


Scaling that Rails App

Christian Amor Kvalheim
Linqia.com





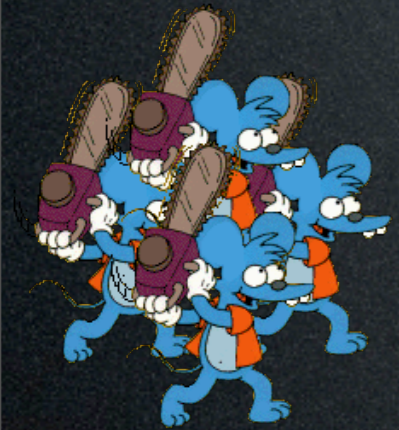
- Search Service for social communities and groups
- Target 10+ million groups index
- ROR/Nginx/Solr/Memcached/Postgres

The Goal

- Help you spend less time f* * *ing around
- Get the gears grinding

Understand your pattern

- Mostly read only (aka content)
- Fully dynamic (aka twitter)
- Mixed dynamic (aka social networks)



Saving Your App

- The traffic is killing me
- Pure Ruby execution not so hot :(
- I need to save my database

WOR

- Uh Webrick
- Hmm Mongrel
- Config Hell Apache + Mongrel
- Better than expected Nginx + Mongrel

Slim it down dude

- Size of Page
- Size of Scripts
- Scripts
- CSS



http://images.google.com/images?q=homer_simpson&prev=/images?

How did I get that object

- What's in that Query
- Repetitive Queries == Caching candidate
- DRY content fetching

Cache_fu Memcached



- Memcached
- Distributed hash map
- Just grow it when needed

Cache_fu Memcached

Cache_fu + Memcached =



- Rails plugin
- Active Record caching
- Fragment and Action caching

ActiveRecord caching

- Save the database
- Write custom finders
- Cache for effect

```
# Caching options
acts_as_cached :include =>
[:languages, :countries, :comments,
:thumbnail, :community_members]
```


Page Caching

- Render content to disk
- Bypass Mongrel, let Nginx/Apache do the lifting
- Not good for dynamic content
- Standard page cache invalidation based on shared file system

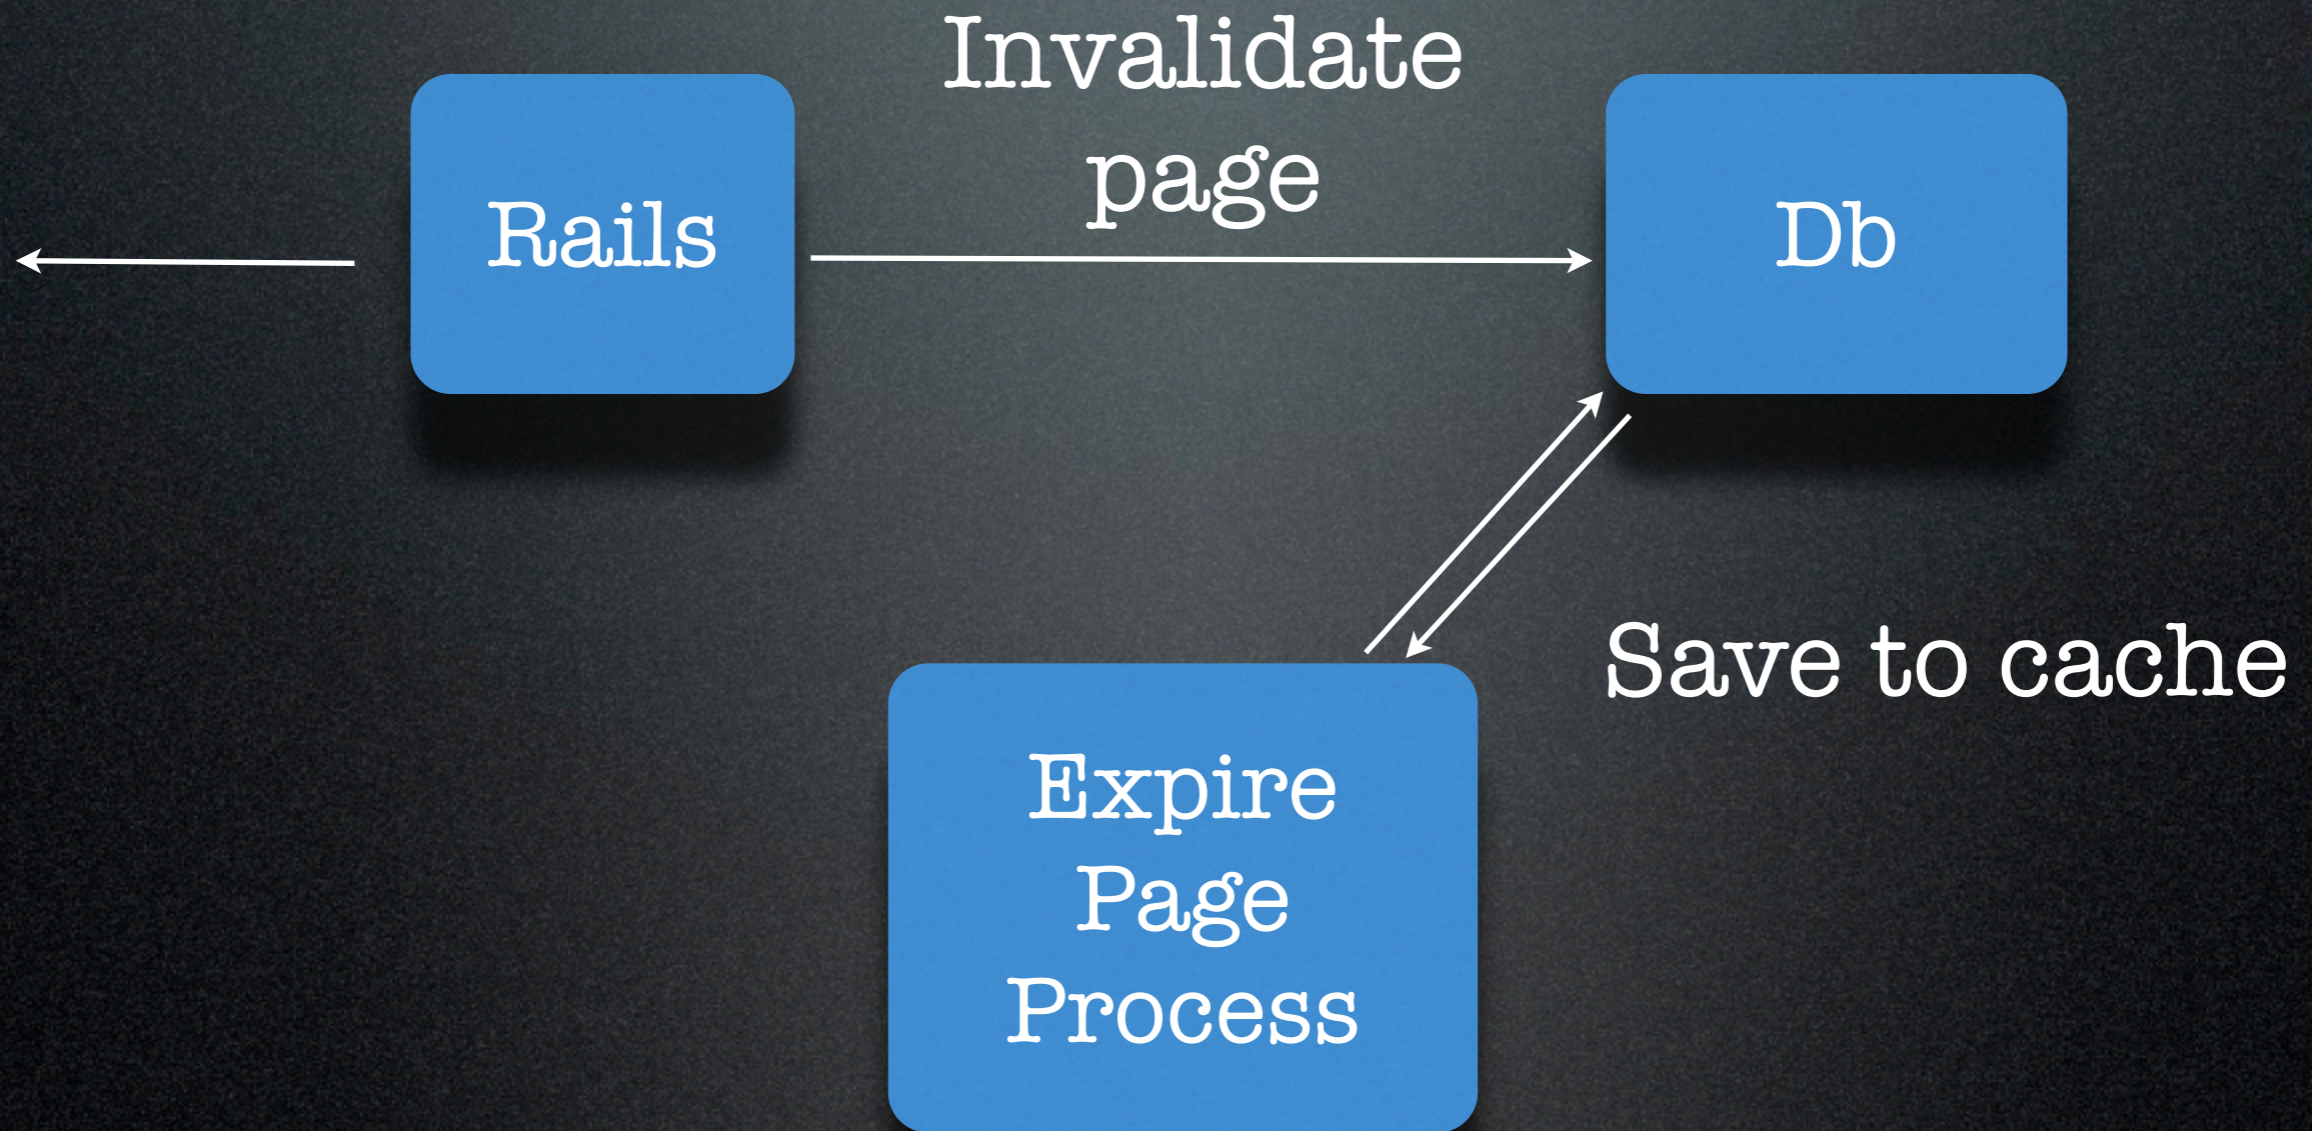
Cache_fu

Invalidation plugin

- Uses the DB for storing list of invalidated pages
- Deletes pages on all servers
- Used at <http://www.linqia.com>

Cache_fu

Invalidation plugin



Action Caching

- Runs filters etc before serving cache
- But it's not all good...what about parameters?

```
/community/index/5?param1=a
```

```
saves as /community/index/5.html
```


Cache_fu

Params plugin

- Takes params into consideration
- Used at <http://www.linqia.com>
- Coming soon to a browser near you



Cache_fu

More for your money



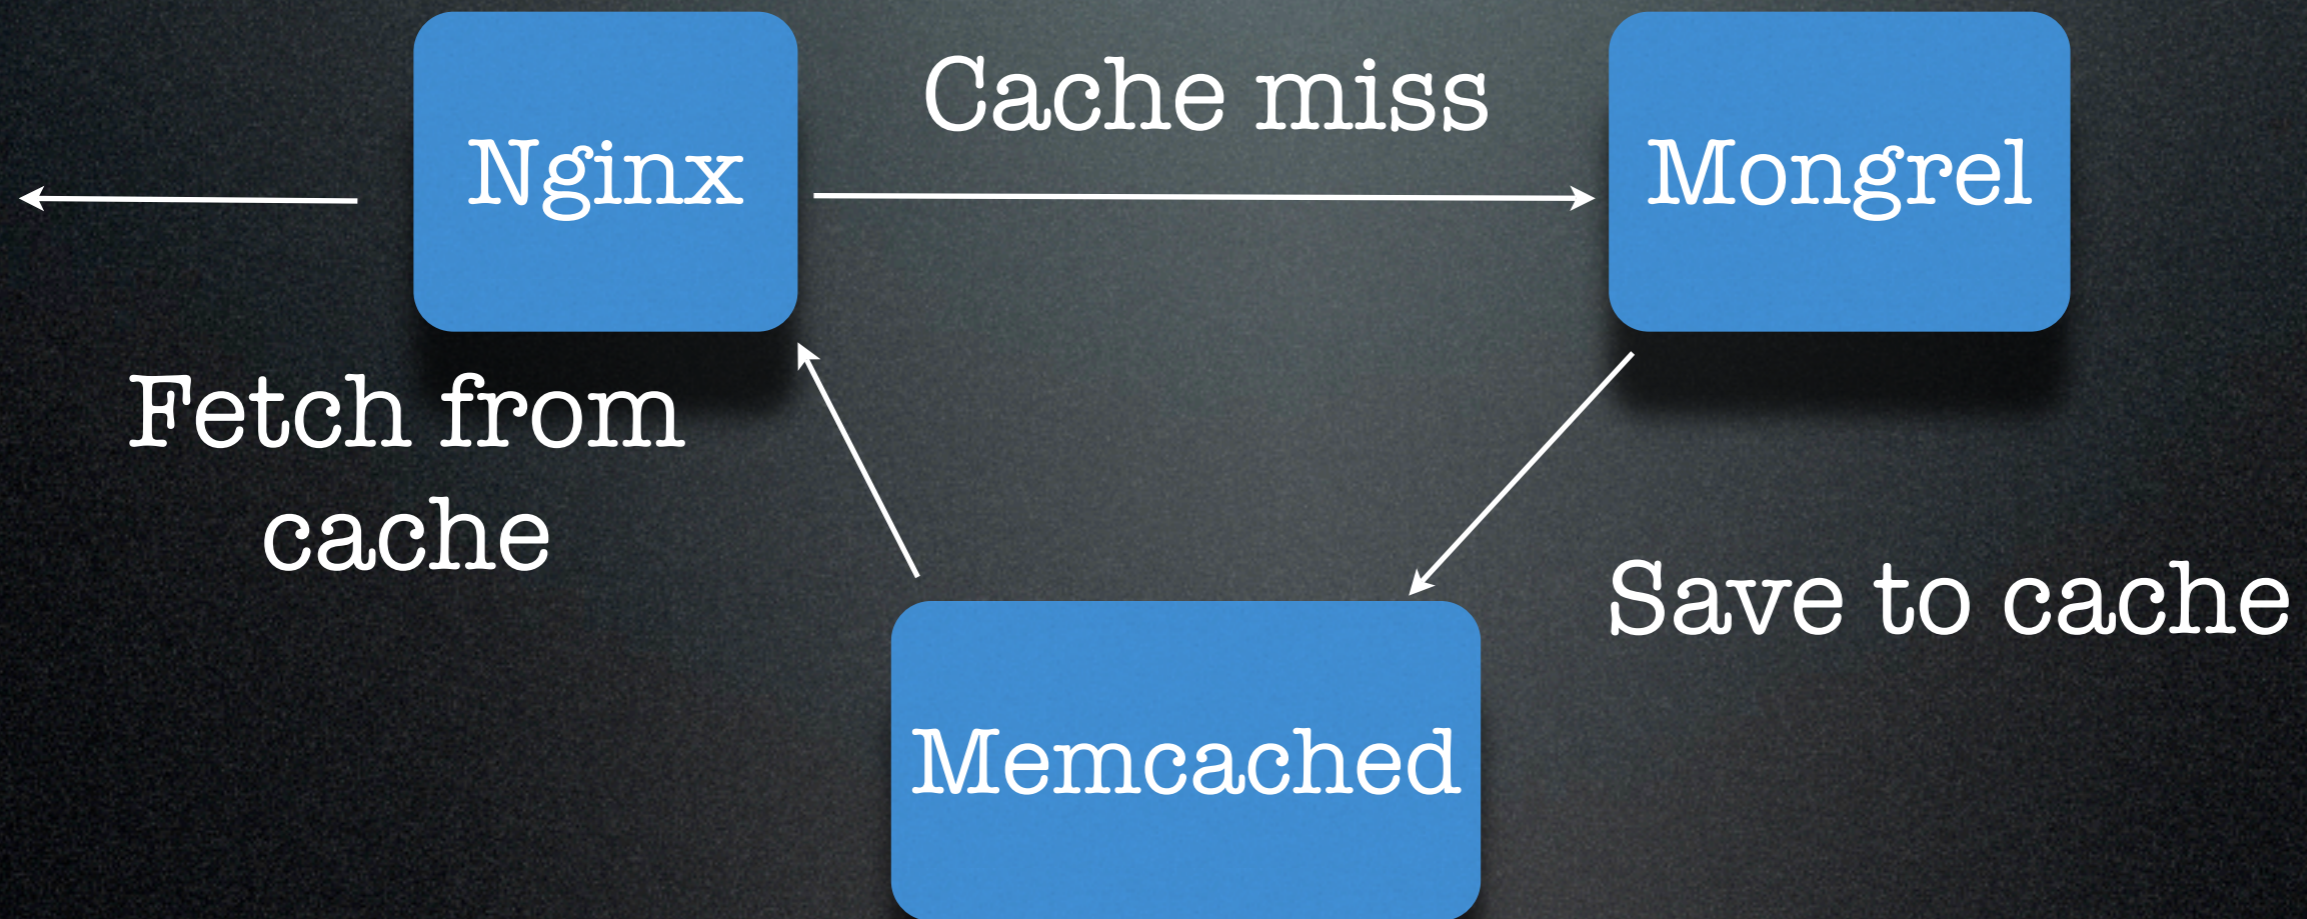
- Mongrel still doing heavy lifting
- Use Evented Mongrel for Gods sake
- Nginx

Strong man seeks server

- Nginx Serving from memcached
 - <http://openhack.ru/nginx-patched>
- Based on params
- Patch for Cache_fu needed
 - No Raw support as far as I can see

Cache_fu

How it works !



Serving Action Cache

- But what about dynamic pages?
- Use cookie to let Nginx serve correct content
- Serve cached page and do dynamic AJAX updates



Cache_fu

Obvious DB Smell

- Any sql “like ‘%faster%’” ?
- Get a search engine (solr, sphinx)
- Using find(:all) ?
- Use it with select, limit and offset

Cache_fu

Learn from Pain



- Need to scale the db
- Read the Mysql story
- <http://www.baselinemag.com/c/a/Projects-Networks-and-Storage/Inside-MySpacecom/>

Cache_fu



Linqia

- All search is GET
- Nginx serving from memcached
- Page Caching for community info
- Dynamic change of cached pages
- Using SOLR for all searching

Cache_fu



Linqia

- Core System hosted servers
- On demand servers from EC2
- S3 for assets
- SQS for asynchronous tasks

Tools

- Railsbench
 - railsbench.rubyforge.org
- Tsung load testing tool
- Firefox + firebug
- YSlow from Yahoo
- Website Optimization
 - www.websiteoptimization.com