

Campfire Loves Erlang

Mark Imbriaco
mark@37signals.com



What is Campfire?



What is Campfire?

- Web based group chat service.



What is Campfire?

- Web based group chat service.
- A few thousand simultaneous users, and growing rapidly.



What is Campfire?

- Web based group chat service.
- A few thousand simultaneous users, and growing rapidly.
- Primarily a Ruby on Rails application backed by MySQL, but ...



What is Campfire?

- Web based group chat service.
- A few thousand simultaneous users, and growing rapidly.
- Primarily a Ruby on Rails application backed by MySQL, but ...
- Part of it was recently rewritten in Erlang.



37signals

[Create a new room](#)

4 people currently chatting

All Talk

Active less than a minute ago

"All I want to know is where I'm going to die so I'll never go there." -Charles Munger

Jason Fried
Jason Zimdars
Mark Imbriaco
Matt Linderman

Small talk

Unoccupied

It was 1980 29 years ago.

Backpack History

Unoccupied

Discussion about Backpack, screenshots, historical decisions, feature discussion.

Highrise History

Unoccupied

System Administration

Unoccupied

Notes and details about the server cluster. Post significant updates [here](#).

Fireside Chat

Unoccupied

Interview room.





12:45 PM

Sam S. [perfection.png](#)

In our time, many of us have been taught to strive for an insane perfection that means nothing. To get wholeness, you must try instead to strive for this kind of perfection, where things that don't matter are left rough and unimportant, and the things that really matter are given deep attention. This is a perfection that seems imperfect. But it is a far deeper thing. (Alexander 1991, emphasis in original)

1:00 PM

Jason Z. has entered the room

Send message

All Talk

"All I want to know is where I'm going to die so I'll never go there." -Charles Munger
[Edit](#)

 Search

Who's here? [Lock room](#) | [Leave](#)

- Jason Fried
- Jason Zimdars
- Mark Imbriaco
- Matt Linderman

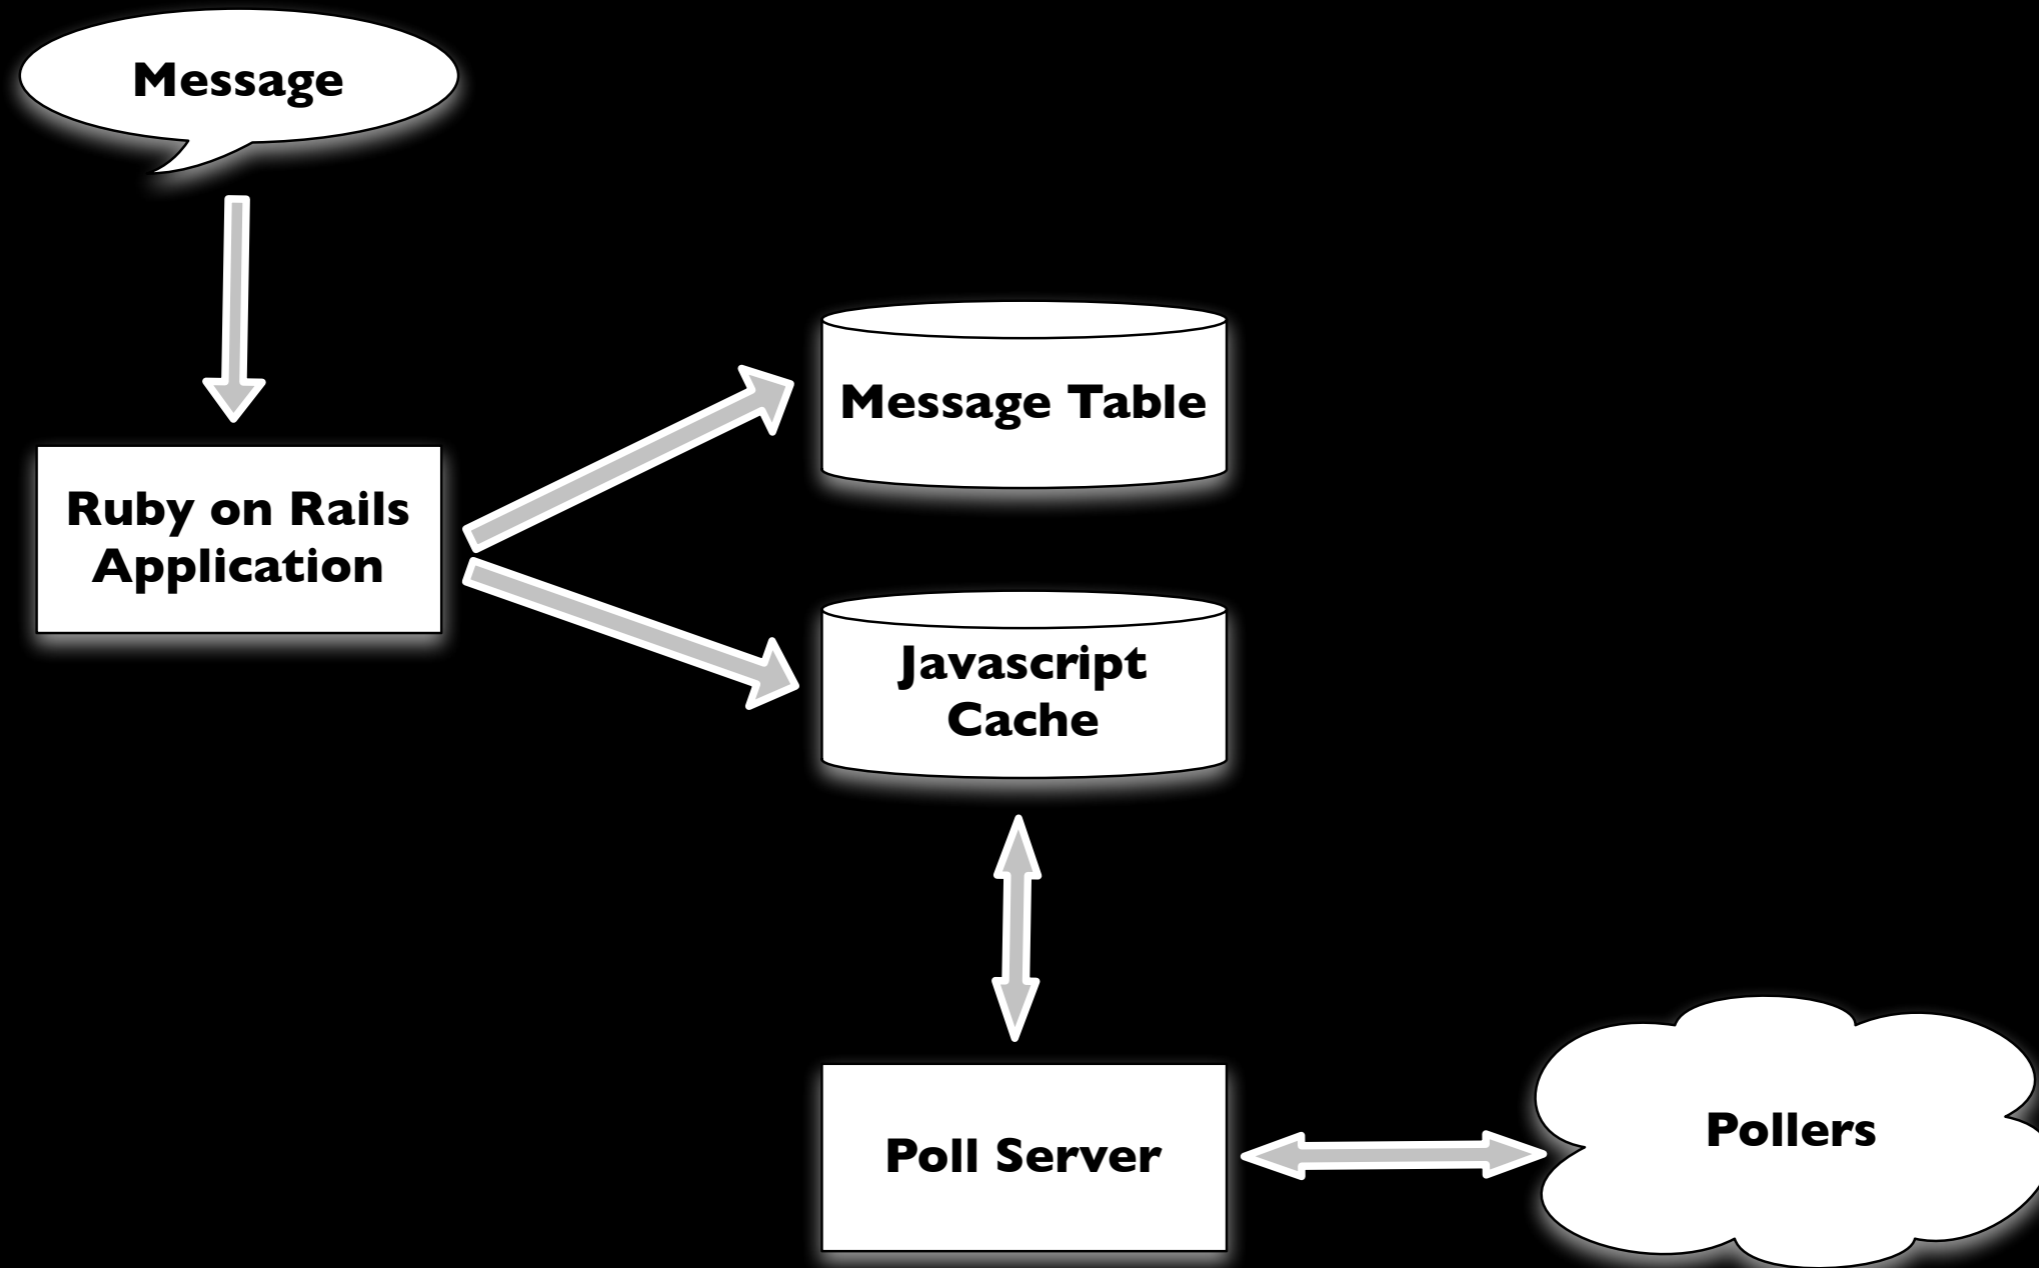
Guest access is off (?) [Turn it on](#)

Latest 5 files (all) [Upload a file](#)

- [perfection.png](#)
- [house of tiles.png](#)
- [FancyZoom loading.mov](#)
- [wrap.png](#)
- [cover6.png](#)



Architecture





Erlang!

12:45 PM

Sam S. [perfection.png](#)

In our time, many of us have been taught to strive for an insane perfection that means nothing. To get wholeness, you must try instead to strive for this kind of perfection, where things that don't matter are left rough and unimportant, and the things that really matter are given deep attention. This is a perfection that seems imperfect. But it is a far deeper thing. (Alexander 1991, emphasis in original)

1:00 PM

Jason Z. has entered the room

Send message

All Talk

"All I want to know is where I'm going to die so I'll never go there." -Charles Munger
[Edit](#)

 Search

Who's here? [Lock room](#) | [Leave](#)

- Jason Fried
- Jason Zimdars
- Mark Imbriaco
- Matt Linderman

Guest access is off (?) [Turn it on](#)

Latest 5 files (all) [Upload a file](#)

- [perfection.png](#)
- [house of tiles.png](#)
- [FancyZoom loading.mov](#)
- [wrap.png](#)
- [cover6.png](#)



Poll Process



Poll Process

- Every 3 seconds.



Poll Process

- Every 3 seconds.
- Client authenticates itself and sends the ID of the last message it has seen.



Poll Process

- Every 3 seconds.
- Client authenticates itself and sends the ID of the last message it has seen.
- Server responds with the contents of Javascript cache table.



Performance Requirements



Performance Requirements

- Normal daily traffic around 1500 requests per second.



Performance Requirements

- Normal daily traffic around 1500 requests per second.
- Spontaneous synchronization.



Performance Requirements

- Normal daily traffic around 1500 requests per second.
- Spontaneous synchronization.
- Memory and CPU footprint.



Timeline



Timeline

- First version was Ruby based, never made it to production.



Timeline

- First version was Ruby based, never made it to production.
- Next, C based poll server using the FastCGI specification.



Timeline

- First version was Ruby based, never made it to production.
- Next, C based poll server using the FastCGI specification.
- Ruby Mongrel handler prototype.



Timeline

- First version was Ruby based, never made it to production.
- Next, C based poll server using the FastCGI specification.
- Ruby Mongrel handler prototype.
- Erlang poll server.



Ruby Servers



Ruby Servers

- Small, only 127 lines of code.



Ruby Servers

- Small, only 127 lines of code.
- Clean, nicely factored code.



Ruby Servers

- Small, only 127 lines of code.
- Clean, nicely factored code.
- Great for developing against.



Ruby Servers

- Small, only 127 lines of code.
- Clean, nicely factored code.
- Great for developing against.
- Throughput and CPU/memory consumption make it unsuitable for production.



C FastCGI Server



C FastCGI Server

- 397 lines of code.



C FastCGI Server

- 397 lines of code.
- Fast, very low memory and CPU footprint.



C FastCGI Server

- 397 lines of code.
- Fast, very low memory and CPU footprint.
- Not designed to be extensible.



C FastCGI Server

- 397 lines of code.
- Fast, very low memory and CPU footprint.
- Not designed to be extensible.
- Needs one OS process for each simultaneous request. *



Story Time



Erlang Server



Erlang Server

- 273 lines of code, 63 of them for logging.



Erlang Server

- 273 lines of code, 63 of them for logging.
- Makes use of existing battle tested code like Mochiweb and OTP behaviors.



Erlang Server

- 273 lines of code, 63 of them for logging.
- Makes use of existing battle tested code like Mochiweb and OTP behaviors.
- Modular and extensible.



Erlang Server

- 273 lines of code, 63 of them for logging.
- Makes use of existing battle tested code like Mochiweb and OTP behaviors.
- Modular and extensible.
- As fast as the C version.



Erlang Server

- 273 lines of code, 63 of them for logging.
- Makes use of existing battle tested code like Mochiweb and OTP behaviors.
- Modular and extensible.
- As fast as the C version.
- Handles spontaneous synchronization well, thanks to Erlang processes. *



Recap

	Ruby	C	Erlang
LOC	127	397	273
Req/sec	250-350	1800	1800
Response Time	20ms	2-3ms	2-3ms
OS Processes	n/a	80	1
Extensible	Yes	No	Yes



Operational Details



Operational Details

- Not much information out there about how exactly to host Erlang applications.



Operational Details

- Not much information out there about how exactly to host Erlang applications.
- Settled on `run_erl ...` and later added `runit`.



Operational Details

- Not much information out there about how exactly to host Erlang applications.
- Settled on `run_erl ...` and later added `runit`.
- Three nodes behind HAproxy.



Operational Details

- Not much information out there about how exactly to host Erlang applications.
- Settled on `run_erl ...` and later added `runit`.
- Three nodes behind HAproxy.
- Use Capistrano to deploy, no hot code reload.



Memory Consumption



Memory Consumption

- Processes climbing to ≥ 1 GB resident.



Memory Consumption

- Processes climbing to \geq 1GB resident.
- Lots of short lived binaries and strings floating around, particularly in the MySQL processes.



Memory Consumption

- Processes climbing to \geq 1GB resident.
- Lots of short lived binaries and strings floating around, particularly in the MySQL processes.
- `erlang:system_flag(fullsweep_after, 0)`



Memory Consumption

- Processes climbing to \geq 1GB resident.
- Lots of short lived binaries and strings floating around, particularly in the MySQL processes.
- `erlang:system_flag(fullsweep_after, 0)`
- Memory usage now hovers around 150-200MB and doesn't exceed 250MB.



runit vs. Heart



runit vs. Heart

- Heart is great.



runit vs. Heart

- Heart is great.
- But what happens when the VM runs out of memory?



runit vs. Heart

- Heart is great.
- But what happens when the VM runs out of memory?
- Had existing runit tooling and automation.



runit vs. Heart

- Heart is great.
- But what happens when the VM runs out of memory?
- Had existing runit tooling and automation.

<http://smarden.sunsite.dk/runit/>



What's next?



What's next?

- Experimenting with Comet.



What's next?

- Experimenting with Comet.
- CouchDB (via Chef)



What's next?

- Experimenting with Comet.
- CouchDB (via Chef)
- RabbitMQ



What's next?

- Experimenting with Comet.
- CouchDB (via Chef)
- RabbitMQ
- Who knows?



Questions?

