

Building Services With webmachine

Kevin Smith



OPSCODE

RULE THE CLOUD



- Infrastructure as code
- Describe server config using Ruby DSL
- Client/Server. Nodes (your servers) run chef-client, talk to Chef server
- It's awesome.



- Merb, Ruby, Unicorn, Nginx
- Stateless, horizontally scalable
- Talks to
 - CouchDB,
 - authorization service (Erlang),
 - Solr



OPSCODE

Typical Chef Server API Request

1. User public key for authentication
2. Node data from CouchDB (median 22K, 3rd Qu. 44K)
3. Authoirzation check
4. **POST, GET, PUT, DELETE**

RULE THE CLOUD



OPSCODE

**How much RAM should it
use?**



$$60 \text{ req/sec} \times 44\text{K} = 2.7\text{MB}$$



**2.7MB data + code +
copies...
27MB?**



OPSCODE

100MB



OPSCODE

Concurrency?
One request per worker.



OPSCODE

204 MB
per unicorn worker



OPSCODE

12 workers per server



OPSCODE

8 servers



OPSCODE

$$12 \times 204 \text{ MB} = 2.4 \text{ GB}$$

$$8 \times 2.4 \text{ GB} =$$

19.2 GB

for pulling JSON out of a database and returning it



OPSCODE

What is webmachine?



- RESTful behavior “engine”
- Graphical Debugger
- Good performance
- Battle tested in commercial products



- Resources
- Callbacks
- Routes



URL




Resource



Init Params



`{"users"}, usermgr_all_users_resource, []}`.





- Corresponds to a single API endpoint
- Implemented as an Erlang module which...
 - Includes `webmachine.hrl`
 - Implements a public `init/1` function
 -at a minimum



- Functions exported by the resource
- Each callback runs at a designated time
- ~ 30 possible callback functions
- Sane defaults automatically used



OPSCODE

Code



- Reduces server-to-client traffic
- Improves performance
- **A REAL PAIN** with webmachine



```
-export([generate_etag/2]).
```

-
-
-

```
generate_etag(Req, #state{data=Data}) ->  
  mochihex:to_hex(crypto:md5(Data)).
```



OPSCODE

Compressed Content

- Reduces traffic
- Faster data transmission times
- Snappier user experience
- **A TOTAL NIGHTMARE** with webmachine

RULE THE CLOUD



```
-export([encodings_provided/2]).
```

```
•  
•  
•
```

```
encodings_provided(Req, State) ->  
  [{“gzip”, fun(X) -> zlib:gzip(X) end},  
   {“identity”, fun(X) -> X end}].
```



OPSCODE

Code



OPSCODE

DRYing Up Endpoints



You write a couple of endpoints and they work great.



OPSCODE

Boilerplate: A Cautionary Tale

Your users are happy.

RULE THE CLOUD



OPSCODE

Boilerplate: A Cautionary Tale

Your boss is happy.

RULE THE CLOUD



OPSCODE

Boilerplate: A Cautionary Tale

You're happy.

RULE THE CLOUD



OPSCODE

Boilerplate: A Cautionary Tale

Life is **GOOD.**

RULE THE CLOUD



OPSCODE

Boilerplate: A Cautionary Tale

Then....

RULE THE CLOUD



OPSCODE

Boilerplate: A Cautionary Tale

The world wakes up and notices your app.

RULE THE CLOUD



User signups increase by an order of magnitude.

API use goes through the roof.



OPSCODE

Boilerplate: A Cautionary Tale

Servers catch fire. Network connections are choked with traffic.

RULE THE CLOUD



OPSCODE

Boilerplate: A Cautionary Tale



RULE THE CLOUD



OPSCODE

Boilerplate: A Cautionary Tale

“We need to monitor API performance”, you say to your team.

RULE THE CLOUD



“We’ll instrument all the endpoints to log to graphite!”



OPSCODE


Boilerplate: A Cautionary Tale

Now you have **N** problems.

RULE THE CLOUD



OPSCODE





- Repetitive
- Copy pasta errors
- Repetitive
- PITA
- Repetitive



- Basic mixins for Erlang
- Reuse entire functions with minimal typing
- No runtime overhead (parse transform)
- Explicit and easy to reason about



OPSCODE

Code



OPSCODE

How did we do?



	Erlang	Ruby
--	---------------	-------------

idle	19MB	100MB
-------------	-------------	--------------

loaded	75MB	204MB
---------------	-------------	--------------




Erlang

Ruby

600MB

19.2GB





OPSCODE

Thank You

Email: k@opscode.com

Twitter: kevsmith

RULE THE CLOUD