

Velkommen til GeekNight

Hands-on introduktion til Riak

Taler: Rune Skou Larsen

 Follow @GeekNight_dk




GOTO Aarhus 2012

Premier software development conference created by developers for developers.

Conference: Oct. 1-3 // Training: Sept. 30, Oct. 4-5

Se mere på www.gotocon.com/aarhus-2012

 Follow @GOTOcon



Kurser

Trifork har et bredt program af kurser og er altid klar til at møde efterspørgslen efter nye.

RIAK Training d. 20. september, 2012

Få et overblik over de forskellige Riak features og en diskussion om fordelene herved.

Se flere kurser på www.trifork.com eller kontakt vores kursusafdeling på training@trifork.com

Hands-on introduktion til



Agenda

- Intro
- Installér Riak på din maskine
- Øvelser med HTTP API'et
- Key concepts
- Øvelser med Siblings

Pause

- Byg et cluster
- Key concepts (Ringen, N, W, R mv.)
- Øvelser for at demonstrere ovenstående
- Hvis tid: backends, map reduce, 2i, Search og link walking.

Hvad er jeres erfaringer med Riak?

Mine erfaringer med Riak

- Har arbejdet med Riak siden version 0.14
- Fælles Medicinkort (FMK)
 - Med til at udvikle FMK's java Riak-klienter
 - FMK har kørt i produktion på Riak over et år.
 - Arbejder pt. på at flytte FMK's auditlog over på Riak.
- Riak arkitekt på talrige nationale og internationale Riak projekter
 - Seoul, Korea
 - Toulouse, Frankrig
 - Leeds, UK
 - København

Introduction to Riak

What is risk?

riak is ...

A database

open Source

Masterless

distributed

Fault tolerant

highly Available

scalable

BUILT for the web

inspired by amazon's dynamo

riak has ...

http api

protocol buffer api

client Libraries written in ...

Erlang

ruby

java

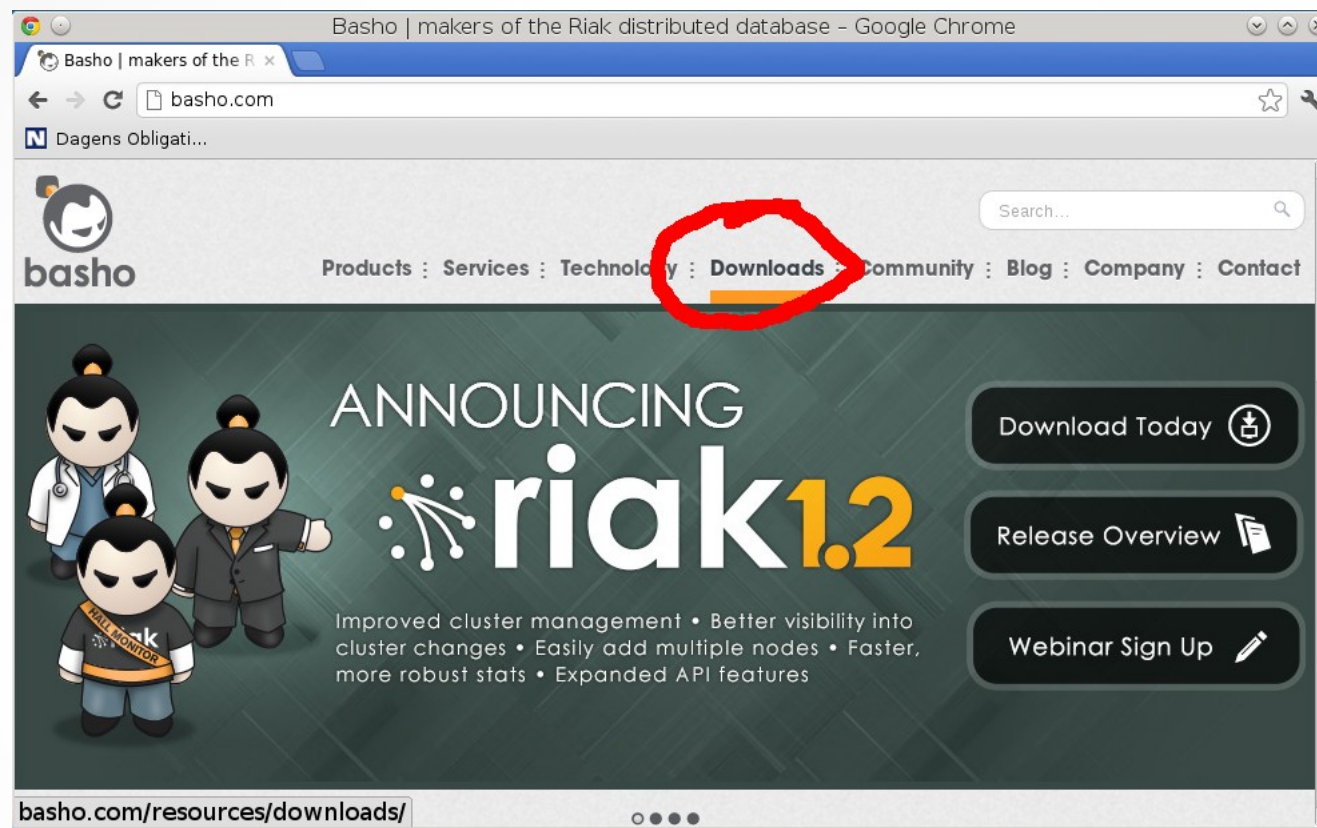
python

javascript/node.js, haskell,
c/c++, clojure, scala, ...

Trifork og Riak

- Bidragyder til Riak
- Riak kurser
 - <http://www.trifork.com/trainingcalendar>
 - Skræder-syede
- Licenser til Riak Enterprise og CS
- Udviklingshjælp
- Produktionssupport

Installér Riak



Riak → 1.2 → 1.2.0 → vælg din distro
(eller tar-arkiv)

Start og test din installation

- riak ping

Node 'riak@127.0.0.1' not responding to pings.

- riak start

- riak ping

Pong

Øvelse 1 – REST API

1) Installer curl

2) curl localhost:8098

Should return an answer

3) PUT data to a key

```
curl -v localhost:8098/buckets/training/keys/my-first-object -X PUT -H "content-type: text/plain" -d "Mit første put"
```

4) GET the data back

```
curl -v localhost:8098/buckets/training/keys/my-first-object
```

5) DELETE the key

```
curl -v localhost:8098/buckets/training/keys/my-first-object -X DELETE
```

6) Check that it is gone

```
curl -v localhost:8098/buckets/training/keys/my-first-object
```

Should return 404

Øvelser 2a – Siblings

1) set Bucket Property allow_mult to true to make Riak track siblings.

```
curl -v -X PUT -H "Content-Type: application/json" -d '{"props": {"allow_mult": true}}' 127.0.0.1:8098/riak/mybucket
```

2) Check that the bucket property was set correctly

```
curl 127.0.0.1:8098/buckets/mybucket/props | json_pp
```

"| json_pp" is just for pretty printing the json response.

3) PUT data to a key

```
curl -v 127.0.0.1:8098/buckets/mybucket/keys/one -X PUT -H "content-type: text/plain" -d "First value"
```

4) GET the data back

```
curl -v 127.0.0.1:8098/buckets/mybucket/keys/one
```

Notice the X-Riak-Vclock header. Riak has assigned a vector clock to the key now. Vector clocks are a means to track causality of changes.

5) PUT data to the key again without specifying vector clock

```
curl -v 127.0.0.1:8098/buckets/mybucket/keys/one -X PUT -H "content-type: text/plain" -d "Second value"
```

6) GET the data back again

```
curl -v 127.0.0.1:8098/buckets/mybucket/keys/one
```

Notice that siblings are returned. Riak did not know which version was the newest and kept both. Also notice that the vector clocks from the individual siblings have been combined into a consolidated vector clock.

Key concepts

cluster

node

bucket

key

value

metadata

vector clock

riak object

Øvelser 2b – Siblings

7) GET all siblings by specifying content-type as multipart

```
curl -v 127.0.0.1:8098/buckets/mybucket/keys/one -H "Accept: multipart/mixed"
```

Should return both values.

8) Resolve the conflict by PUTting using X-Riak-Vclock header. REMEMBER TO REPLACE THE VCLOCK WITH THE ONE FROM THE ABOVE GET REQUEST.

```
curl -v 127.0.0.1:8098/buckets/mybucket/keys/one -X PUT -H "content-type: text/plain" -H "X-Riak-Vclock: a85hYGBgzGDKBVlcR4M2cvst/bLYEpky2NleP3C+iRfFgA=" -d "Third value"
```

9) GET the data back again, and see that the siblings were resolved

```
curl -v 127.0.0.1:8098/buckets/mybucket/keys/one
```

This should return "Third value" (at the very bottom).

Vi bygger cluster...

- Stop riak

riak stop

- Firewall

- Enten disable firewall..

- ..eller

```
angiv portrange i /etc/riak/app.config { kernel, [  
    {inet_dist_listen_min, 6000},  
    {inet_dist_listen_max, 7999}  
]},
```

og sørg for, at din maskine ikke firewall'er riaks porte:

- epmd's listener: TCP:4369
- handoff_port listener: TCP:8099
- angivet range 6000-7999

- Find din maskines IP-adresse

- angiv ip i /etc/riak/vm.args i stedet for localhost

fx

-name [riak@172.30.252.221](#)

- Slet ring filen

```
sudo rm -rf /var/lib/riak/ring/
```

- Start riak

riak start

Join noderne sammen

- Se cluster configuration

riak-admin member-status

- Planlæg join request

riak-admin cluster join [riak@123.123.123.123](#)

- Verificer evt. planen

riak-admin cluster plan

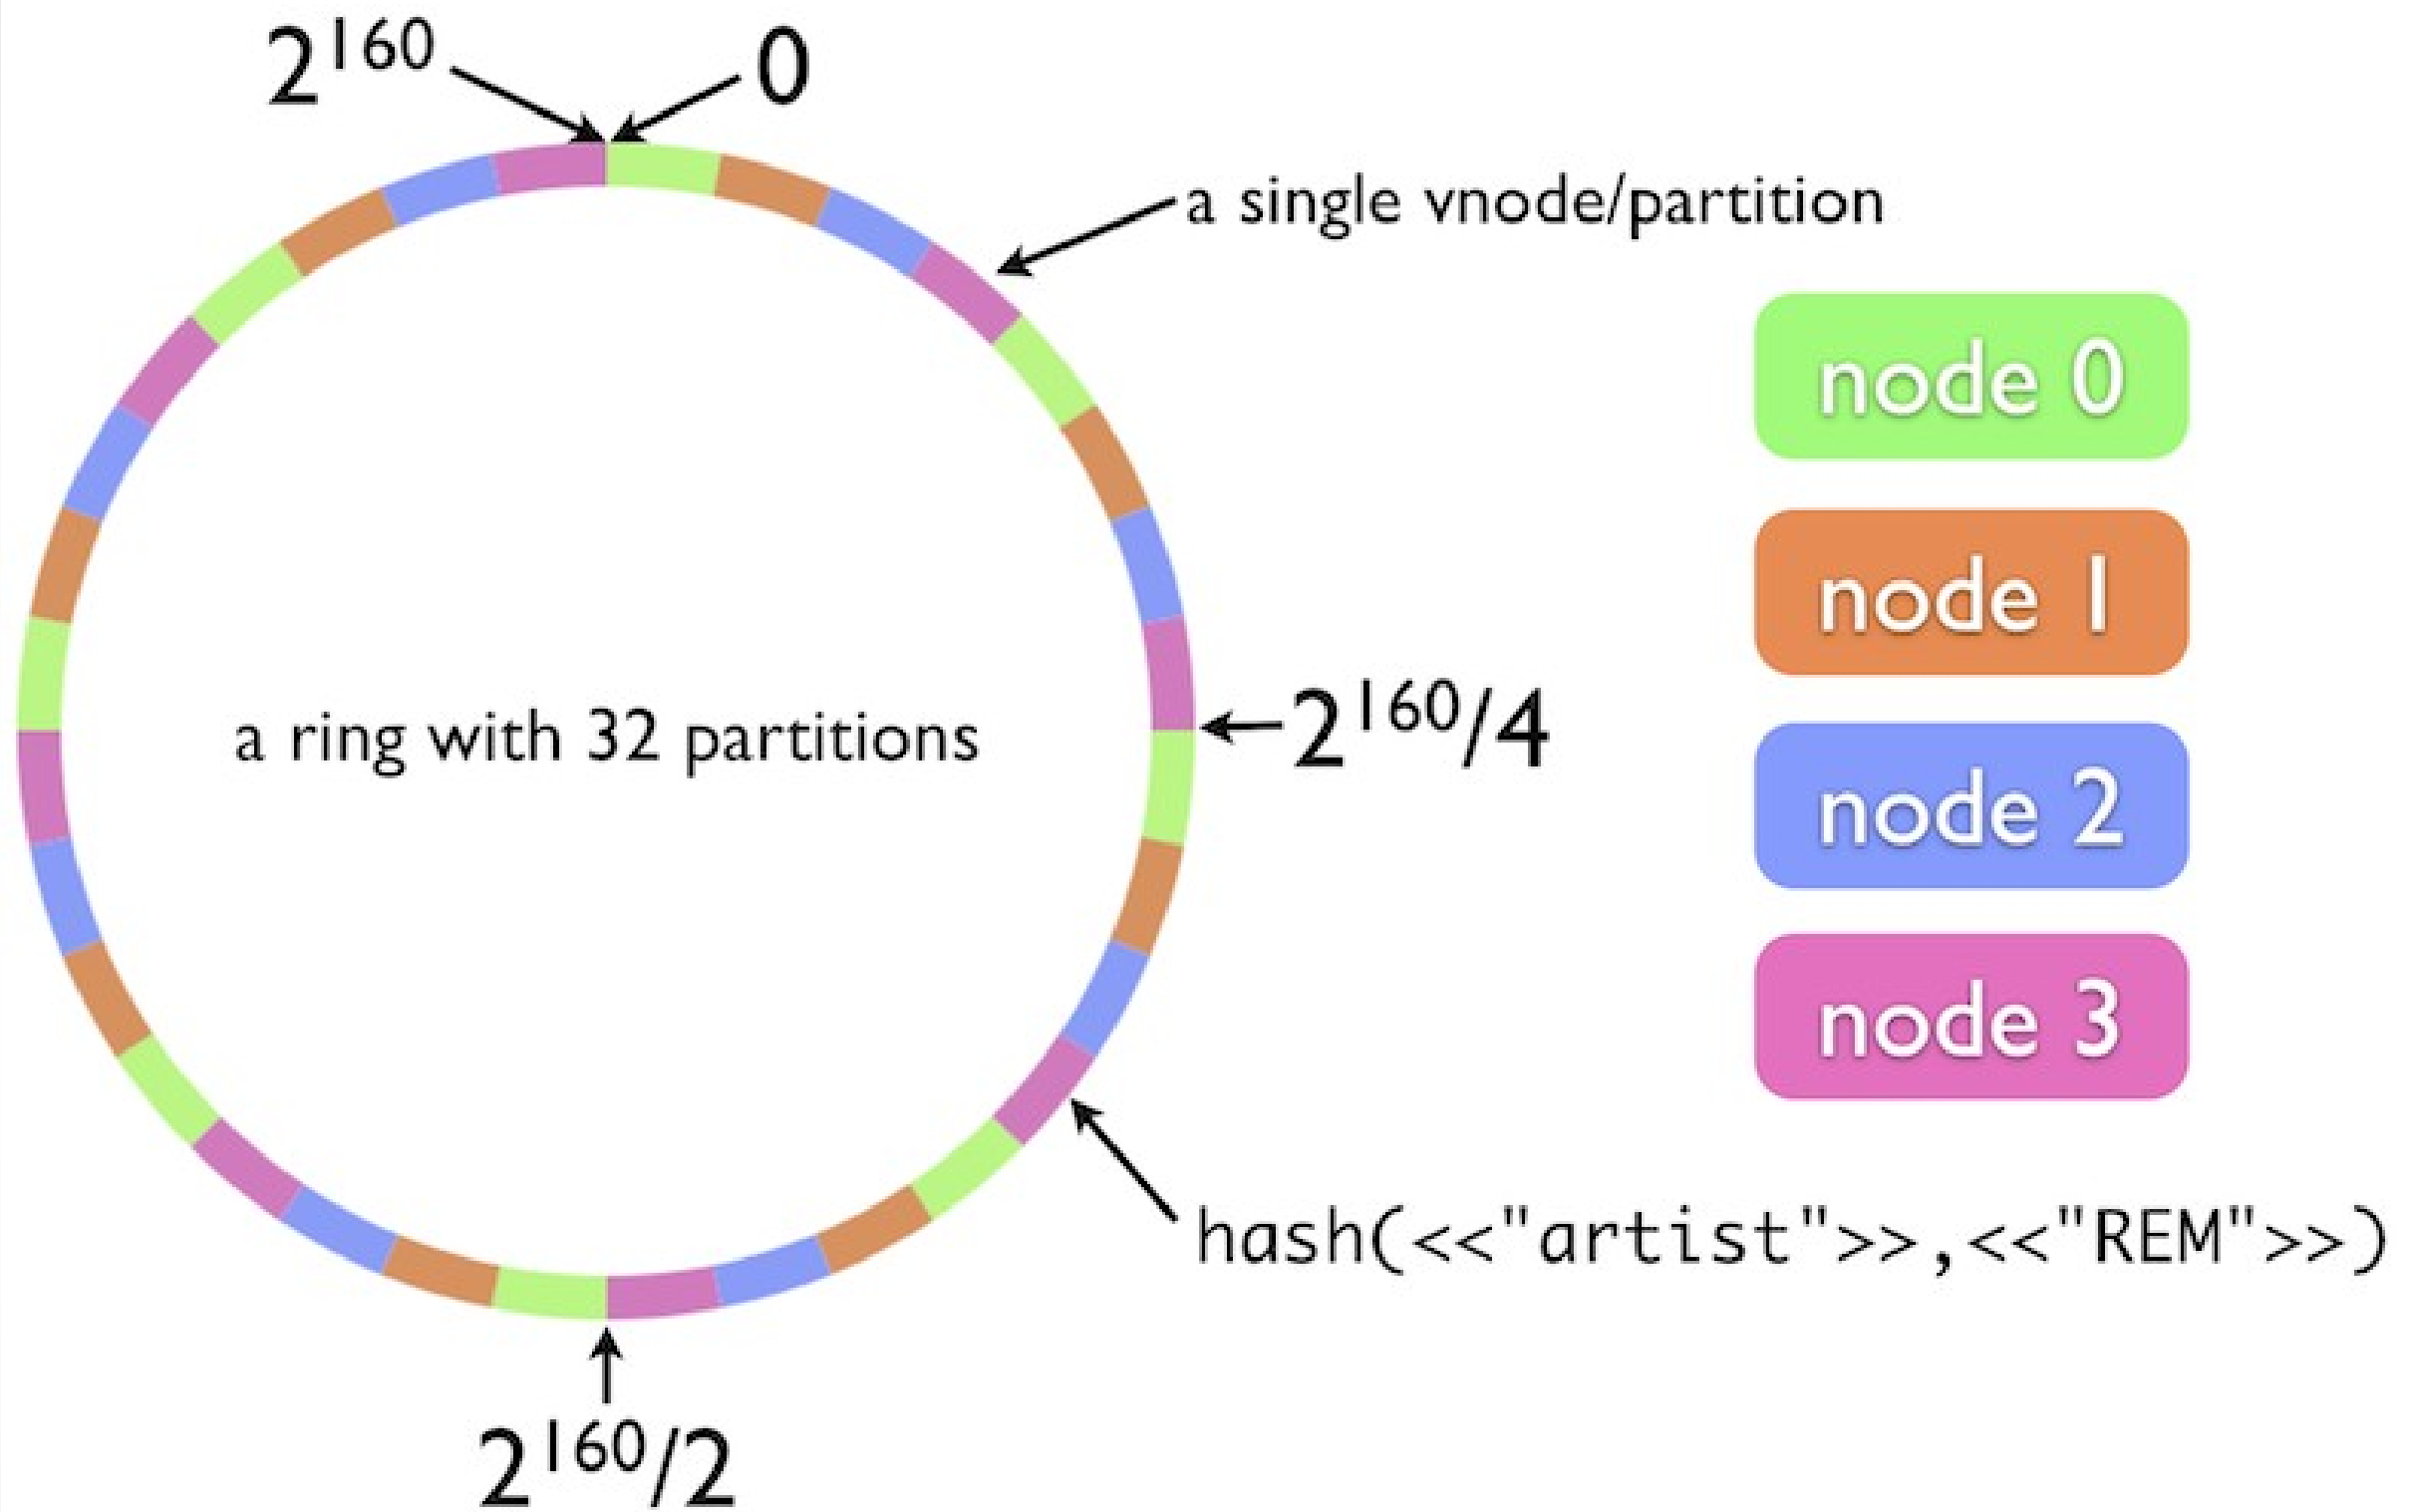
- Eksekver planen

riak-admin cluster commit


- Se ny cluster configuration

riak-admin member-status

the Ring



replication



partition

not to be confused
with

a network partition

Vnode

handoff

ownership handoff

hinted handoff

quorum

r value

w value

dw

pr/pw

storage backend

map reduce

link walking

Riak Search

secondary indexes (2i)