



Bringing Riak to the Mobile Platform

Kresten Krab Thorup
Hacker

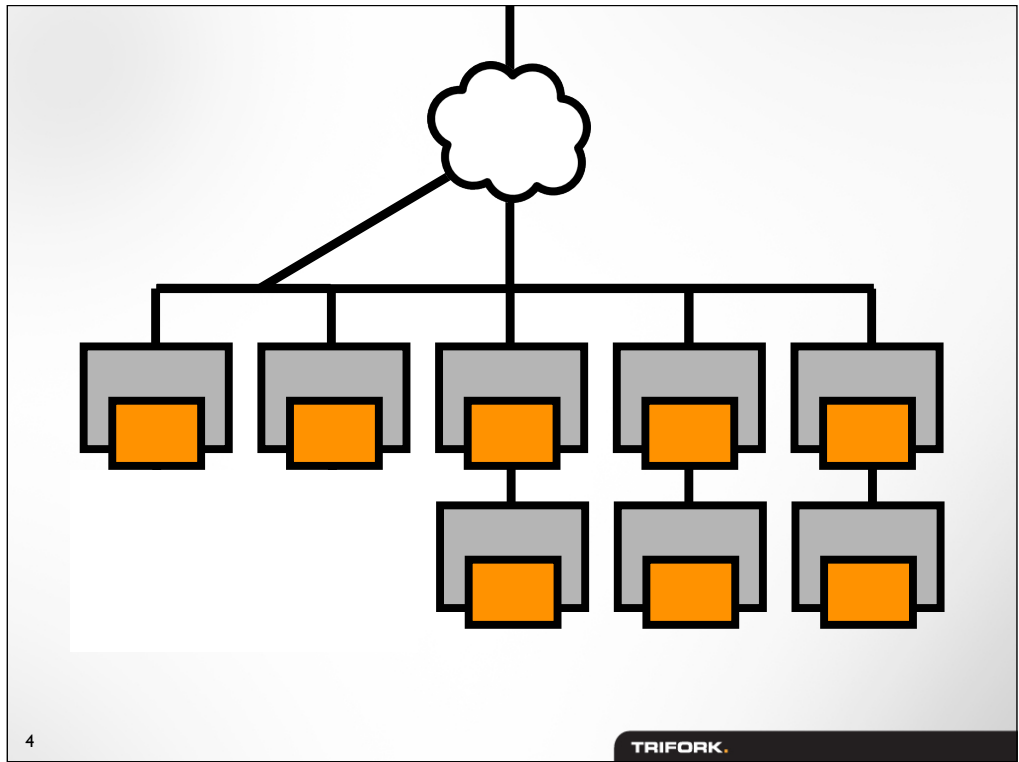



Outline

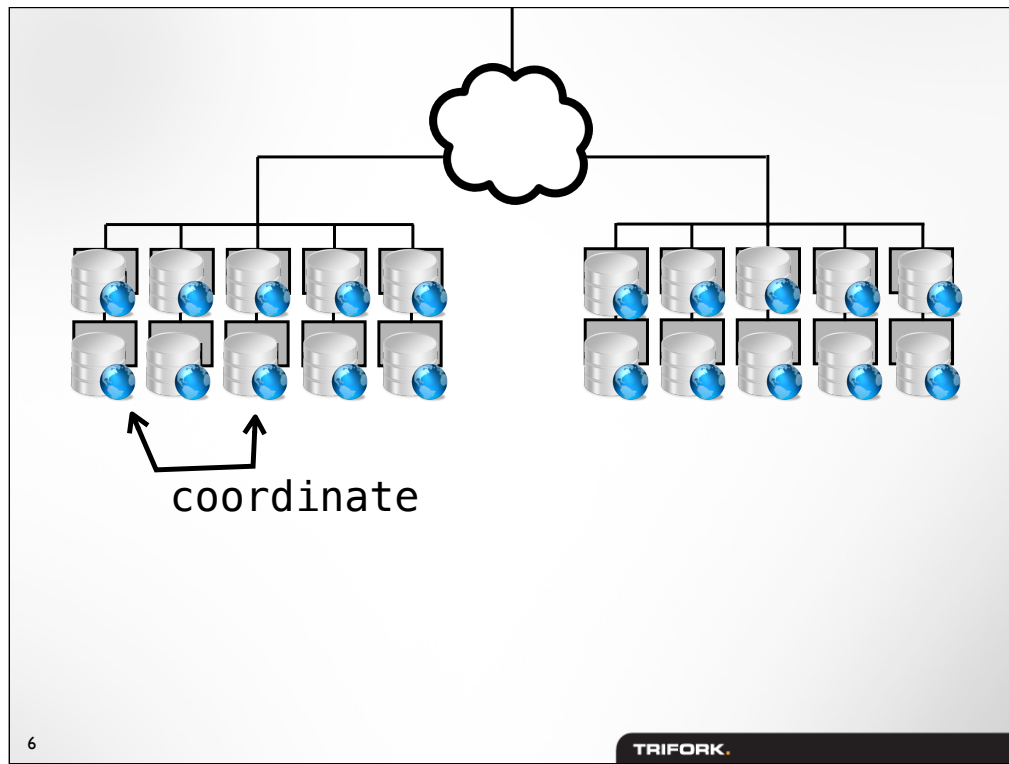
- Riak and it's Data Model
- RiakSync, a protocol for Key/Value synchronization
- RiakMobile, Riak clients for mobile

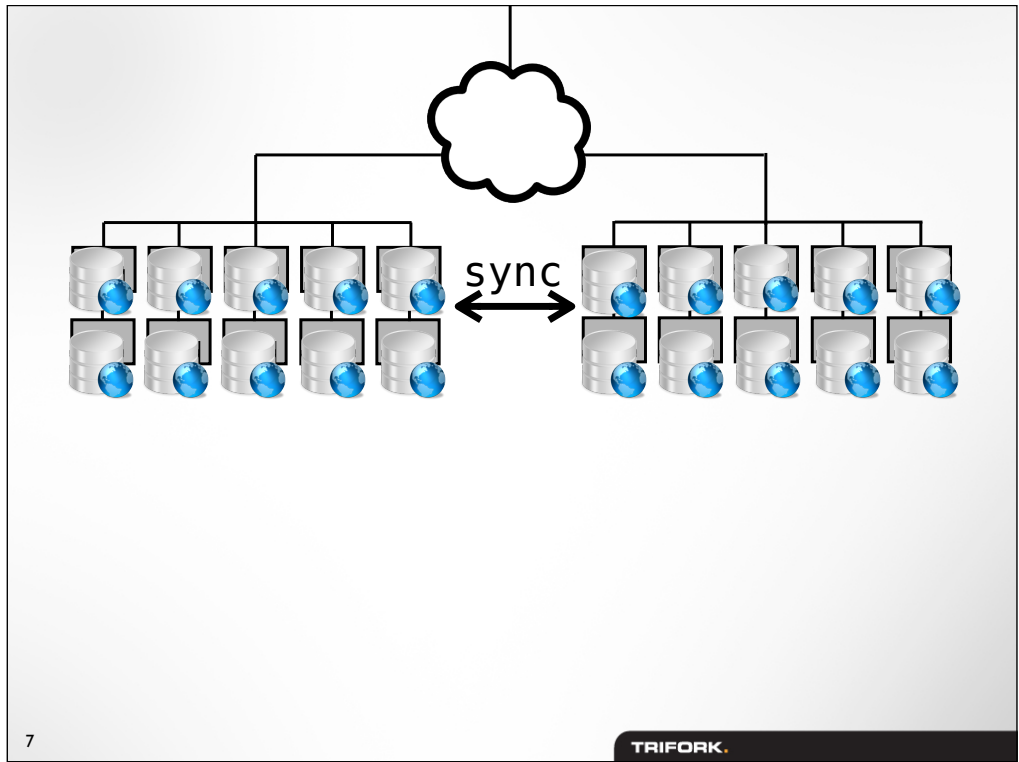
What is Riak, anyway?

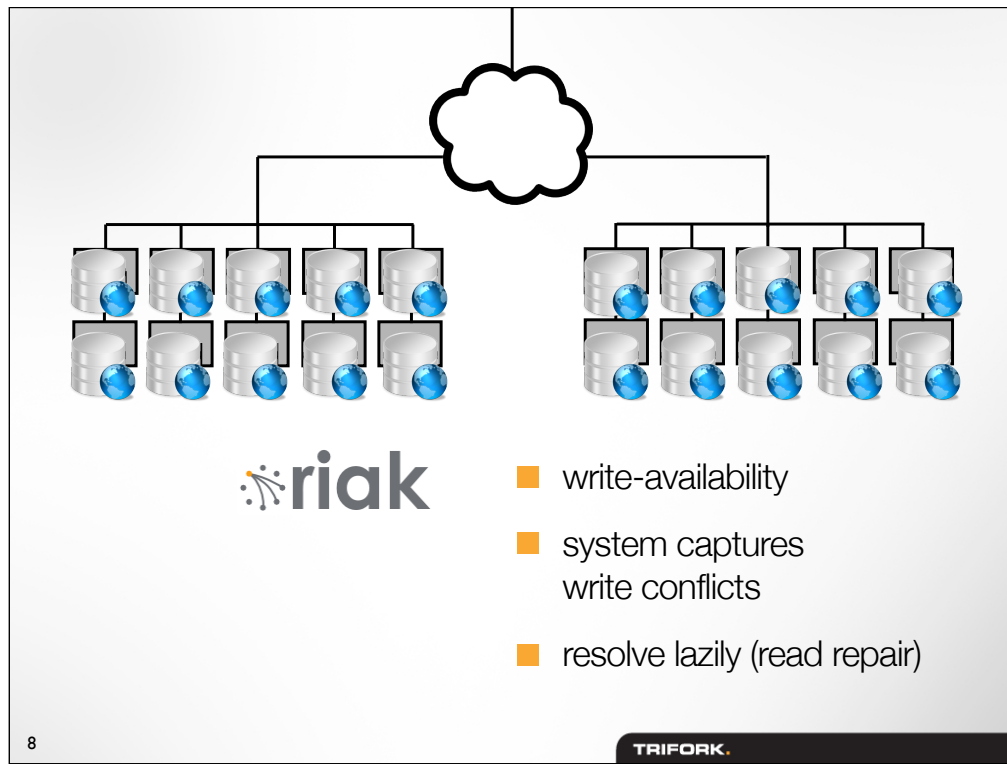













Shared Medicine Card



foo



ola: foo=4



ola: foo=6




foo



ola: foo=4



joe: foo=7




foo




ola: foo=4




ola: foo=6



joe: foo=7



merge



merge

Riak API review

`connect(ClientID) -> Client`


```
key()    :: {Bucket, Key}
datum()  :: {ContentType, RawData}
```

`Client:insert(key(), datum()) -> {ok, VClock} | ...`

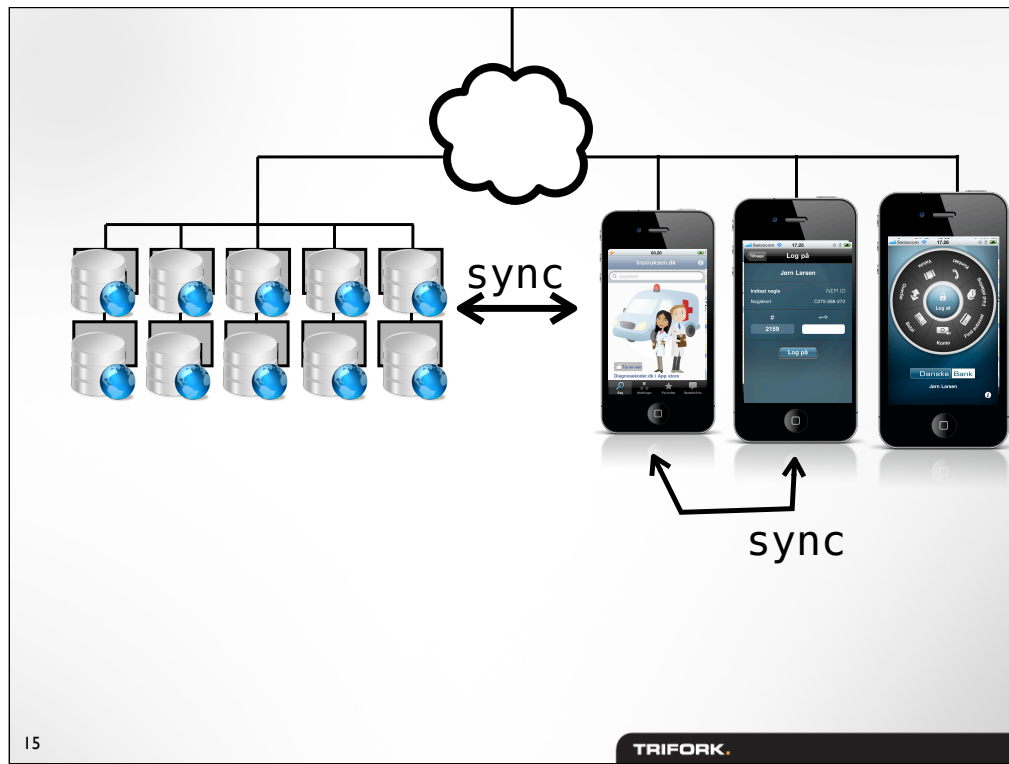
`Client:lookup(key()) -> {ok, VClock, [datum()]}` | ...

`Client:update(key(), VClock, datum()) -> ok | ...`

`Client:delete(key(), VClock) -> ok | ...`

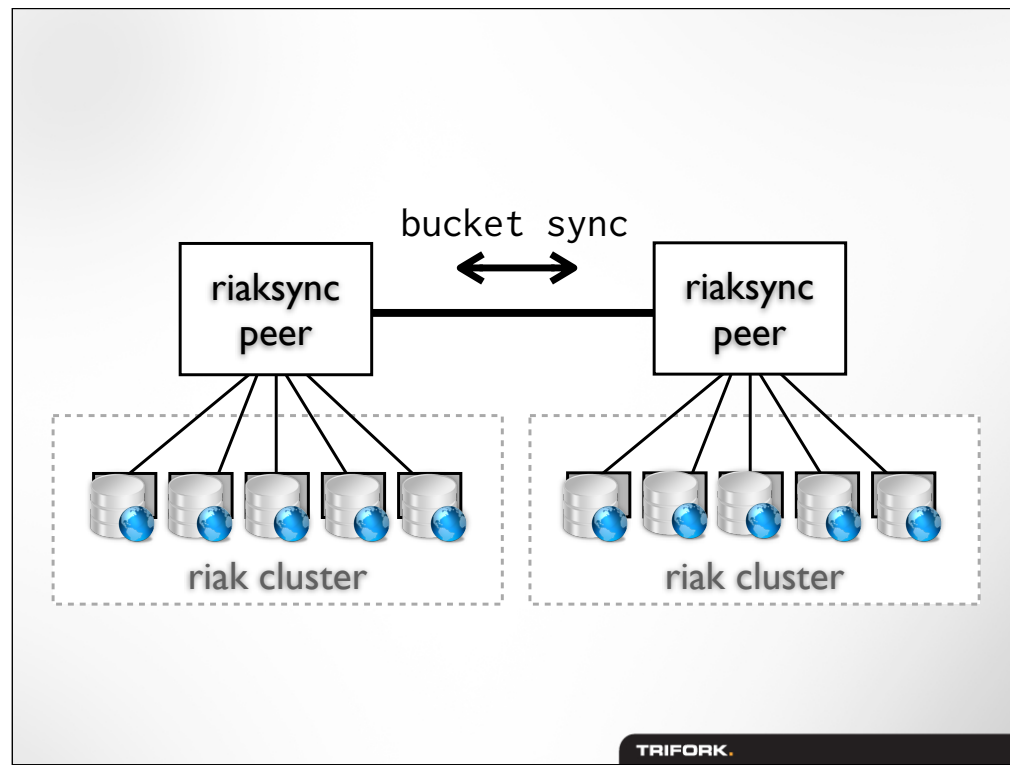


- write-availability
- system captures write conflicts
- resolve lazily (read repair)



Today: Two Things

- **RiakSync**: Protocol for Riak bucket replication
- **RiakMobile**: Java & C++ based client implementations



```
LocalStore = riak:client_connect('riak@127.0.0.1')
```

```
%% sync via protobuf socket
```

```
> riak_sync:sync(LocalStore, <<"bucket">>,
                 "172.1.35.204", 8082)
```

```
%% Other site has riak_sync running...
```

```
(riak@172.1.35.204)1> riak_sync:start()
```

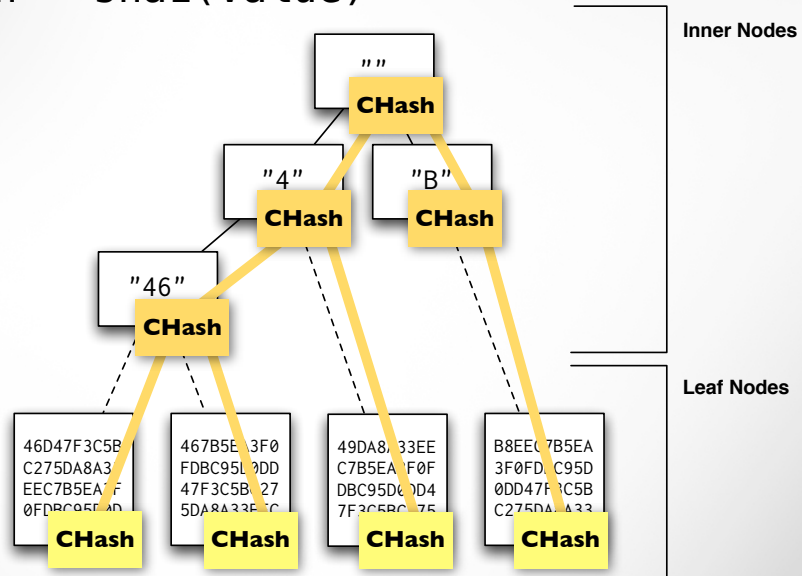
```
{riak_sync, [{pb_ip, "172.1.35.204"},
              {pb_port, 8082},
              {riak, 'dev@127.0.0.1'}]}
```

Riak Sync Protocol

- Riak bucket replication
- Suitable for up to ~100.000 objects/bucket
- Asymmetric work load
- Designed for high-latency networks

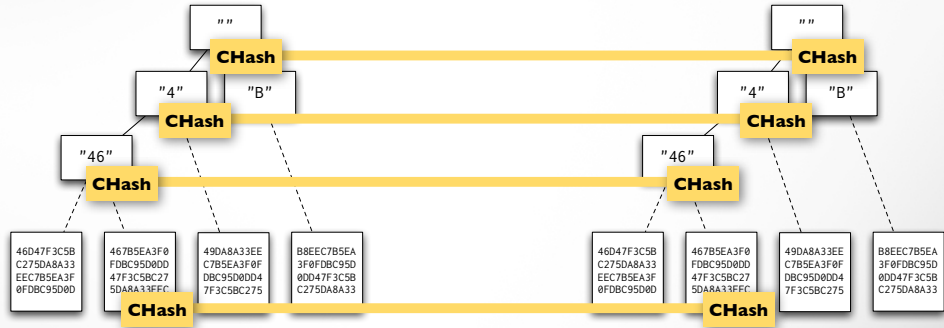
- **Very different design criteria than normal Riak cluster-cluster replication**

StorageKey = sha1(Key)
CHash = sha1(Value)



riksync peer

riksync peer

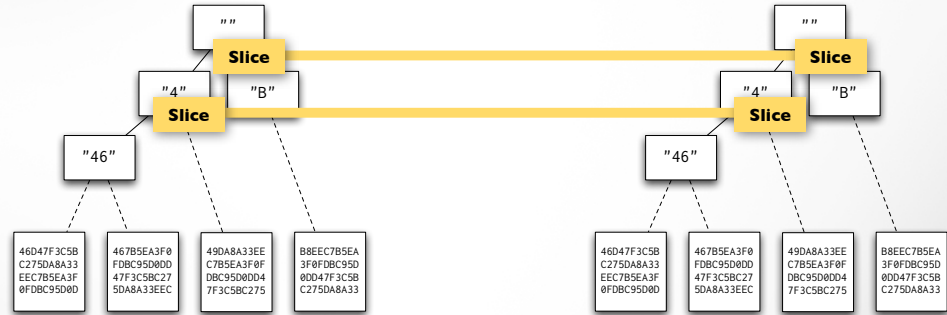


HashTree Slices

```
slice_msg() ::  
  {slice, Path, CHash,  
   [{Nibble :: 0..16, shallow_node()}]}.  
  
shallow_node() ::  
  {inner, CHash}  
  | {leaf, [{Key, VClock, CHash}, ...]}.
```

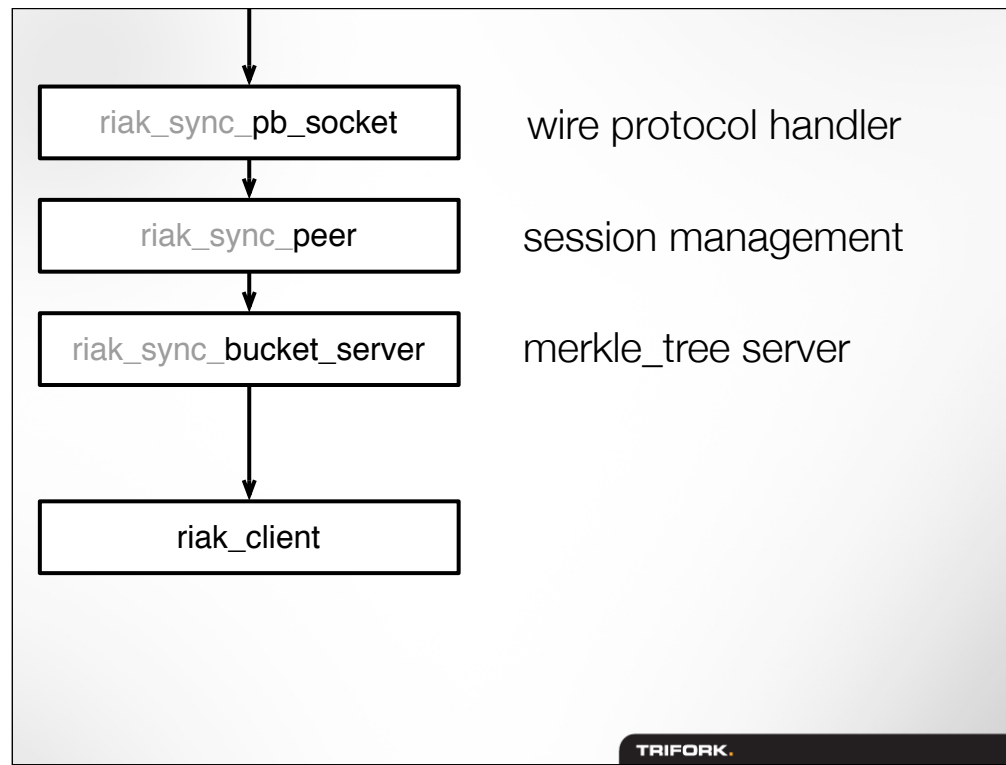
riksync peer

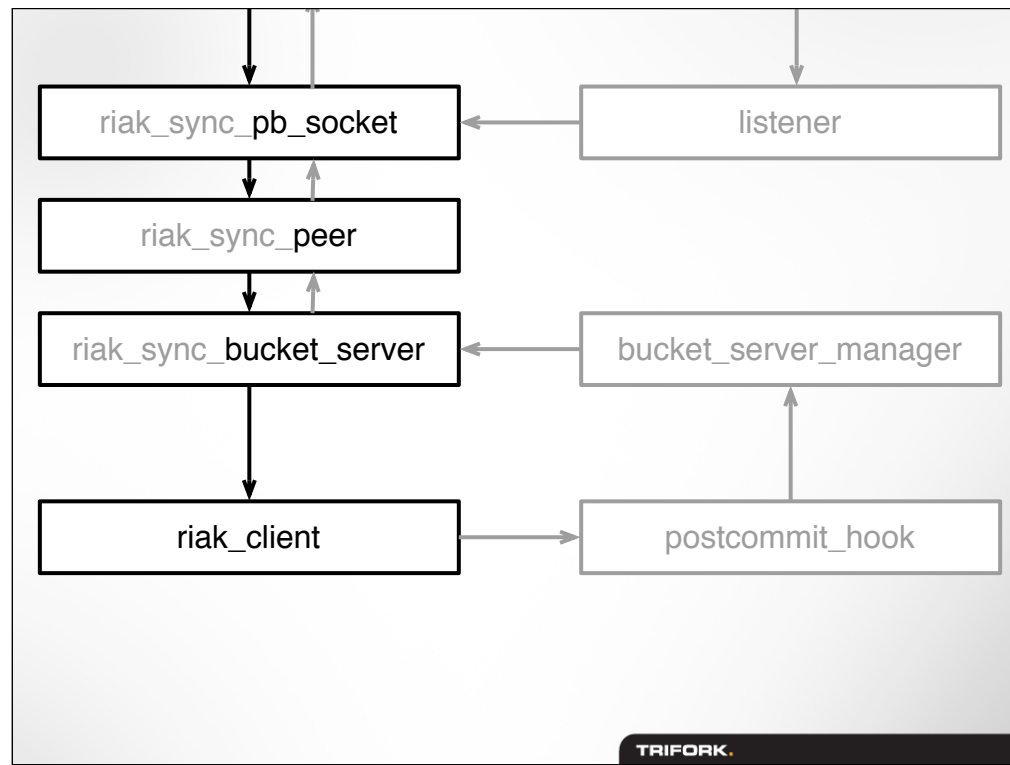
riksync peer



RiakSync Protocol

- Asynchronous
- Parallelizes quickly (client does breadth first)
- Server peer reorders requests to reduce session state
 - **First do any put/get requests**
 - **The execute slice requests in depth first order**
- Packet size (slice messages) ~ 1000 bytes
- Client needs (almost) no session state







Interesting Issues

- Capturing deletes
 - RiakSync's `postcommit` hook creates records of deleted data to maintain vector clocks for those
 - If a client comes along a week after the delete, it will thus not re-create the deleted record.
- Cheating riak's vector clocks
 - In Riak, you can only "put and increment vclock" but RiakSync needs "put sans vclock increment".


Riak Mobile



Content Distribution



Reliable Queueing

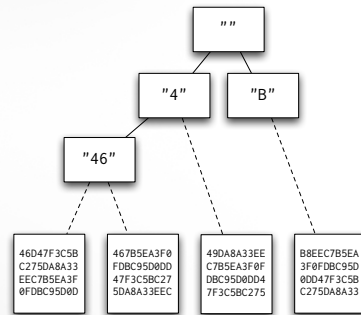


riakmobile

riaksync peer



riakmobile



- Client stores data in local data store as a HASH TREE
- Always “ready to sync” (no session state, no cpu/memory intensive computation)

DATA

ROW_ID	KEY	CONTENT TYPE	DATA

TRIGGER



MERKLE_LEAF

HKEY	VCLOCK	CHASH	KEY

MERKLE_INNER

PATH	CHASH	CHILDREN

```

class RiakClient {
    bool insert(const string& key,
                /*out*/ VClock& version,
                const Datum& value);

    bool lookup(const string& key,
                /*out*/ VClock& version,
                /*out*/ Data& values);

    void update(const string& key,
                const VClock& version,
                const Datum& value);
}

class Data {
    Datum[] datas;
    int count;
}

class Datum {
    string content_type;
    string body;
}

```

```
class RiakClient { ...continued...  
    // create client for bucket  
    RiakClient(string& bucket);  
  
    // initiate RiakSync, and receive live updates  
    void connect(string& host, int port);  
  
    // disconnect from RiakSync server  
    void disconnect();  
  
    // observe updates  
    void set_commithook(update_hook callback);  
  
}  
  
typedef void (*update_hook)(const string& key,  
                             const Data& data);
```

Summary

- Riak Data Model
- RiakSync, a protocol for Key/Value synchronization
- RiakMobile, Riak clients for mobile

You need to restart your computer. Hold down the Power button for several seconds or press the Restart button.

Veillez redémarrer votre ordinateur. Maintenez la touche de démarrage enfoncée pendant plusieurs secondes ou bien appuyez sur le bouton de réinitialisation.

Sie müssen Ihren Computer neu starten. Halten Sie dazu die Einschalttaste einige Sekunden gedrückt oder drücken Sie die Neustart-Taste.

コンピュータを再起動する必要があります。パワーボタンを数秒間押し続けるか、リセットボタンを押してください。

Thank You

@drkrab



Thank You

