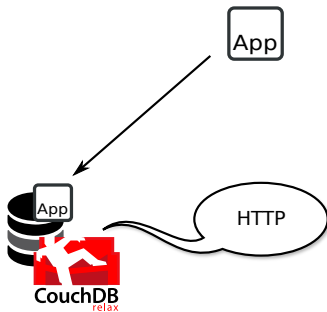


The idea

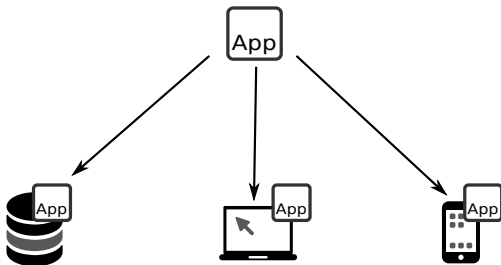
App



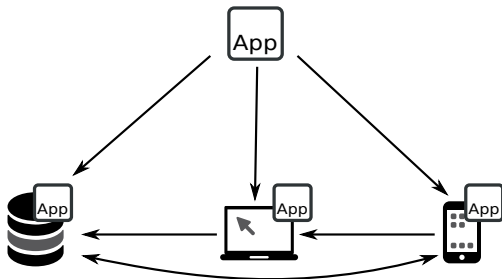
The idea



The idea



The idea



Eventual consistency

- ▶ Eventual consistency
 - ▶ CouchDB does not enforce relation integrity

Eventual consistency

- ▶ **Eventual consistency**
 - ▶ CouchDB does not enforce relation integrity
 - ▶ CouchDB servers may use delayed synchronization

Eventual consistency

- ▶ **Eventual consistency**
 - ▶ CouchDB does not enforce relation integrity
 - ▶ CouchDB servers may use delayed synchronization
 - ▶ Servers will *eventually* be consistent

Eventual consistency

- ▶ **Eventual consistency**
 - ▶ CouchDB does not enforce relation integrity
 - ▶ CouchDB servers may use delayed synchronization
 - ▶ Servers will *eventually* be consistent
- ▶ **Applications**
 - ▶ Mirror database into userspace

Eventual consistency

- ▶ **Eventual consistency**
 - ▶ CouchDB does not enforce relation integrity
 - ▶ CouchDB servers may use delayed synchronization
 - ▶ Servers will *eventually* be consistent
- ▶ **Applications**
 - ▶ Mirror database into userspace
 - ▶ Offline usage and synchronization of Browser applications

Eventual consistency

- ▶ **Eventual consistency**
 - ▶ CouchDB does not enforce relation integrity
 - ▶ CouchDB servers may use delayed synchronization
 - ▶ Servers will *eventually* be consistent
- ▶ **Applications**
 - ▶ Mirror database into userspace
 - ▶ Offline usage and synchronization of Browser applications
 - ▶ TouchDB developed by CouchBase
 - ▶ <https://github.com/couchbaselabs/TouchDB-Android>
 - ▶ <https://github.com/couchbaselabs/TouchDB-iOS>

Eventual consistency

- ▶ **Eventual consistency**
 - ▶ CouchDB does not enforce relation integrity
 - ▶ CouchDB servers may use delayed synchronization
 - ▶ Servers will *eventually* be consistent
- ▶ **Applications**
 - ▶ Mirror database into userspace
 - ▶ Offline usage and synchronization of Browser applications
 - ▶ TouchDB developed by CouchBase
 - ▶ <https://github.com/couchbaselabs/TouchDB-Android>
 - ▶ <https://github.com/couchbaselabs/TouchDB-iOS>
 - ▶ PouchDB – JavaScript CouchDB implementation based on HTML 5 Indexed DB <https://github.com/mikeal/pouchdb>

Design documents

► Defining the “logic” of a CouchApp

```
1 {
2   _id           : "_design/app",
3   _rev         : "1-0139ca4e37f873b846fd37714a191e1a",
4   _attachments : { ... },
5   views        : { ... },
6   rewrites     : [ ... ],
7   filters      : { ... },
8   lists        : { ... },
9   shows        : { ... },
10  modules       : { ... },
11  validate_doc_update : "...",
12 }
```

Part II

Consistency

Outline

Consistency

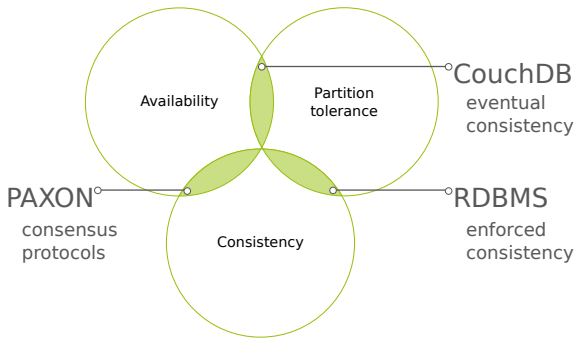
Replication

Remember:

- ▶ Multi-Version Concurrency Control (MVCC)
 - ▶ All documents in the database are versioned
- ▶ There is no ensured inter document consistency in CouchDB (relational integrity)

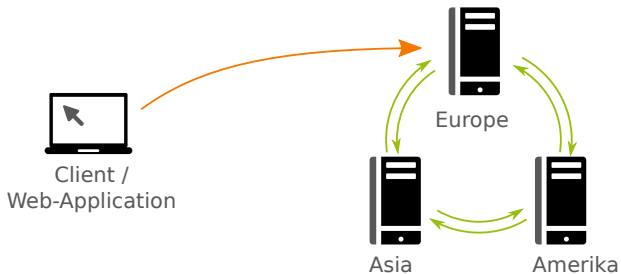
Scaling: The CAP theorem

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

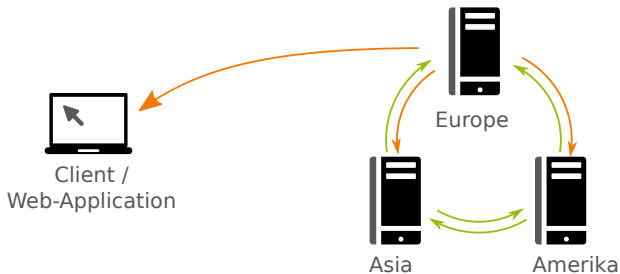


- ▶ CouchDB employs “Eventual Consistency” [Vog09]

Eventual consistency

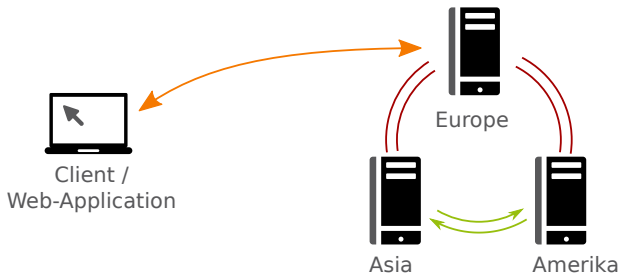


Eventual consistency



- ▶ Delayed, triggered synchronization (push, pull)
 - ▶ Deterministic (manual) conflict resolution on replication on all nodes

Eventual consistency



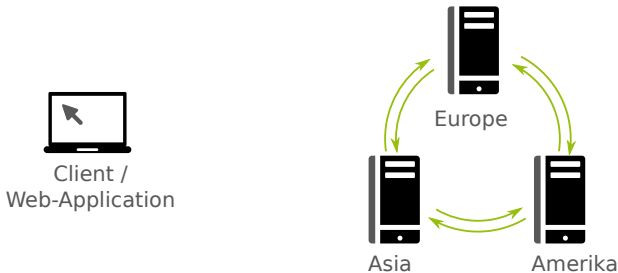
- ▶ Delayed, triggered synchronization (push, pull)
 - ▶ Deterministic (manual) conflict resolution on replication on all nodes

Eventual consistency



- ▶ Delayed, triggered synchronization (push, pull)
 - ▶ Deterministic (manual) conflict resolution on replication on all nodes
- ▶ Scales well for seldom concurrent writes

Eventual consistency



- ▶ Delayed, triggered synchronization (push, pull)
 - ▶ Deterministic (manual) conflict resolution on replication on all nodes
- ▶ Scales well for seldom concurrent writes
 - ▶ Structure your documents accordingly

Outline

Consistency

Replication

Replication

► Replication is trivial

```
1 $ curl -X POST http://localhost:5984/_replicate \  
2   -H 'Content-Type: application/json' \  
3   -d '{"source": "ipc", \  
4     "target": "http://user:pass@192.168.1.3:5984/ipc"}' \  
5 \  
6 { "ok": true, \  
7   "no_changes": true, \  
8   "session_id": "73d69e7b5cdaea059e55ed1db7802151", \  
9   "source_last_seq": 141, \  
10  "history": [ { \  
11    "session_id": "73d69e7b5cdaea059e55ed1db7802151", \  
12    "start_time": "Thu, 23 Sep 2010 08:47:05 GMT", \  
13    "end_time": "Thu, 23 Sep 2010 08:51:53 GMT", \  
14    "start_last_seq": 135, \  
15    "end_last_seq": 141, \  
16    "recorded_seq": 141, \  
17    "missing_checked": 0, \  
18    "missing_found": 1, \  
19    "docs_read": 1, \  
20    "docs_written": 1, \  
21    "doc_write_failures": 0 \  
22  } ] \  
23 }
```

Replication details

- ▶ Source and target can be any combination of remote and local URLs

Replication details

- ▶ Source and target can be any combination of remote and local URLs
- ▶ Set {"continuous": true} to have CouchDB keeping the replication alive.

Replication details

- ▶ Source and target can be any combination of remote and local URLs
- ▶ Set {"continuous": true} to have CouchDB keeping the replication alive.
 - ▶ Persisted during restarts since CouchDB 1.1

Replication details

- ▶ Source and target can be any combination of remote and local URLs
- ▶ Set `{"continuous": true}` to have CouchDB keeping the replication alive.
 - ▶ Persisted during restarts since CouchDB 1.1
 - ▶ Can be cancelled again using `{"cancel": true}`

Replication details

- ▶ Source and target can be any combination of remote and local URLs
- ▶ Set `{"continuous": true}` to have CouchDB keeping the replication alive.
 - ▶ Persisted during restarts since CouchDB 1.1
 - ▶ Can be cancelled again using `{"cancel": true}`
- ▶ Potential replication failures:
 - ▶ Crashed node
 - ▶ Network failure
 - ▶ `validate_docs_update` does not allow writing

Replication details

- ▶ Source and target can be any combination of remote and local URLs
- ▶ Set `{"continuous": true}` to have CouchDB keeping the replication alive.
 - ▶ Persisted during restarts since CouchDB 1.1
 - ▶ Can be cancelled again using `{"cancel": true}`
- ▶ Potential replication failures:
 - ▶ Crashed node
 - ▶ Network failure
 - ▶ `validate_docs_update` does not allow writing
 - ▶ Replication will be resumed once the error is fixed.

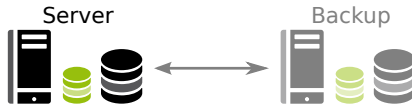
Checking replication status

```
1 $ curl -X GET http://user:pass@localhost:5984/_active_tasks
2
3 [ { "type": "Replication",
4     "task": "228689:␣http://koredn:*****@192.168.1.3:5984/ipc/␣->␣ipc",
5     "status": "W_Processed␣_source␣_update␣_#20",
6     "pid": "<0.273.0>"
7   }, {
8     "type": "Replication",
9     "task": "0444e5:␣ipc␣->␣http://koredn:*****@192.168.1.3:5984/ipc/",
10    "status": "MR_Processed␣_source␣_update␣_#141",
11    "pid": "<0.292.0>"
12  }
13 ]
```

Outline

Replication
 Filtered Replication

Filtered Replication



Client

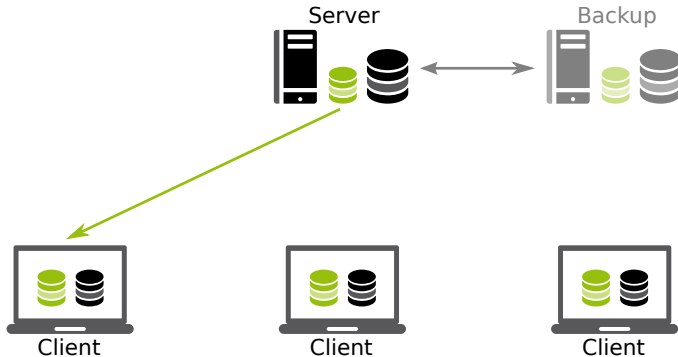


Client

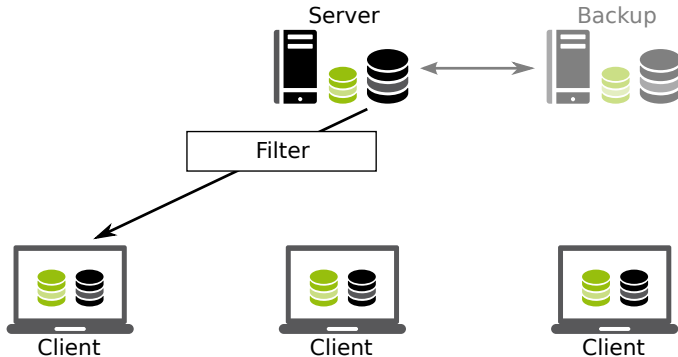


Client

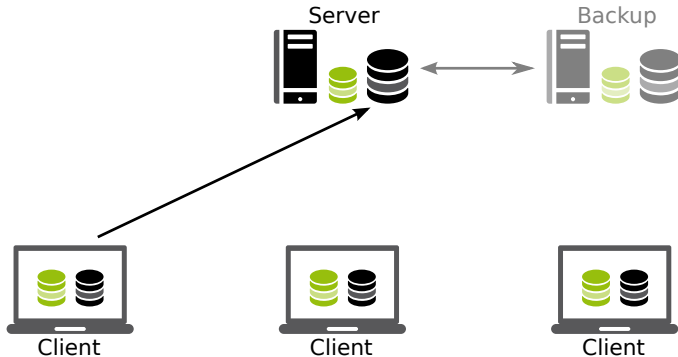
Filtered Replication



Filtered Replication



Filtered Replication



Filtered Replication

- ▶ Append a filter function to design document

```
1 { "_id": "_design/app",  
2   "language": "javascript",  
3   "filters": {  
4     "for_user": "function(_doc, _req) {_return _false ;_}" ,...  
5  
6   }...  
7  
8 }
```

Filtered Replication

- ▶ Common filtering function

```
1 function( doc, req ) {
2   if( !req.userCtx.name ) {
3     throw( "Unauthorized!" );
4   }
5
6   if( doc.recipients &&
7       doc.recipients.indexOf( req.userCtx.name ) !== -1 )
8     {
9       return true;
10    }
11
12   return false;
13 }
```

Filtered Replication

► Usage during replication

```
1 $ curl -X POST http://localhost:5984/_replicate \  
2   -H 'Content-Type: application/json' \  
3   -d '{"source": "ipc", "\\  
4   "target": "http://user:pass@192.168.1.3:5984/ipc", "\\  
5   "filter": "app/for_user"}'...
```

Part III

Mechanics

Outline

Restrospect

vHosts & Rewrites

Documents

- ▶ CouchDB stores arbitrary JSON documents
 - ▶ Deep structures are fine
 - ▶ You can index on properties in deep structures

Attachments

- ▶ CouchDB allows you to attach files to documents

Attachments

- ▶ CouchDB allows you to attach files to documents
- ▶ Files are replicated

Attachments

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

Views

- ▶ Index the documents in the database
 - ▶ Provide structured / filtered access to documents
 - ▶ Can calculate statistics on documents
 - ▶ Can provide JOIN-Views on documents

Views

- ▶ Index the documents in the database
 - ▶ Provide structured / filtered access to documents
 - ▶ Can calculate statistics on documents
 - ▶ Can provide JOIN-Views on documents
- ▶ Usually written in JavaScript / Erlang

Views

- ▶ Index the documents in the database
 - ▶ Provide structured / filtered access to documents
 - ▶ Can calculate statistics on documents
 - ▶ Can provide JOIN-Views on documents
- ▶ Usually written in JavaScript / Erlang
- ▶ Small simple functions executed for every document in database
 - ▶ Results are stored in an iteratively built append-only disk-serialized B-Tree

Views

- ▶ Index the documents in the database
 - ▶ Provide structured / filtered access to documents
 - ▶ Can calculate statistics on documents
 - ▶ Can provide JOIN-Views on documents
- ▶ Usually written in JavaScript / Erlang
- ▶ Small simple functions executed for every document in database
 - ▶ Results are stored in an iteratively built append-only disk-serialized B-Tree
- ▶ **Employing the map-reduce pattern**

Validate doc updates

- ▶ JavaScript function, which accepts or rejects document updates

Outline

Restrospect

vHosts & Rewrites

Rewrites

- ▶ Allows to rewrite URLs
- ▶ Specified in design documents

```
1 { "_id": "_design/app",  
2   "rewrites": [  
3     { "from": "/blog",  
4       "to": "../.. /blog/index.html",  
5     }, ...  
6   ]  
7 }  
8 }
```

Rewrites

- ▶ Allows to rewrite URLs
- ▶ Specified in design documents

```
1 { "_id": "_design/app",  
2   "rewrites": [  
3     { "from": "/blog",  
4       "to": "../.. /blog/index.html",  
5     },...  
6   ]  
7 }  
8 }
```

- ▶ Requested as:
`http://localhost:
5984/db/_design/app/_rewrite/blog`
- ▶ Rewrites to:
`http://localhost:5984/db/blog/index.html`

Rewrites

- ▶ Parameters are possible

```
1 { "_id": "_design/app",  
2   "rewrites": [  
3     { "from": "/images/:image",  
4       "to": ".../images/:image",  
5     },...  
6   ]  
7 }  
8 }
```

Rewrites

- ▶ Parameters are possible

```
1 { "_id": "_design/app",  
2   "rewrites": [  
3     { "from": "/images/:image",  
4       "to": ".../images/:image",  
5     },...  
6   ]  
7 }  
8 }
```

- ▶ Requested as:
`http://localhost:5984/db/_design/app/_rewrite/images/favicon.png`
- ▶ Rewrites to:
`http://localhost:5984/db/images/favicon.png`

Rewrites

- ▶ Match anything, including query parameters
 - ▶ Especially useful for views, will be covered later

```
1 { "_id": "_design/app",  
2   "rewrites": [  
3     { "from": "/*",  
4       "to": "../.. / blog/404.html",  
5     }, ...  
6   ],  
7 }  
8 }
```

Rewrites

- ▶ Match anything, including query parameters
 - ▶ Especially useful for views, will be covered later

```
1 { "_id": "_design/app",  
2   "rewrites": [  
3     { "from": "/*",  
4       "to": "../.. /blog/404.html",  
5     },...  
6   ],  
7 }  
8 }
```

- ▶ Requested as:
`http://localhost:
5984/db/_design/app/_rewrite/something`
- ▶ Rewrites to:
`http://localhost:5984/db/blog/404.html`

vHosts

- ▶ Those are still pretty ugly URLs. . .

vHosts

- ▶ Those are still pretty ugly URLs...

- ▶ **vHosts to the rescue:**

```
1 $ cat /etc/couchdb/local.ini
2 [...]
3 myhost:5984 = /db/_design/app/_rewrite
4 [...]
```


vHosts

- ▶ Those are still pretty ugly URLs...

- ▶ vHosts to the rescue:

```
1 $ cat /etc/couchdb/local.ini
2 [...]
3 myhost:5984 = /db/_design/app/_rewrite
4 [...]
```

- ▶ Requested as:

`http://myhost:5984/blog`

- ▶ Rewrites to:

`http://localhost:5984/db/blog/index.html`

vHosts

- ▶ Those are still pretty ugly URLs...

- ▶ vHosts to the rescue:

```
1 $ cat /etc/couchdb/local.ini
2 [...]
3 myhost:5984 = /db/_design/app/_rewrite
4 [...]
```

- ▶ Requested as:

`http://myhost:5984/blog`

- ▶ Rewrites to:

`http://localhost:5984/db/blog/index.html`

- ▶ A rewrite rule for "/" is also possible, of course...

Tampering

Demo

Part IV

Graceful degradation

Outline

Server side formatting

Server side formatting

- ▶ Dataformatting usually happens in the client using JavaScript

Server side formatting

- ▶ Dataformatting usually happens in the client using JavaScript
- ▶ ... but we *can* also do this on the server
 - ▶ *show* functions display a single document
 - ▶ *list* functions display a view result

Show function

- ▶ Displays a specified document
- ▶ Extend your design document:

```
1 { "_id": "_design/app",
2   "language": "javascript",
3   "shows": {
4     "wiki": "function (_doc, _req) { /* ... */ return _responseObject; }",
5   }, ...
6
7 }
```


Show function

- ▶ Displays a specified document
- ▶ Extend your design document:

```
1 { "_id": "_design/app",
2   "language": "javascript",
3   "shows": {
4     "wiki": "function (_doc, _req) { /* ... */ return _responseObject; }",
5   }, ...
6 }
7 }
```

- ▶ Request:
`/db/_design/app/_show/wiki/someDocId`

Show function

- ▶ Displays a specified document
- ▶ Extend your design document:

```
1 { "_id": "_design/app",
2   "language": "javascript",
3   "shows": {
4     "wiki": "function (_doc, _req) { /* ... */ return _responseObject; }",
5   }, ...
6
7 }
```

- ▶ Request:
`/db/_design/app/_show/wiki/someDocId?param=value`

Simple show function

```
1 function( doc, req ) {
2   if ( doc ) {
3     return {
4       body: "Hello_World"
5     }
6   } else { // document not found
7     if( req.id ) {
8       // handle unused doc id
9     } else {
10      // handle unspecified doc id
11    }
12  }
13 }
```

Simple show function

```
1 function( doc, req ) {
2   if ( doc ) {
3     return {
4       body: "Hello World"
5     }
6   } else { // document not found
7     if( req.id ) {
8       // handle unused doc id
9     } else {
10      // handle unspecified doc id
11    }
12  }
13 }
```

- ▶ Request and response:
<https://wiki.apache.org/couchdb/ExternalProcesses>

List function

- ▶ Displays a specified view
- ▶ Extend your design document:

```
1 { "_id": "_design/app",
2   "language": "javascript",
3   "views": {
4     "by_name": {
5       "map": "function(doc) { if (doc.type == 'user') emit(doc.name, null) }",
6     },
7   },
8   "lists": {
9     "posts": "function(_head, _req) { /* ... */ }",
10  }, ...
11 }
12 }
```

List function

- ▶ Displays a specified view
- ▶ Extend your design document:

```
1 { "_id":      "_design/app",
2   "language": "javascript",
3   "views": {
4     "by_name": {
5       "map":    "function(doc) { if (doc.type == 'user') emit(doc.name, null) }",
6     },
7   },
8   "lists": {
9     "posts": "function(_head, _req) { /* ... */ }",
10  }, ...
11 }
12 }
```

- ▶ Request:
/db/_design/app/_list/posts/by_name

List function

- ▶ Displays a specified view
- ▶ Extend your design document:

```
1 { "_id": "_design/app",
2   "language": "javascript",
3   "views": {
4     "by_name": {
5       "map": "function(doc) { if (doc.type == 'user') emit(doc.name, null) }",
6     },
7   },
8   "lists": {
9     "posts": "function(_head, _req) { /* ... */ }",
10  }, ...
11 }
12 }
```

- ▶ Request:
`/db/_design/app/_list/posts/by_name?viewParams`

List function example

```
1 function(head, req) {
2   start( {
3     "headers": {
4       "Content-Type": "text/html"
5     }
6   } );
7   while( row = getRow() ) {
8     send( row.value );
9   }
10 }
```


List function example

```
1 function(head, req) {
2   start( {
3     "headers": {
4       "Content-Type": "text/html"
5     }
6   } );
7   while( row = getRow() ) {
8     send( row.value );
9   }
10 }
```

- ▶ start() call to send headers

List function example

```
1 function(head, req) {
2   start( {
3     "headers": {
4       "Content-Type": "text/html"
5     }
6   } );
7   while( row = getRow() ) {
8     send( row.value );
9   }
10 }
```

- ▶ start() call to send headers
- ▶ getRow() call to receive next row from view

List function example

```
1 function(head, req) {
2   start( {
3     "headers": {
4       "Content-Type": "text/html"
5     }
6   } );
7   while( row = getRow() ) {
8     send( row.value );
9   }
10 }
```

- ▶ start() call to send headers
- ▶ getRow() call to receive next row from view
- ▶ send() echo something to the client

Outline

Server side formatting Modules

Server side code reuse

- ▶ It is possible to reuse code on server side
- ▶ CouchDB can use JavaScript code from CommonJS modules, attached to your design document

```
1 { _id: "_design/test",  
2   language: "javascript",  
3   modules: {  
4     jquery : "exports.jquery = /*_Source_code_*/",  
5     something: "exports.answer = 42;",  
6   },  
7   shows: {  
8     post: "function () {_var lib =_require ('modules/something ');_return lib .answer;_};"  
9   },  
10 }
```

- ▶ The name `modules` is arbitrary

Server side code reuse

- ▶ It is possible to reuse code on server side
- ▶ CouchDB can use JavaScript code from CommonJS modules, attached to your design document

```
1 { _id: "_design/test",
2   language: "javascript",
3   modules: {
4     jquery : "exports.jquery = /*_Source_code_*/",
5     something: "exports.answer = 42;",
6   },
7   shows: {
8     post: "function () { var lib = require('modules/something'); return lib.answer; }",
9   },
10 }
```

- ▶ The name `modules` is arbitrary
 - ▶ You can even put the library below `views` or `shows`
 - ▶ https://wiki.apache.org/couchdb/CommonJS_Modules

Summary

- ▶ Eventually Consistent Offline Replicated JavaScript Frontend, Backend & Data

Summary

- ▶ Eventually Consistent Offline Replicated JavaScript Frontend, Backend & Data
- ▶ Example application: <http://github.com/Qafoo/Lounge>
- ▶ Application stub: <http://github.com/Qafoo/Stub>

Thanks for listening

- ▶ More:
 - ▶ <http://kore-nordmann.de>
 - ▶ <http://qafoo.com>
 - ▶ <http://github.com/Qafoo>
- ▶ Feedback:
 - ▶ <https://joind.in/6650>

Bibliography I

- [JCA09] Noah Slater J. Chris Anderson, Jan Lehnardt, *Couchdb: The definitive guide*, O'Reilly Media, Inc., 2009.
- [Vog09] Werner Vogels, *Eventually consistent - revisited*, http://www.allthingsdistributed.com/2008/12/eventually_consistent.html, December 2009.