Living in a Polyglot World Ruby on the Client, Erlang on the Server

Christopher Brown Chief Technology Officer OPSCODE Twitter: @skeptomai, Email: <u>cb@opscode.com</u>

Early guy, Amazon EC2





Wednesday, January 9, 13



Director of Engineering Microsoft Edge Computing Network



Copyright © 2010 Opscode, Inc - All Rights Reserved

Wednesday, January 9, 13







http://dlutzy.wordpress.com/2011/06/16/velocity-2011-day-1/



Copyright © 2010 Opscode, Inc - All Rights Reserved

Wednesday, January 9, 13

4

Agenda

- What is Chef
 - from an implementor's perspective?
- Why Ruby and Erlang?
- Porting the server from Ruby to Erlang

What is Chef?

"Chef is like a little system admin robot... you tell it how you want your system configured and it will do all the dirty work."

- Probably Interactive



Architecture



- Code Repository
- Chef Server
- Chef Clients
- Data Bags
- Recipes and Cookbooks
- Roles and Run Lists

Recipe

Recipes contain lists of resources evaluated in order

package "haproxy" do action :install end

template "/etc/default/haproxy" do
 source "haproxy-default.erb"
 owner "root"
 group "root"
 mode 0644
 notifies :restart,
"service[haproxy]"
end

service "haproxy" do action [:enable, :start] end

Resource

- Have a type
- Have a name
- Have parameters
- Take action to put the resource in the declared state

```
package "apache2" do
    version "2.2.11-2ubuntu2.6"
    action :install
end
```

```
template "/etc/apache2/
apache2.conf" do
   source "apache2.conf.erb"
   owner "root"
   group "root"
   mode 0644
   action :create
end
```

What is Chef?

... from an implementor's perspective

- Client Side Tools
 - command line
 - convergence engine
- Server API Implementation
 - publishing platform
 - scalable web service
 - fault tolerant

Client Design Goals

- easy to create command line tools
- easy to hack on (we've got a big open source community which is very important to us)
- startup time / resource consumption secondary
- nearly ubiquitous platform support (wide and deep)
- metaprogramming / internal / external DSL

Languages to consider

...on the client

• JVM

- Java
- Clojure
- Scala
- Perl
- OCaml / ML
- Haskell
- Scheme / Lisp

Goals of a DSL



- Translate Domain expertise to language of implementation
- Reduce syntactic / cognitive load

External vs Internal DSL



- Problem: How much of the original language do you carry along?
- What standard language constructs are necessary?
- Are they easy to understand / teach your DSL users?
- Cost of creating & maintaining grammar parser, etc.



http://mitpress.mit.edu/sicp/



Structure and Interpretation of Computer Programs

Harold Abelson and Gerald Jay Sussman with Julie Sussman

Lisp in Small Pieces





Metaprogramming Ruby



All I see are blondes, brunettes and redheads...

	e					ι.														0	1			Č				*			0		V.	0						
	Q	*				۲.		T			0	T			C	5			>		1.			e.				x		C.	ð		e	2	**				Ç,	
	'n					τ.	÷ 7	Ŧ			õ	à	ε		5	á			i i	N.		-		à.				D		2	£	ũ	ĩ	я	8				i.	
÷.	č	6		0		à.	à.	ċ.		•	÷	÷.	Ť			1			2	Ĥ		÷	ē.	ž.				à		2	R.		à.	ö.	F.				1	
2	à	R.		1	÷.		ĩ.	è		M	ĩ.					÷			ĩ	F	Ę.	-						÷.		2	2	2		С.	÷.					
	2	t			Ξ.	2	-	2		ē.	à	2	r.	ĕ	- 2			-		2		М	1	2				2		Τ.	2	2	5	2					ř.	
	2	γ			2	Ξ.		۰.		2	×.	5	4	2				2	12		2	1	1	Ξ.						2		4	Υ.	5					Ŀ	
	ę.	2		1	2			2		ш.	ų.	•	2	1		2		2		-11	7	2	1		2.					2	2	1	9	9						
	ł.	n		Ϋ́.	2	2	Υ.			2	ň	÷	2	Ï.		U.		2	2	α.	4	2	-		1			Ţ		•	2	τ.	-	9					9	
	÷	Ľ		0	1	<u>×</u>	Σ.	2			9	σ.	Σ	2		1		E.			4	2						1		Y	Ξ	1	÷	2	٤.					
4	ē,	0		н	Т		Q,			X	H	-	н	ō				-10			E	Н	0					0		Ξ.	٩Ū	×	×	1					3	
	Ÿ.			4	1.		Ø.			đ	7	2		н	2			2	0		-	T	ę.	ζ.	64					2	Ø.	>	<u>L</u>	7					0.1	
	8	£		2	Ħ.	5	e			3	ŧ.	×	>	ŧ.	1 4			2	+		÷	1	1 0	2	ζ (ę	T	٠	х	2	κ.	ε,				Q (
	ε	х			>		0			I.		*	11	7	C.		ε	4	2			ε	U	۰.	74					Ŧ.	0	ę		۴	2			>	ŧ.	
ž	5	ð	>		ð	ę.	в			1	Ŧ	L.	L.		¢		>	Y	*			٠	1	8	EII					>		6		1	6	Ċ,	2	2	۲.1	Ċ,
2	۲	9	7			0	3			2	7	х	ð	3	65		14	*	6			\mathbf{A}	1	41	40	9.6	Ŧ		Ve	T	5	>		¥.,	1	>	3	2		
	x				\mathbf{A}	£.	ŧ.			U	7	0		A	111		5		0			0		3	64		0		t.	x	¢.	2		¢,	н	20	2	0	t: •	6
	г	x	0		20	۰.	A				à	14	2	ŧ.	œ 🖗		ι	L.	ŧ			0		Ú.	1 2		2		7	*		-		A	r	ε	-		t	1
1	k				a	λī i	11				x	t	0		0		-					Ū	ch-		1 2				ĩ.			Ŧ		x	*	5		τ		
1			1		T	0			7		1	÷	c.	÷	č.		i.		ς.			20	a						*	÷.	1.	ŵ.		20		Ť		T		
			6		5	Ă	-		n		5	÷	Ŧ	-								-		0	26					i.		0				x	-	÷		
	÷		6			F			Ř		4	ú		£.				F.				5					Ŧ		ŝ.		1	5			ă,	æ			ĥ	
2	à.				Ξ.	ă.			e		à		à	è			-					ē					-		6	ς.	÷	ź		ę.	3					
	5		÷.		2	1					2	2	ř.					1				2			- 1				ā	à.	÷	4		2		1	3			
	2		2		â.	2			ų.		S		4	5			5	4		5		4		2					2		2	2		Δ.		2	$\widehat{}$			
	Ξ		à		ę.	2			ł.		2	S.	2	g.				2		2		2		2			12		2		2			γ.		S.				
	÷				^	ç.			2		2	2		2				Ч				9		5		9	4		2	7	2	e.		2		2				
	۲.		ų			Ļ.						2				X		H		7		2				14	đ		н		2	÷		c.		2	\geq			
	1		đ			1			1		1	٦		2			2	ć		1		٤.								0	Ó	7		V		2				
						1					Ŧ	4		2		đ	7			C		t.				1				2	ſ	0				V				
	8		>			11			ł		8	1		Ó		1	C)					40								8	2	T		ę		V	ы			
	¢,		ę			¢,			Ċ.		Ц	1		£.		9	X					2		0		11		2		2	>	1		7		2	L1			
١.	4	E	Ш						L		>	7		1		4	t			÷		*						ε		11	9	>		ŧ.		9	0			
	r	х	X	X		11			11		ę	Ŀ		T		T	7			11	2	+		T		+		0		é.	Ŀ	2		ę		ſ				
2	3	L	X			7					H	V		2		0	*			ę	Ŧ	4		ð.		7		9		2	V.	8		1	2	0				
																																					A COLUMN A			

Internal DSL

- Exploit Metaprogramming
 - Heritage of Lisp, Smalltalk, Ruby
 - Requires Flexible syntax
- May inherit types, data structures

External DSL

- You are responsible for grammar
 - Lex / Yacc / ANTLR
- You are responsible for types, data structures and execution



"I think I realized a long time ago that Java is not a good enough language to implement applications."

http://ola-bini.blogspot.com/2008/01/language-explorations.html



XML?

```
<?xml version="1.0" encoding="UTF-8"?>
<!-- generated by N1 SPS -->
<component label='1.0.0' xmlns='http://www.sun.com/schema/SPS' name='tuxedo81.cont' version='5.1' description='For deploying</pre>
tuxedo81 domain for ocsccm 1.0.0' xmlns:xsi='http://www.w3.org/2001/XMLSchema-instance' path='/ocs/ccm/1.0'
xsi:schemaLocation='http://www.sun.com/schema/SPS component.xsd'>
    <extends>
         <type name='system#container'></type>
    </extends>
    <varList>
         <var name='execUser' default=':[target:execUser]'></var>
         <var name='execGroup' default=':[target:execGroup]'></var>
         <var name='installRoot' default=':[target:installRoot]'></var>
         <var name='installPath' default=':[target:installRoot]'></var>
         <var name='platform' default=':[target:platform]'></var>
         <var name='platformPath' default=':[installRoot]/:[platform]'></var>
         <var name='environment' default=':[target:environment]'></var>
         <var name='environmentPath' default=':[platformPath]/:[environment]'></var>
    </varList>
    <componentRefList>
         <componentRef name='deploy_tuxedo81_dir'>
              <argList installPath=':[installRoot]' installGroup=':[execGroup]' installDiffDeploy='FALSE' installPermissions='754'</pre>
installName='tuxedo81' installUser=':[execUser]' installDeployMode='REPLACE'></argList>
              <component name='tuxedo81' path='/ocs/ccm/1.0/components' version='1.2'></component>
         </componentRef>
    </componentRefList>
    <installList>
         <installSteps name='default'>
              <install blockName='default'>
                   <allNestedRefs></allNestedRefs>
              </install>
              <execNative userToRunAs='root'>
                   <inputText><![CDATA[</pre>
#cd :[installRoot]
#find :[platform] -type d -print | xargs chown -R :[execUser]::[execGroup]
        ]]></inputText>
                   <exec cmd='sh'></exec>
              </execNative>
         </installSteps>
    </installList>
</component>
```

Erlang?

```
#!/usr/bin/env escript
%% -*- erlang -*-
%%! -smp enable -sname factorial -mnesia debug verbose
main([String]) ->
    try
        N = list_to_integer(String),
        F = fac(N),
        io:format("factorial \sim w = \sim w \setminus n", [N,F])
    catch
        _:_ ->
             usage()
    end;
main( ) ->
    usage().
usage() ->
    io:format("usage: factorial integer\n"),
    halt(1).
```

Perl?

```
use Chef;
resource file => '/tmp/foo', sub {
  my $r = shift;
  $r->owner('adam');
  $r->action('create');
};
```

Common Lisp?

```
(require 'chef)
(resource :file "/tmp/foo")
          :owner "cb"
          :action :create)
(resource :file (concatenate 'string
                              "/tmp/"
                               (node-attributes :hostname)
                              "-made-with-lisp")
          :action :create)
;; Or
(file :path "/tmp/foo"
      :owner "cb"
      :action :create)
(file :path (concatenate 'string
                          "/tmp/"
                          (node-attributes :hostname)
                          "-made-with-lisp")
      :action :create)
```

Wednesday, January 9, 13

Server Design Goals

- Horizontally scalable
- Highly concurrent
- Stateless request handling
- Fault tolerant

Languages to consider

...on the server

- JVM
 - Java
 - Clojure
 - Scala
- Perl
- OCaml / ML
- Haskell
- Scheme / Lisp

Chef Server API

CouchDB

- Merb, Ruby, Unicorn, Nginx
- Stateless, horizontally scalable
- Talks to



- CouchDB,
- authorization service (Erlang),





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

Average Chef Server API Response Times

500 ms



stats.timers.chefAPI.application.upstreamRequests.authenticate_every.mean

stats.timers.chefAPI.application.upstreamRequests.check_rights.mean

stats.timers.chefAPI.application.upstreamRequests.databaseCalltimePerReq.mean

stats.timers.chefAPI.application.alRequests.mean

stats.timers.erchefAPI.application.allRequests.mean

Nervous sweat



Slow, Irregular, and Out of Control

CouchDB Uptime





Heavy on system resources

How much RAM should it use?

60 req/sec × 44K = 2.7MB

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

IOOMB

at rest, after startup

Concurrency? One request per worker.
204 MB per unicorn worker

under load

Wednesday, January 9, 13

12 workers per server

8 servers

$12 \times 204 \text{ MB} = 2.4$ GB 8 × 2.4 GB = 9,2 G B

for pulling JSON out of a database and returning it

Wednesday, January 9, 13

Unicorns Eat RAM







How It Works Quickstart Demo Dispatching Request Resource

Streamed Body

Debugging

low Diagram

Referential Transparency

Home > Webmachine

REST Toolkit

Webmachine is not like the web frame in Erlang, Webmachine is an application management provided by <u>mochiweb</u>, a

Resources

A Webmachine application is a set of r object-methods, infinite-server-loops, are relatively easy to understand and

These functions give you a place to de that the first-class things on the Web a constrained.



How did we do?

Ruby Erlang idle 19MB 100MB loaded 75MB 204MB

Erlang

Ruby

600MB 19.2GB

Wednesday, January 9, 13

We win!



But wait! There's more.

Where is Ruby API spending time?

DB calls?

JSON parsing/ rendering?



Garbage Collection?

Garbage Collection!



>40% CPU in GC



CPU Usage on Chef Server



Frequent GET/ PUT of node JSON



No concurrency accessing a single database

(until recently)

Motivation: Why not CouchDB?

Database replication unreliable for 1000s of databases.

Motivation: Why not CouchDB?

File handle and memory resource leaks

Motivation: Why not CouchDB?

WORKEDENEIN **OPS PROBLEM NOW**

It became an operations "thing"

What we need in a data store

- Happy with write heavy load
- Support for sophisticated queries
- Able to run HA

Did you consider NoSQL database X?

Yes, but we also asked: Why not SQL?

Measure!



Elapsed Secs

Table of Contents

- Testing Node Operations in MySQL
 - Test system
 - Insert 40K nodes test
 - insert 80 workers, no compression
 - insert 80 workers, gzip
 - insert 80 workers "no blobs"
 - Insert, Fetch, Update tests
 - all ops 80 workers, no compression
 - all ops 80 workers, gzip
 - all ops 80 workers "no blobs"
 - Software Used

Testing Node Operations in MySQL

We tested two scenarios: insertion only and a combined insert, fetch, and update test that was run after the insertion test with a base of twice current production node count).

Tests were performed using basho_bench with custom drivers and generators (see <u>Software Used</u> below). The back-end Private C was used to host the database. The load generation was run on the front-end Private Chef PoC box (172.28.5.24).

Test system

Private Chef PoC boxes. Test client on front-end, MySQL on back-end.

MySQL installed from Percona's apt repo:

mysql Ver 14.14 Distrib 5.5.15, for Linux (x86_64) using readline 5.1

This was installed from Percona as:

percona-server-server-5.5/lucid uptodate 5.5.15-rel21.0-158.lucid

Insert 40K nodes test

The insert test was run for five minutes or until 40K nodes were added to the database.

Generated nodes are randomly assigned to one of 500 orgs and one of four environments. Node names and IDs are randomly generate

basho_bench





So we replaced CouchDB with MySQL

...while the system was running



Live Migration: Starts out easy!

Live Migration in 3 Easy Steps

1.Put org into read-only mode

2.Copy from CouchDB to MySQL

3.Route org to Erchef
It Gets Harder

Real World Hard



Wednesday, January 9, 13

Migration Tool

- 1. Coordinate feature flippers and load balancer config
- 2. Move batches of orgs through migration
- 3. Track status of migration and individual orgs
- 4. Resume after crash



Migration Tool

- 1. Track inflight write requests
- 2. Put org into read-only mode
- 3. Wait for inflight write requests to complete
- 4. Migrate org data
- 5. Reconfig/HUP load balancer
- 6. Handle errors



Scripting with gen_fsm

- Helper methods \rightarrow states
- Server state and supervision tree make crash recovery easier
- Free REPL

OTP + gen_fsm =:= Happy Migration Tool

Organization Robustness

state functions		
state record	\checkmark	
manager/worker processes		
supervision tree		

DETS local store

No migration plan survives contact with production

http://en.wikiquote.org/wiki/Helmuth von Moltke the Elder

Wednesday, January 9, 13

Database CPU CouchDB MySQL





Database Load Average CouchDB MySQL





API Average Latency



Chef Server Roles Endpoint 90th Latency









CouchDBWrite Requests



CouchDB Network Traffic



Network traffic on Chef Server





The Road Ahead?



Copyright © 2010 Opscode, Inc - All Rights Reserved

Wednesday, January 9, 13

88



Christopher Brown Chief Technical Officer OPSCODE Twitter: @skeptomai, Email: <u>cb@opscode.com</u>