

A Change in Protocol

Exploring XMPP in Ruby

Brian Dainton

brian.dainton@gmail.com

brian@fiveruns.com

twitter.com/bdainton

Certainly, as web developers, the protocol we use most frequently is HTTP.

HTTP



HyperText Transfer Protocol

We're all familiar w/HTTP network chatter picture.

HTTP



client



server

A client....

Opens a connection

Makes a request

server does some work

Sends a response

And the connection is closed.

Image sources:

http://www.flickr.com/photos/bakelite_phones

<http://manalang.com/archives/2004/09/19/the-home-computer-of-2004/>

XMPP


XML Messages Persistent Connection

XMPP is proving itself more and more to be a complementary protocol on the web today. This makes our network chatter picture look quite a bit different:

XMPP



client



server

A client...

Opens a connection

xml messages are sent back and forth

And then the connection is closed. [click]

This capability to send messages over a persistent connection means that we can do some really neat things with XMPP that we can't do (or do well, anyway) with HTTP.

Announcing

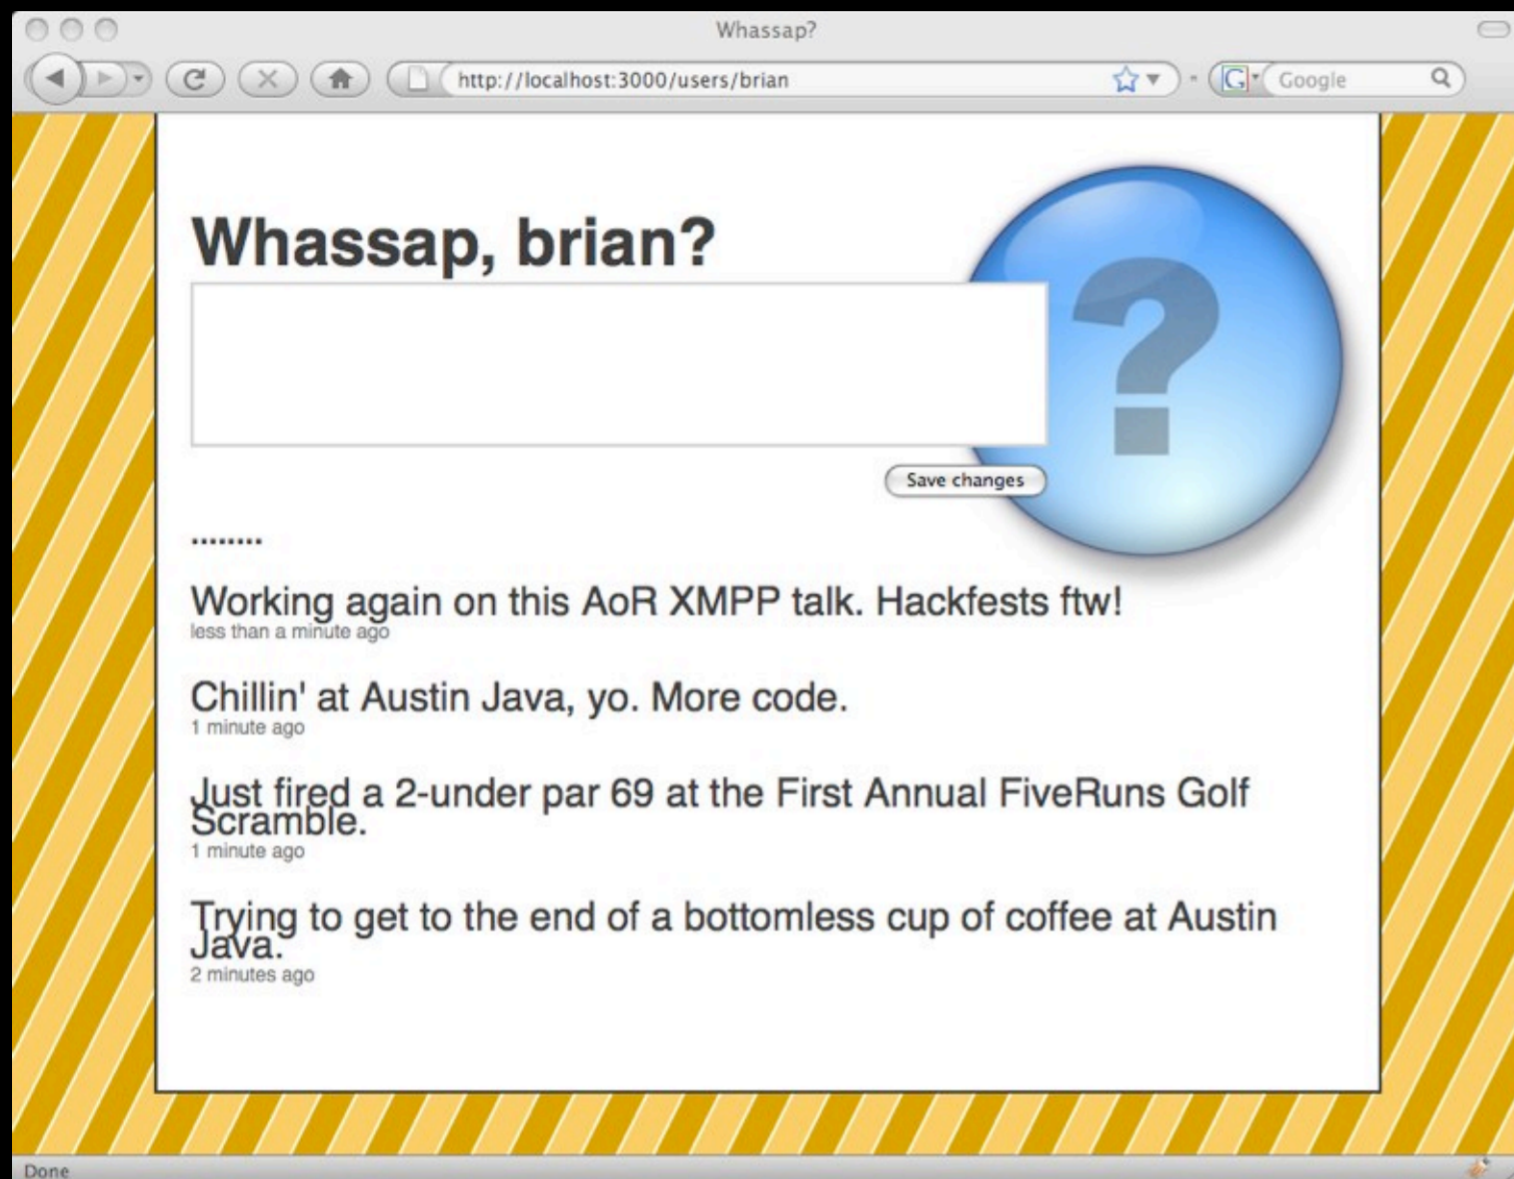


I'm very happy to announce my latest social networking app...

Image source:

<http://258marketing.wordpress.com/category/advertising-campaign/>

Whassap



Like Twitter.
No following.

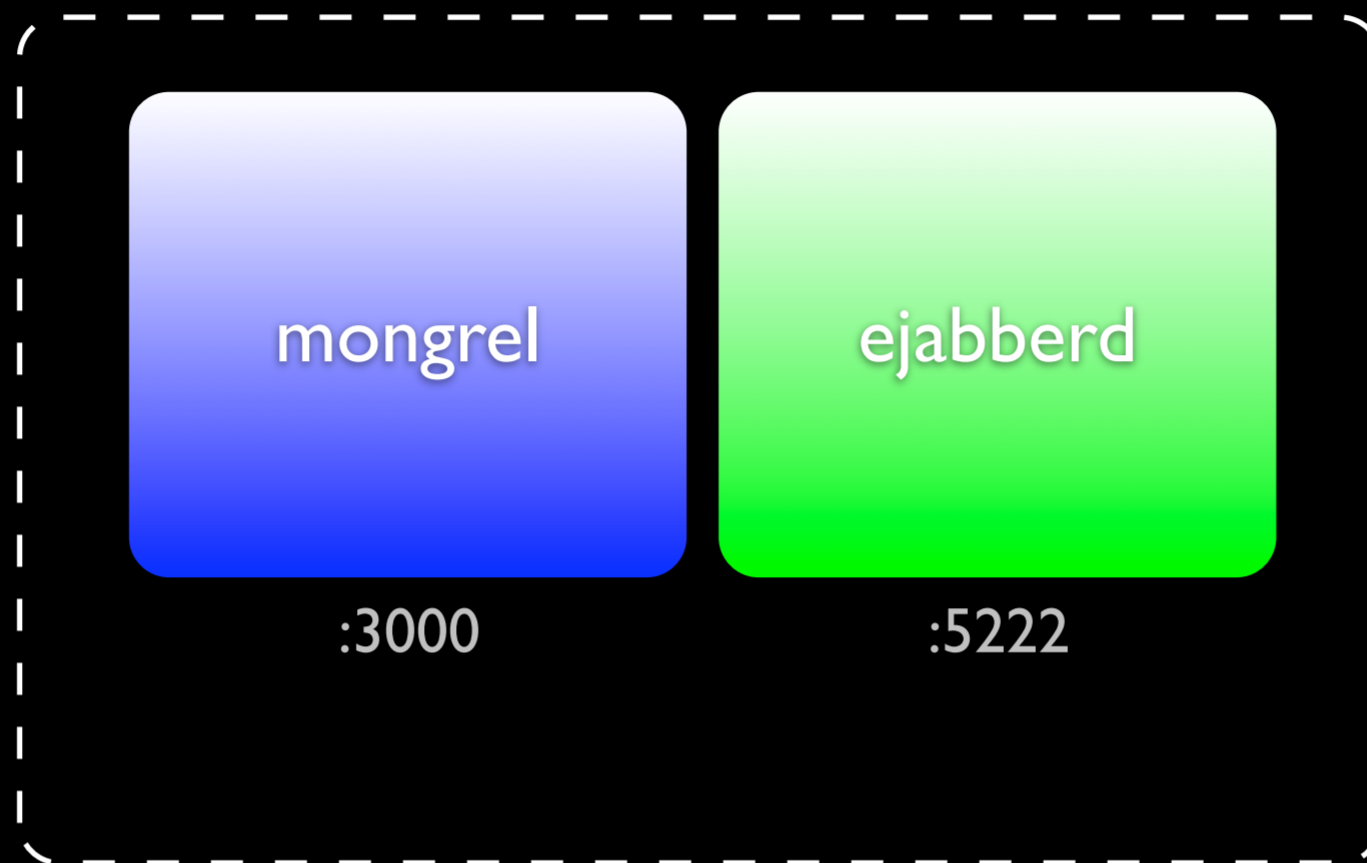
Polite IM.

Whassap.

Like Twitter, but without following.

New key feature: Whassap IMs you to thank you for setting your status.


Whassap



tiamat.local

Here's the app-level architecture for delivering the Whassap service. I've got a Rails app served by mongrels on port 3000 and an XMPP server (ejabberd) accepting connections on port 5222.

What does XMPP have to do with IM? We need a history lesson.



1998



1998. Internet boom. Top hits of the day.

Image sources:

http://www.amazon.com/Titanic-Music-Motion-Picture/dp/B0000029YC/ref=sr_1_1?ie=UTF8&s=music&qid=1222457439&sr=1-1

http://www.amazon.com/Backstreet-Boys-ENHANCED-CD/dp/B00000053X/ref=sr_1_5?ie=UTF8&s=music&qid=1222457465&sr=1-5

http://www.amazon.com/Come-Over-International-Shania-Twain/dp/B000031VR1/ref=pd_bbs_sr_2?ie=UTF8&s=music&qid=1222457317&sr=8-2

http://www.amazon.com/Genie-Gets-Wish-Christina-Aguilera/dp/B000IFSGUI/ref=sr_1_26?ie=UTF8&s=music&qid=1222457523&sr=1-26

http://www.amazon.com/Lets-Talk-About-Love-Celine/dp/B000002C4J/ref=sr_1_1?ie=UTF8&s=music&qid=1222457564&sr=1-1

http://www.amazon.com/Baby-More-Time-Britney-Spears/dp/B00009YNG9/ref=sr_1_2?ie=UTF8&s=music&qid=1222457593&sr=1-2



| 1998



We were also Instant Messaging our friends and work colleagues.
There were problems. Vendor lock-in. No cross-protocol communication.
This sucked.
Enter Jabber.

Image sources:

<http://webmessenger.msn.com/>

<http://www.icq.com/>

http://www.makeuseof.com/wp-content/uploads/2008/06/aim_logo.jpeg

<http://www.intuitive.com/blog/images/goofy-yahoo-logo.gif>



Jabber

jabberd
protocol

Jabber Project focused on creating an open platform for instant messaging.
Jabber protocol would later for the basis for...

Image source:
<http://florianjensen.com/2008/02/11/jabber-and-data-portability/>

XMPP

Extensible Messaging and Presence Protocol

XMPP

Extensible **Messaging**
and
Presence Protocol

XMPP

Extensible Messaging and Presence Protocol

Presence info: available, offline, in a meeting, etc.

XMPP

Extensible Messaging and Presence Protocol

XML: inherently extensible.

Indeed, there have been many many extensions both proposed and accepted to the XMPP spec.

XMPP

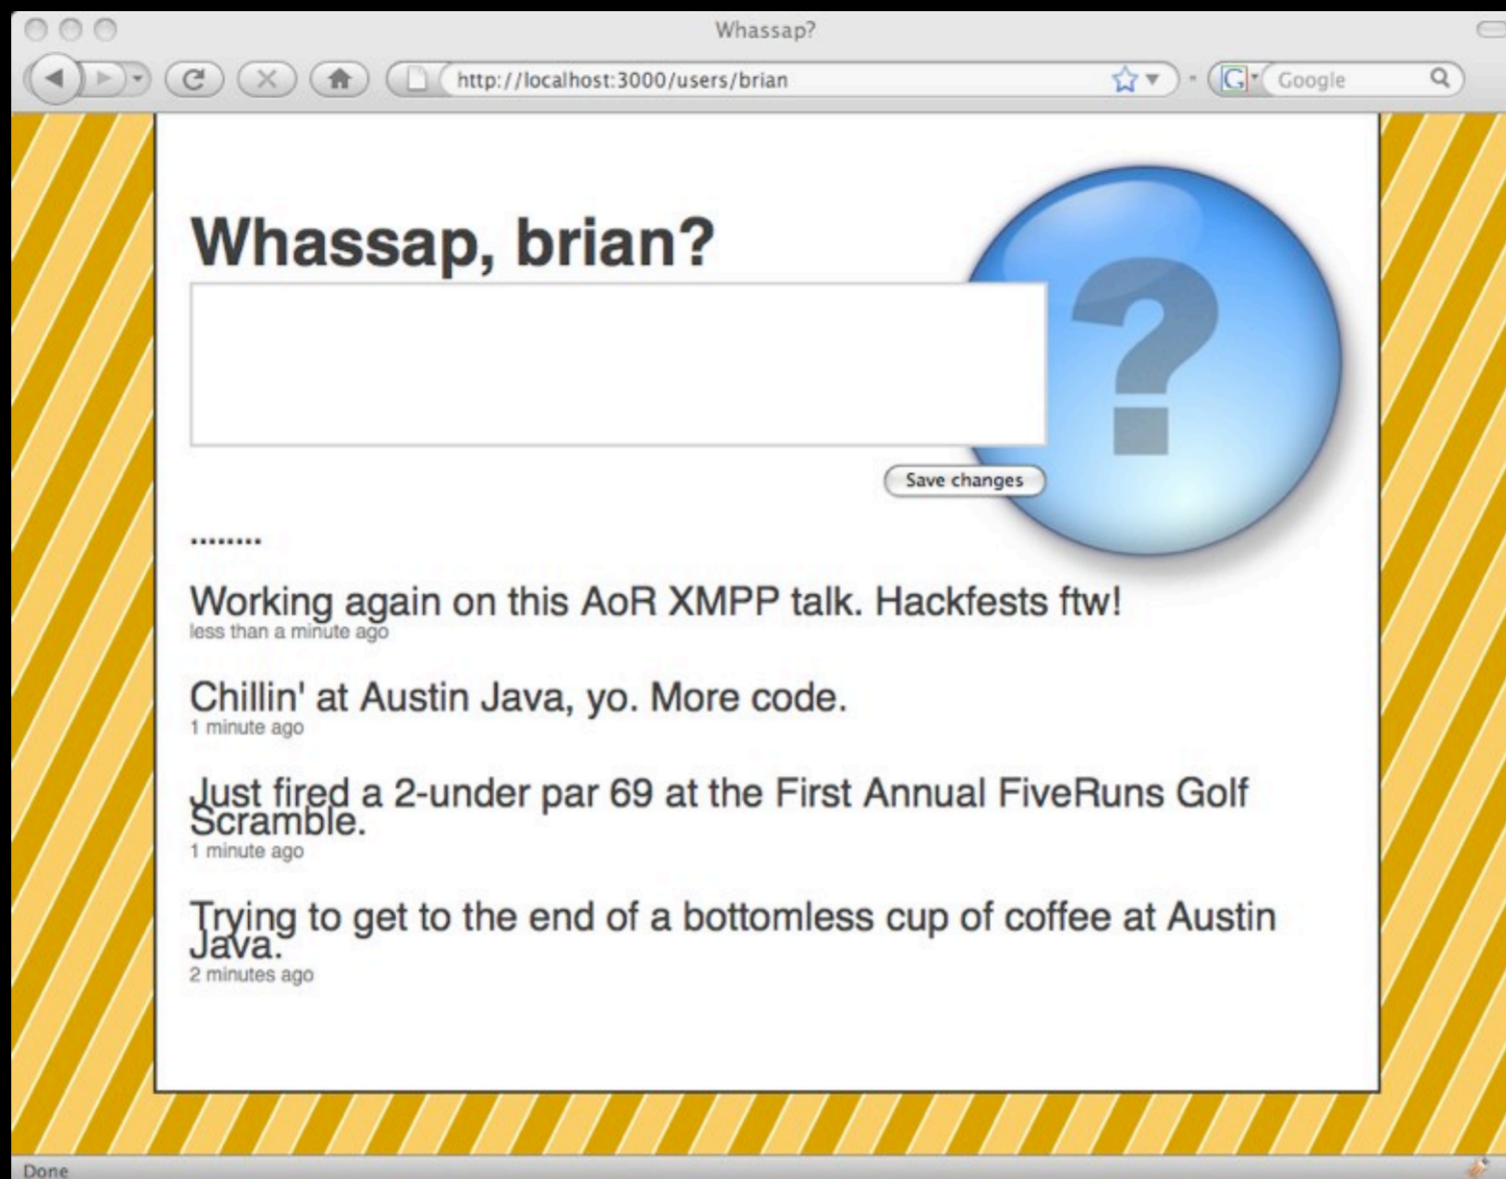
Core IM
VoIP

Embedded Images & Video
Multi-user Chat

The core IM functions comprise 'XMPP-Core', but there are a wide range of extensions today.

Ok, back to the present...

Whassap



Whassap!?

(demo: Whassap Rails app. running: ejabberd, rails, psi)

JIDs

brian@tiamat.local

brian@tiamat.local/iphone

resource

JIDs are a lot like email addresses. XMPP addressing.
Resources allow you to make a device/location-specific distinction.

Gems

xmpp4r-simple
xmpp4r

Useful gems:

xmpp4r-simple is feature-poor but dirt simple to use.

xmpp4r, by contrast, is NOT SIMPLE but is incredibly feature-rich.



Not pretty.

Gems



-sir
p4i



The simple gem was written by Blaine Cook, former chief architect at Twitter. If you're doing basic IM over XMPP stuff, use his gem.

Image sources:

<http://cache.valleywag.com/assets/resources/2008/05/BlaineCook.jpg>

http://photos.igougo.com/pictures-photos-j14184-c4-s2-p61061-Shark_Attack.html

Your Dashboard — GitHub

git https://github.com/ shark attack

Most Visited ▾ Smart Bookmarks ▾ Getting Started Latest Headlines 📡 Apple ▾ Amazon eBay Yahoo! News ▾

github
SOCIAL CODE HOSTING

bdainton account | profile | guides | log out
repositories: all | search

News Feed

for you | from you

- chuyeow** commented on rails/rails about 3 hours ago
Comment in b8e8be8:
I think karmi hit the nail on the head – the documentation ""promises"" some functionality that is insufficient.
- svenfuchs** committed to rails/rails about 3 hours ago
e826c992014b24475203bf1a88b3f7a2356f7f5
add activerecord tests to make sure the deprecated %s and %d interpolation syntax still works
- svenfuchs** committed to rails/rails about 3 hours ago
31dcd7811185134f16a52744ad50f87d197ca438
add activerecord tests for deprecation of %s and %d in error messages (and translations in general)
- svenfuchs** committed to rails/rails about 3 hours ago
79c8b104d676ce40268caaa530b7e4dcc0a8ff20
change activerecord validation tests to not use the deprecated interpolation syntax any more
- josh** committed to rails/rails about 6 hours ago
70a4f6c2dd66d855ff66a176076ad45caf3ed6ab
Switched computed public paths cache over to a simple hash w/ mutex

Your Repositories (create a new one)

all | public | private | sources | forks

- camptweet
- community
- rails
- xmpp4r-simple

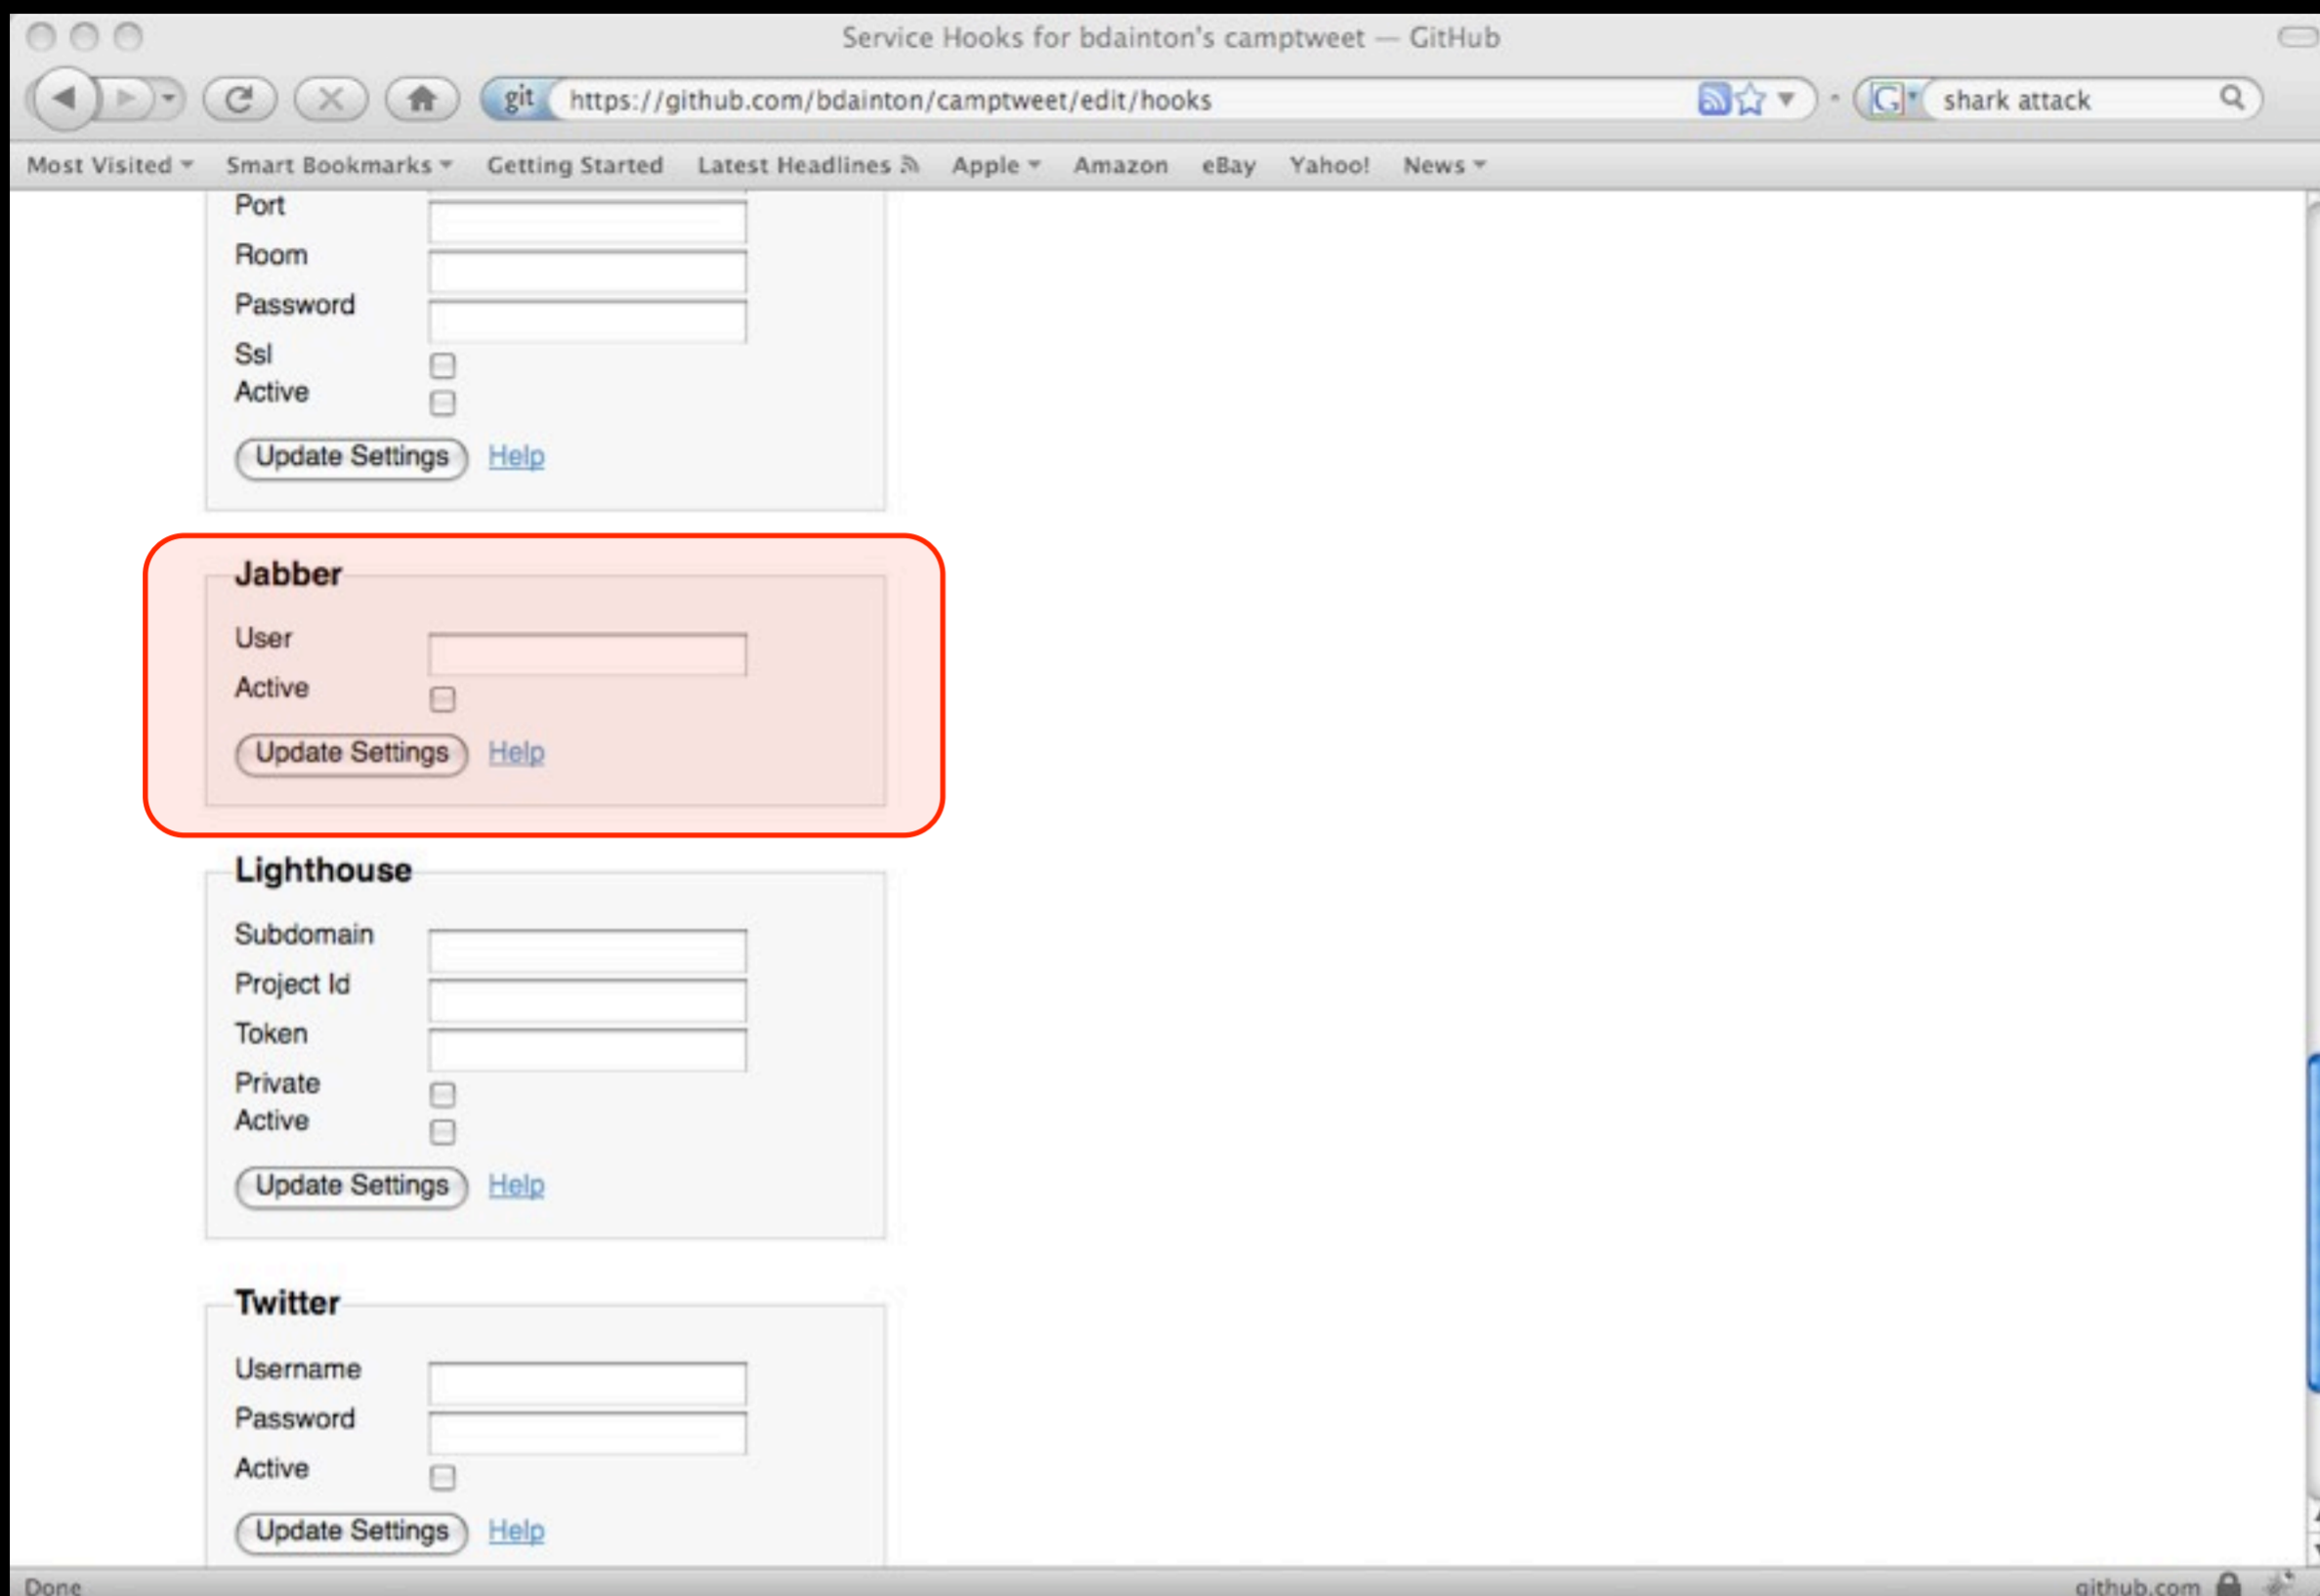
Watched Repositories

- nakajima / acts_as_nakajima
- five runs / altered_beast
- gilesbowkett / archaeopteryx
- five runs / brigit
- bigfleet / community
- gvaughn / community
- tradeking / community
- five runs / data_fabric
- five runs / five runs_tuneup
- five runs / instrument
- terrbear / leetwit

Done github.com

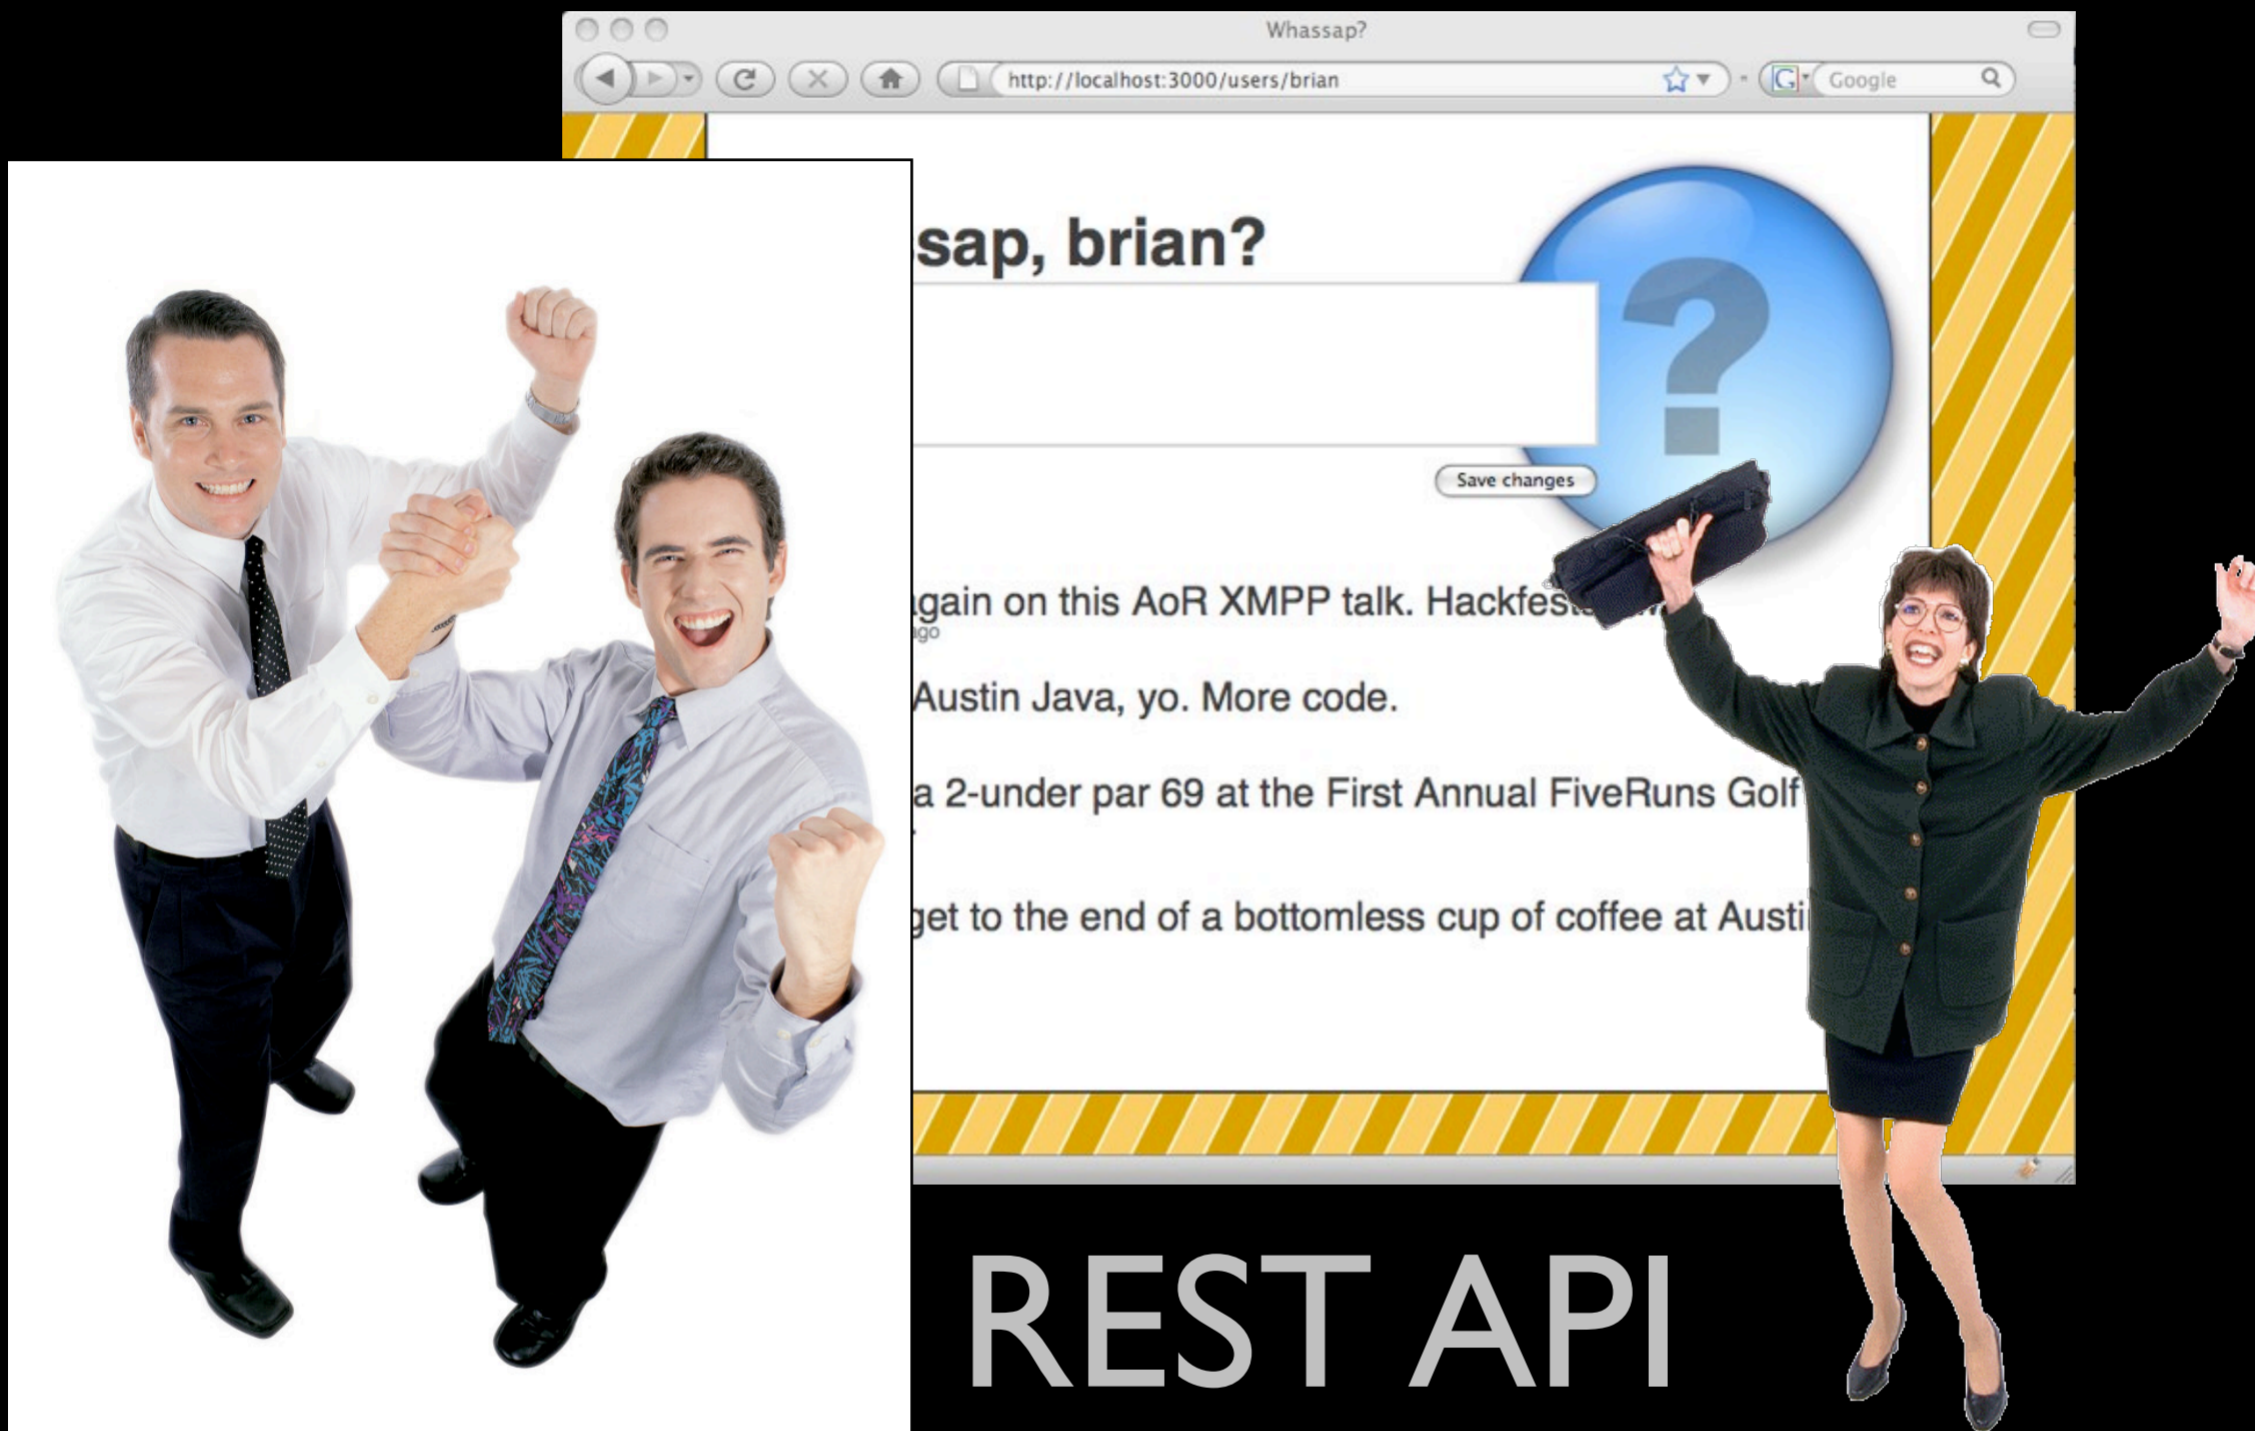
GitHub does Jabber notifications.

Source:
<https://github.com>



Source:
<https://github.com>

Whassap



REST API

So, after a few weeks in production, Whassap starts to get a small -- but excited and loyal -- userbase. The begin to take advantage of the Whassap REST API to poll for updates.

(demo whassap_rest_client w/httparty)

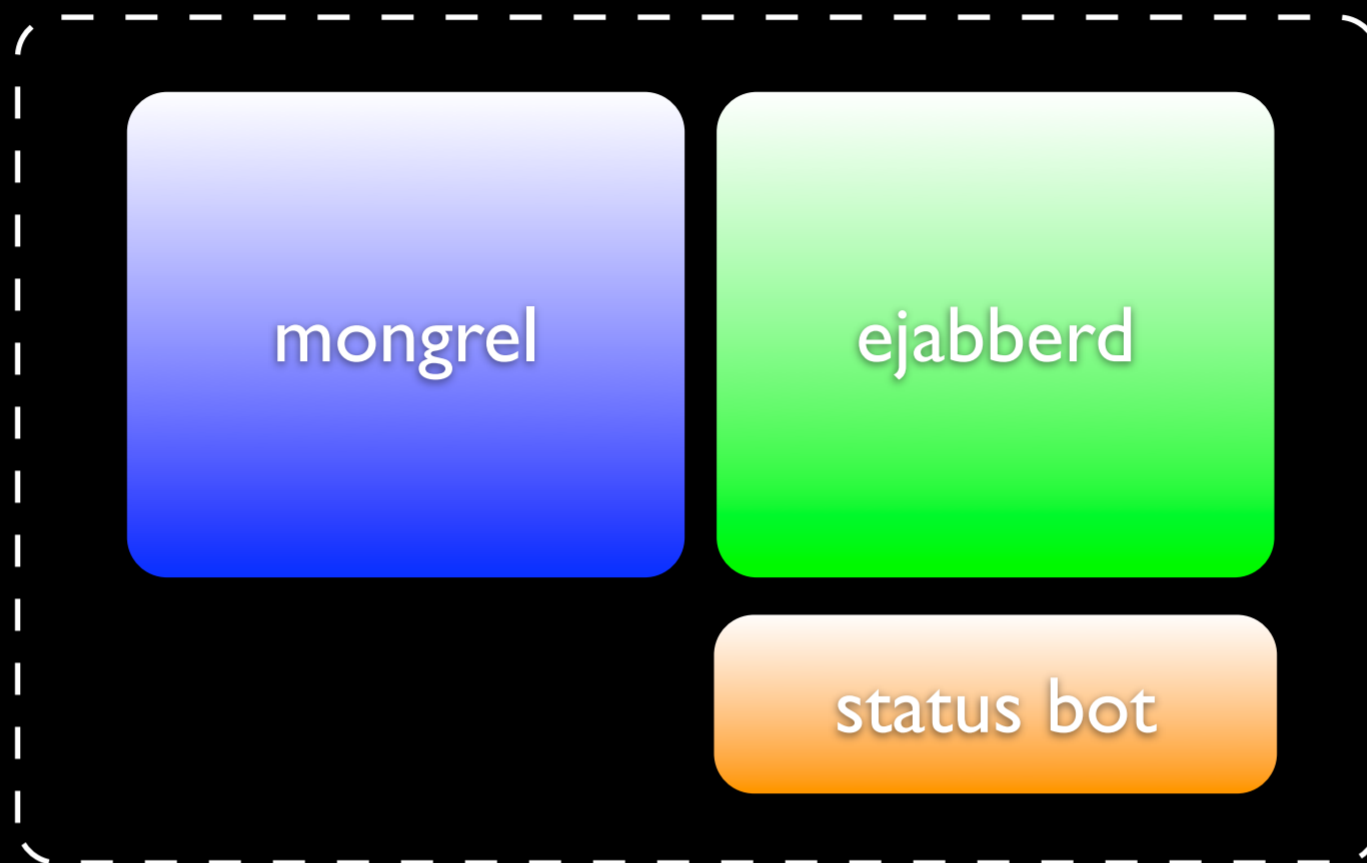
Things are good. REST is awesome; we should all be designing and writing our apps RESTfully.

Image sources:

http://www.chelseapiers.tv/2008/01/hold_the_press_its_wigan_away.html

<http://www.healingwithnutrition.com/vitamin.html>

Whassap



tiamat.local

Clearly I've got traction with Whassap. It's time to offer a new service: the capability to submit Whassap statuses via IM.

Introducing the Whassap Status Bot.

(demo Whassap XMPP Status Bot; show status submission via IM, which uses the REST API to actually submit the status, same resulting IM)

Whassap

http://www.techcrunch.com

Most Visited Smart Bookmarks Getting Started Latest Headlines Apple Amazon eBay Yahoo! News

See your signal strength in street level detail before you sign-up.

Tech Gadgets Mobile Enterprise CrunchBase TechCrunch50 More

TechCrunch

Whassap is the Twitter We Always Wanted

by Erick Schor

Most social network impressions), \$76.50. Text card that mar networks. Th

LinkedIn claim visited the sit of its membe decision mak

Done

http://daringfireball.net/2008/07/not_yet_squirrely

Most Visited Smart Bookmarks Getting Started Latest Headlines Apple Amazon eBay Yahoo! News

DARING FIREBALL

future, and it is Whassap

week regarding the striking improvement in version of WebKit in the iPhone 2.0 update, or this was perhaps attributable to Apple SquirrelFish, the all-new JavaScript interpreter in early June. Announcing the project,

h faster than WebKit's previous numbers. On the SunSpider JavaScript is 1.6 times faster than WebKit's

y unlikely that a project that new — having would be included in the iPhone 2.0 OS. And, the SquirrelFish announcement show, JavaScript interpreter is itself 2.7 times faster JavaScript performance had improved even before SquirrelFish.

When in doubt, look at the source, and the is available from Apple's Darwin sources

Done

javascript.document.body.contentEditable='true'; document.designMode='on'; void 0

keynote z axis

Most Visited Smart Bookmarks Getting Started Latest Headlines Apple Amazon eBay Yahoo! News

Daring Fireball: Not Yet Squirrely keynote element on top - Goog... Scobleizer - Tech geek blogge...

scobleizer

Sponsored by Norton from symantec

Advertisement:

Sept. 7, 2007 Norton Internet Security 2008 No purchase necessary. Click for full details.

Scobleizer Archives

Google Reader Shared Items

Twitter Feed

Tech Events Calendar

Travel Schedule

« CERN: Congrats for... Use Twitter or... »

Whassap: a Paradigm Shift in Personal Eventing

What's the most successful new company out there of the past two years?

I'd argue it's Whassap, a music sharing and discovery service that came out on Facebook first, but has moved other places.

When I quit Microsoft they had no users.

Today they have 30 million.

How did THAT happen? It wasn't because they got bloggers all hot and bothered about them. They didn't win any contests. I don't even think they showed up at Demo or TechCrunch or other industry conferences. If they did, they didn't cause any headlines on blogs about how they were going to be THE next hot thing.

What did they do? They were first out the gate on Facebook's application platform. That was a real paradigm shift. It was the first application platform I remember where I could see the apps YOU had loaded.

Second datapoint?

Yesterday I ran into the founders of Posimotion. Their company didn't

Done


BOOM. An exponential increase in users. In traffic. In hits to the site, and in usage of the REST API.

Whassap

The Facebook logo, consisting of the word "facebook" in white lowercase letters on a blue rectangular background.The Twitter logo, featuring the word "twitter" in a light blue, lowercase, rounded font on a white background with a light blue border.The FriendFeed logo, with the word "friendfeed" in a blue, lowercase, rounded font on a white background with a blue border.The Pownce logo, featuring the word "Pownce" in a bold, white, lowercase font on a dark brown rectangular background.

The Whassap Facebook app hits the scene. Third-party services propagate Whassap statuses to Twitter and Pownce. FriendFeed (social content aggregator) starts polling for Whassap updates.

Whassap



tiamat.local

(server on fire)

Another recent talk on this XMPP topic (given at OSCON this year)...

Beyond REST? Building Data Services with XMPP PubSub

http://www.slideshare.net/rabble/beyond-rest-building-data-services-with-xmpp

Most Visited Smart Bookmarks Getting Started Latest Headlines Apple Amazon eBay Yahoo! News

slideshow BETA

Home My Slidespace Upload Community Tags Widgets

Latest | Most Viewed | Most Embedded | Featured | Most Favorited | Most Downloaded | Slidecasts

It's free! Upload

XMPP/Jabber vs SIP/SIMPLE
In-depth comparison of XMPP/Jabber and SIP/SIMPLE.
www.Jabber.com

Data Loss Prevention
Proven Protection for Data in Motion, at Rest and the Endpoint
www.Orchestria.com

Beauty Rest
Save 50% On All Simmons. Free In-Home Delivery. Order Today!
www.US-Mattress.com/Simmons

SwiftMQ High Availability
Test the most advanced JMS HA messaging system today!
www.swiftmq.com

Ads by Google

Beyond REST? Building Data Services with XMPP PubSub

From [rabble](#), 1 month ago

Decentralization, increased participation, real-time expectations, [more](#)

10812 views | 6 comments | 42 favorites | 288 downloads | 12 embeds ([Stats](#))

Share Favorite + Group / Event Download file

Embed `<div style="width:425px;text-align: options`

Categories
[Add Category](#)

Tags
[dataservices](#) [webservice](#) [oscon08](#) [oscon](#) [rest](#) [pubsub](#) [xmpp](#) [jabber](#) [oauth](#) [scalability](#) [more](#)

Groups / Events

More by user

- [liberating location](#) 2817 views
- [Testing Legacy Rails Applications](#) 3338 views
- [Introduction to Active Record at MySQL Conference 2007](#) 5304 views

Related Slideshows

- [Liberating Location - Fire Eagle - Ecomm 2008](#) 2817 views
- [Testing Legacy Rails Apps](#) 3338 views
- [Introduction to Active Record at MySQL Conference 2007](#) 5304 views

REST is Newtonian physics.

Its like REST is Newtonian physics. For every day problems, its good enough. It makes sense, its coherent, and its well understood. But it breaks down at scale. Its breaks down when you're talking about really small things, and really fast things, and really really huge things, it doesn't explain quarks and quasers.

share 5 / 72 full

Post: [h5](#) [smiley](#) [x](#) [f](#) [b](#) [fw](#) [Post To Twitter](#)


BOOKMARK

Transferring data from cdn.slideshare.net...

...likened REST to Newtonian physics; the analogy is accurate.
REST APIs work great for a long while, but break down in extremely high-scale environments.

Sources:
<http://www.slideshare.net>
Referenced slide copyright Evan Henshaw-Plath (ENTP.com), Kellan Elliot-McCrea (Flickr.com)

Social Web



Apps in the middle of the social web are getting slammed.

- there are 10-20 new Twitter statuses added every second
- there are 60 new photos uploaded to Flickr every second

And those are *just* the POSTs. API polling traffic is order of magnitudes higher.

(demo the GETs)

Image source:

<http://code.google.com/apis/socialgraph/>

The Situation

many clients

interested in specific topic

topic update interval: unknown

data freshness: important

What we've got is a situation where:

- many, many clients are interested to know when
- information on a specific topic (my Whassap status, for example) has been updated
- we can't make any assumptions about **when** the topic will be updated
- and timeliness of the information is important

XMPP PubSub

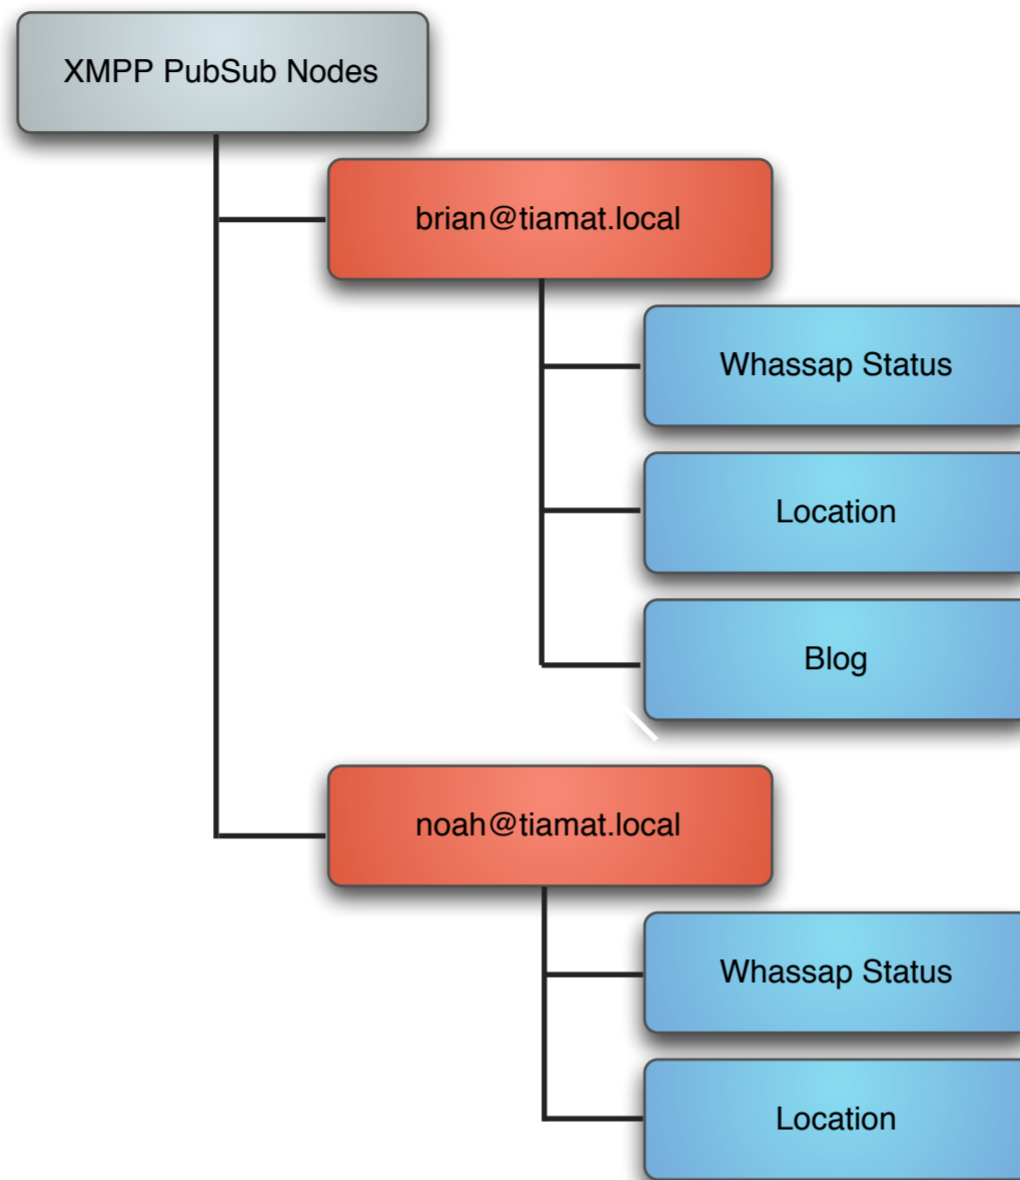
Publish/Subscribe

Enter XMPP PubSub.

The concept of the publish/subscribe model is nothing new.

With XMPP, though, things operate a bit differently because of the persistent connection.

XMPP PubSub



The XMPP PubSub spec: collection of nodes for each JID.
Permissions-based controls for who can publish/subscribe.
Published data is immediately available to a subscriber, manifests itself as a message from the pubsub service to the subscriber.

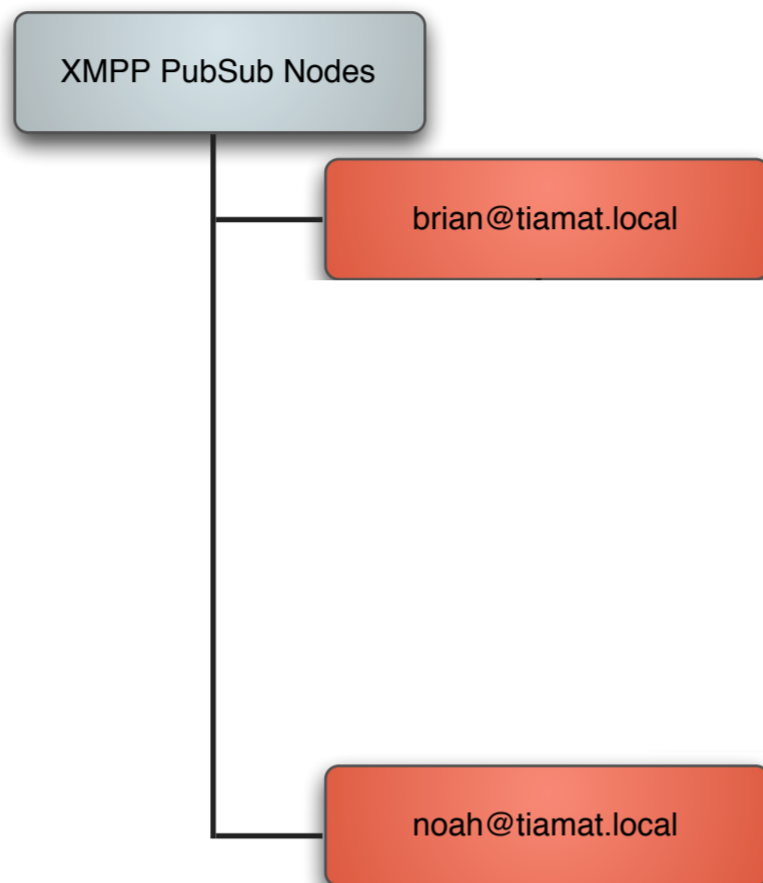
To handle merely single personal updates, a simpler, pared-down spec was created...

PEP

Personal Eventing via PubSub

Personal Eventing via PubSub.

PEP



It constrains the set of nodes to JUST 1 node per Jabber ID.

Users may ONLY publish to their own nodes.

But...users may subscribe to anyone else's node. This is basically Twitter via XMPP.

The screenshot shows a web browser window displaying a GitHub repository page. The browser's address bar shows the URL `http://github.com/bdainton/xmpp4r-simple/tree/master`. The page header includes the GitHub logo and the text "SOCIAL CODE HOSTING". The user profile for "bdainton" is visible in the top right corner, showing a profile picture, name, and links for "account", "profile", "guides", and "log out". Below the profile, it indicates "repositories: all" and a "search" button. The repository navigation tabs include "Source", "Commits", "Graphs", "Wiki (1)", "Watchers (1)", "Network (5)", and "Admin". The "Source" tab is selected, and the current branch is "master". The repository name "bdainton / xmpp4r-simple" is displayed with buttons for "edit", "pull request", "unwatch", and "download". Below this, the description "Fork of blaine/xmpp4r-simple" is shown, along with fields for "Description", "Homepage", "Public Clone URL", and "Your Clone URL". A search box for the repository is also present. The commit history section shows a commit by "blaine (author)" on "July 30, 2008" with the message "add xmpp4r-simple.gemspec". The commit details include the commit hash "37d466f703a71dd624319bb272731661aa24497c", the tree hash "8d7c64b4fcb03625446811902eb7a149e76f2693", and the parent hash "dfde2c5bdfabba0878cd6c6fb41092257c5132689". Below the commit history, a table lists the repository's files and their commit messages:

name	age	message	history
CHANGELOG	July 30, 2008	add xmpp4r-simple.gemspec [blaine]	
COPYING	November 08, 2006	oops, i was already in trunk. move things aroun... [romeda]	
README	July 30, 2008	update copyright [blaine]	
Rakefile	July 30, 2008	add xmpp4r-simple.gemspec [blaine]	

At the bottom of the page, there is a search bar with the text "Find: diagram styles" and buttons for "Next", "Previous", "Highlight all", and "Match case". The browser's status bar shows "Done".

I've forked xmpp4r-simple it on github to add some PEP support.

(demo pep publisher. intermittently subscribe some other users)

Benefits: no incessant polling, and we get the updates that we want as soon as they exist. Cool.

When you extrapolate this to the non-Newtonian scale, the savings in bandwidth and CPU usage are incredible.

Federation



Old IM networks would not play nicely together.
XMPP does away with those walled gardens.

Federation

juliet@capulet.org

romeo@montegue.net

Concept of open federation is baked directly into the core of XMPP servers.

The analog here is with email.

One server knows how to communicate with the other. The net result is a distributed, but federated, userbase.

Federation

bdainton@twitter

damon@twitter

wbruce@whassap

Federation + PEP: federated social networks bring cross-server addressability.



Random Etc. - Firehose

http://www.tom-carden.co.uk/2008/05/16/firehose/

Most Visited Smart Bookmarks Getting Started Latest Headlines Apple Amazon eBay Yahoo! News

Firehose

In Weblog on 16 May 2008 tagged Flash, jabber, openfire, stamen, twitter, Work, xiff, xmpp with 6 comments

Update (June 20th): the 'Firehose' app described in this post relied on Twitter's public timeline IM support, which has been down for several days in a row. Firehose will return when Twitter gets their stuff back online.

going to sleep now to make up for Saturday.

Firehose achieved. Unknown Error

マイメロがタイムラインにあるとなごみますね

@jenetics Got a new camera. Been learning more about Flickr, really loving it! No more botched picture-taking by the time I get to NZ:)

寒い雨だった

@tiabel OK, I'm going....thank you for the nudge :-)

誰か...ワタシにやる気を下さい...

Thanks for trying this out. This is Twitter's public timeline in real-time. It's just a tech demo. Silly, really. But there's always more to do!

ature as I Stamen!

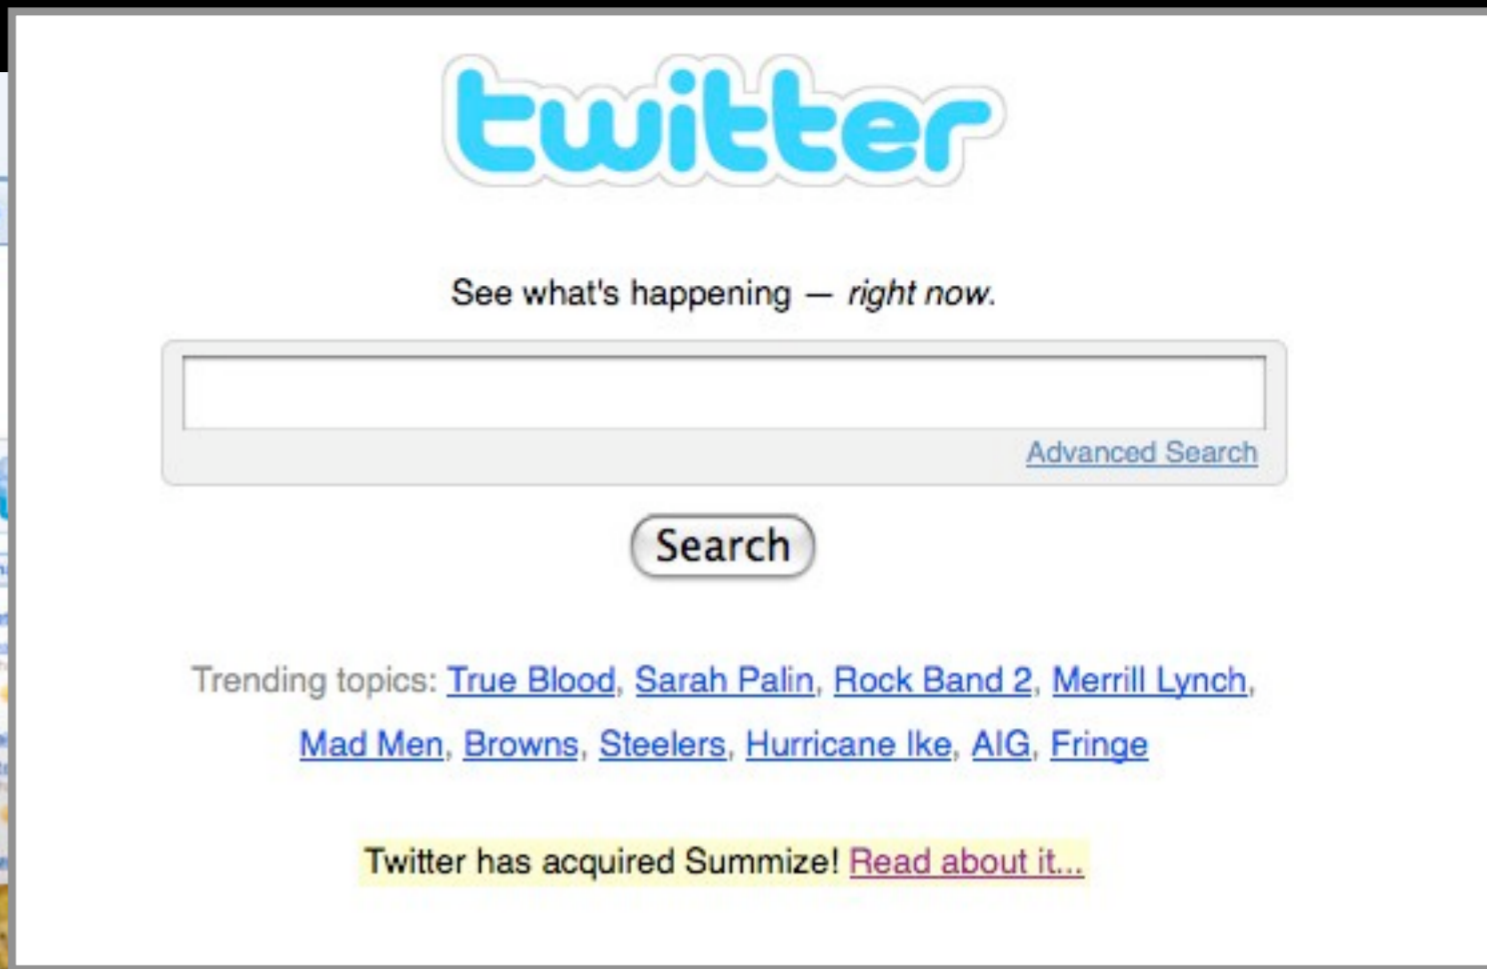
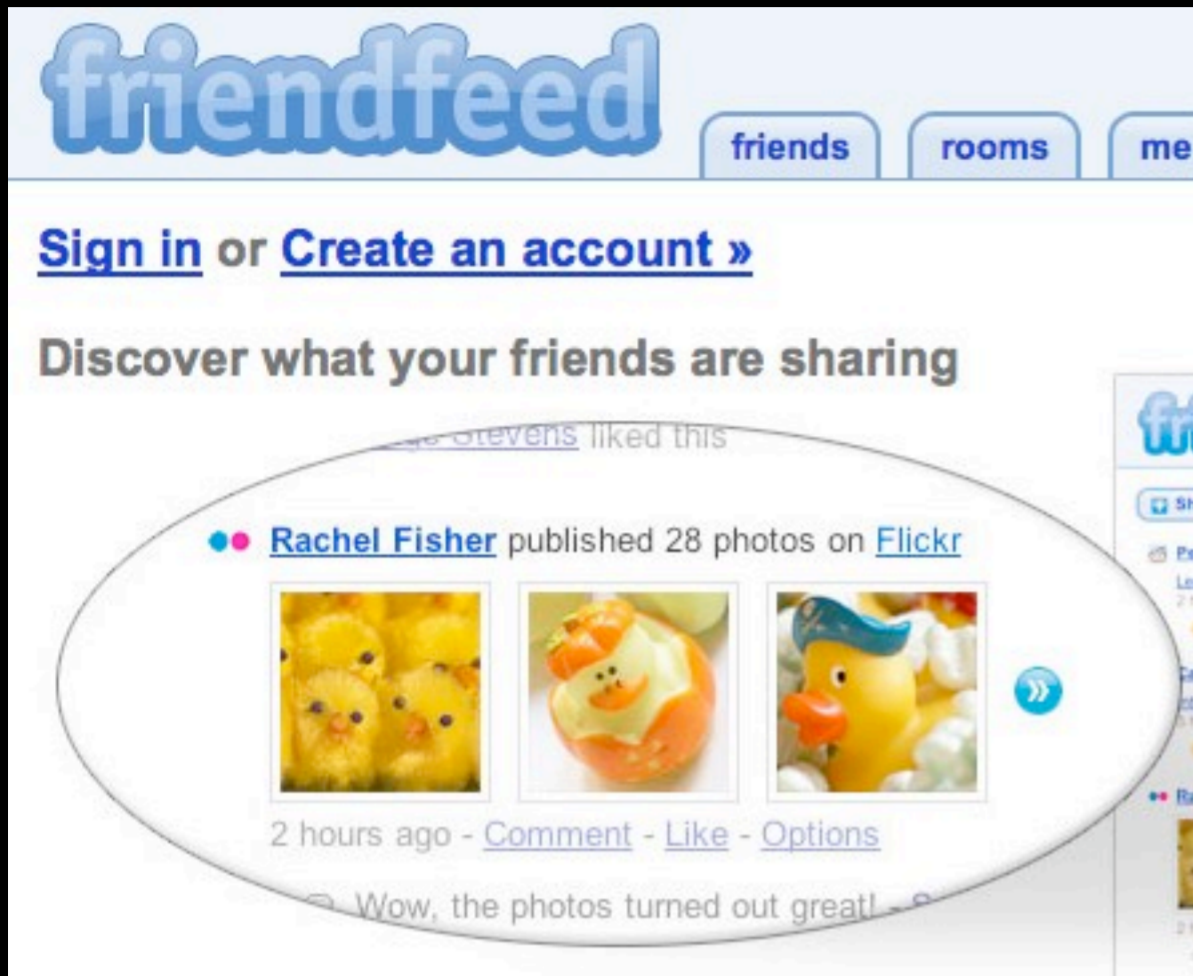
About a week ago I made a proof of concept Flash site that I'm calling Firehose. It shows every post from Twitter's public timeline as fast as it can, in a BigSpy style.

Done

For a while, Twitter made its user updates available via XMPP PubSub. Stamen firehose (now defunct)

Source:
<http://www.tom-carden.co.uk/2008/05/16/firehose/>

twitter



Gnip 
(guh-nip)

[API](#) | [create account](#) | [login](#)

Making data portability suck less.

A few select services gained access to the Twitter firehose. FriendFeed and Summize (later became Twitter Search).

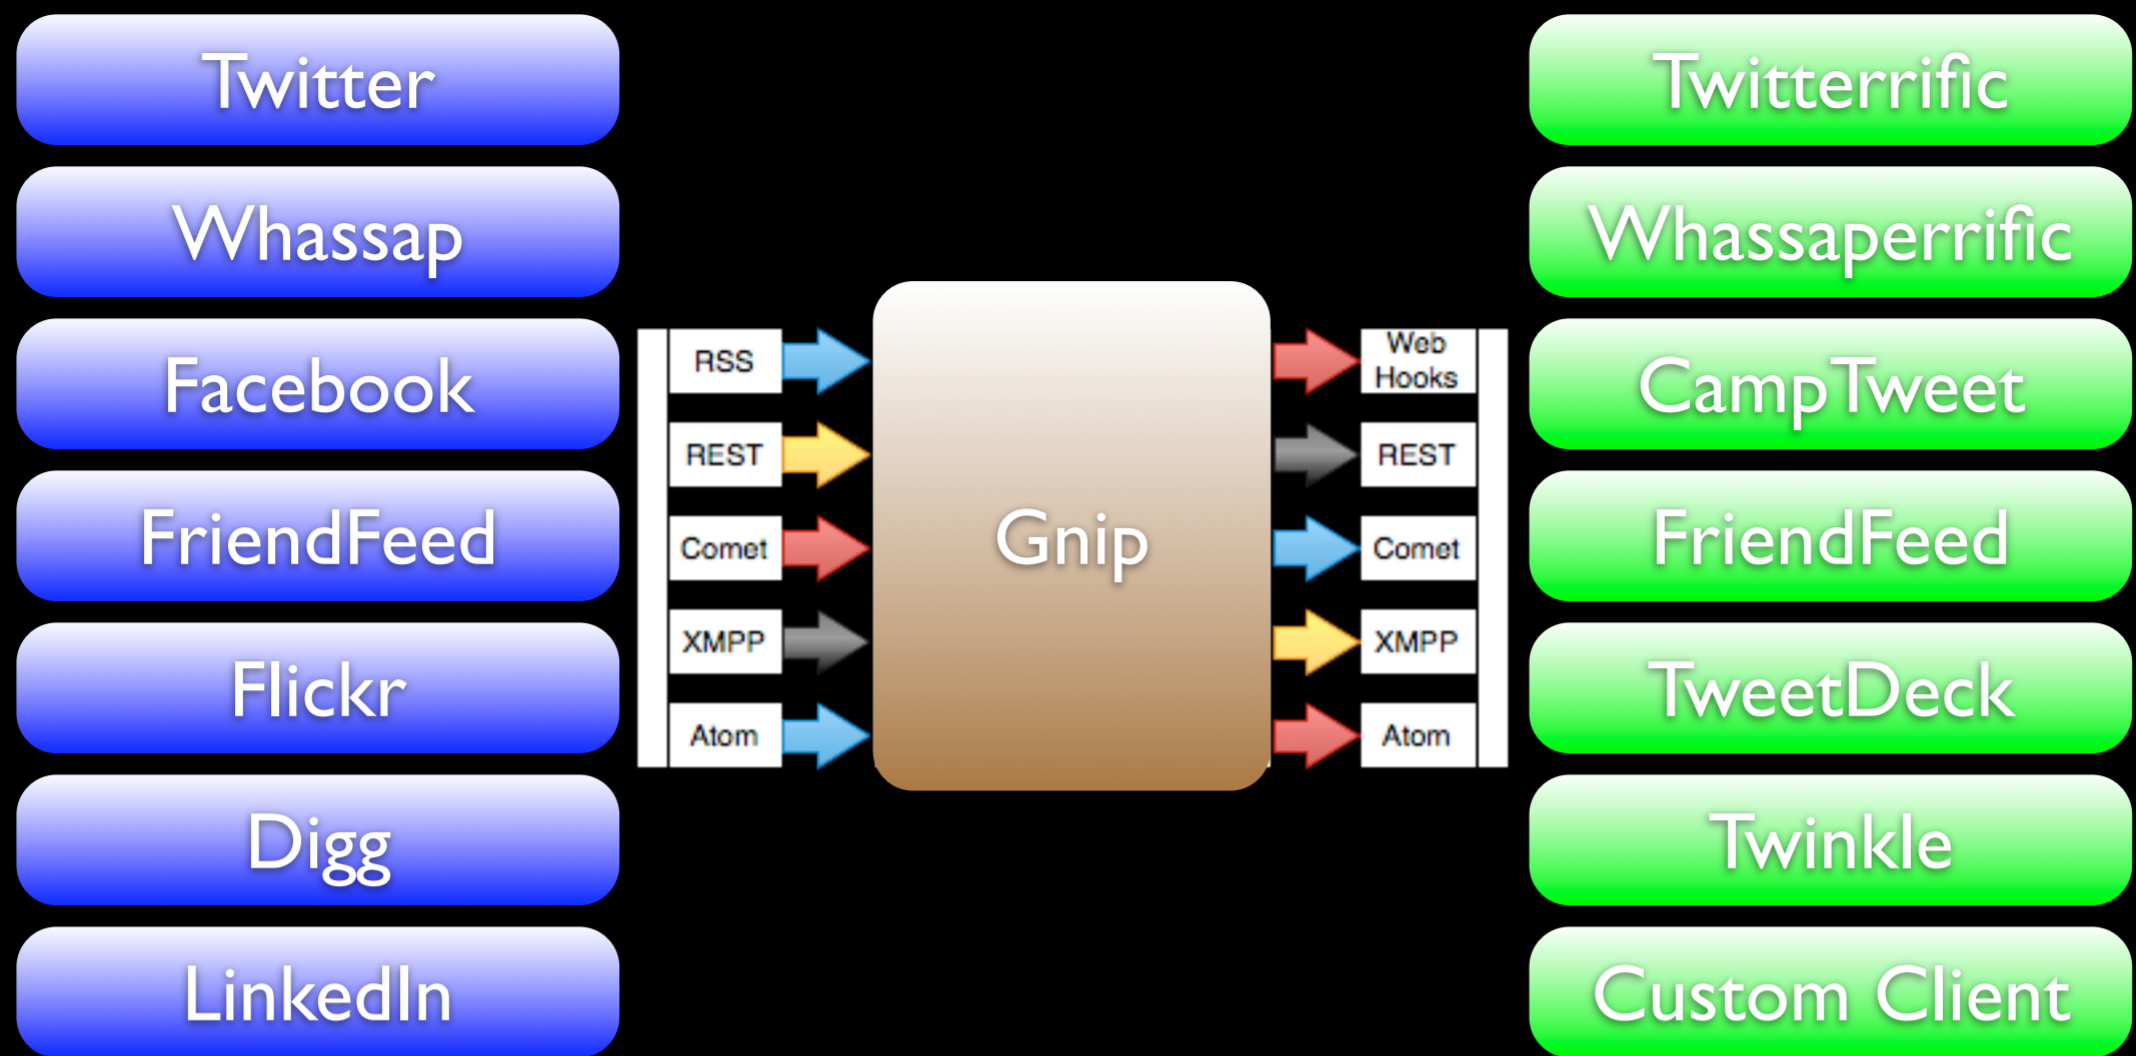
A new service has access to the Twitter firehose as well: Gnip.

Image sources:

<http://www.friendfeed.com>

<http://search.twitter.com>

<http://www.gnipcentral.com/>



[API](#) | [create account](#) | [login](#)

Making data portability suck less.

Gnip sits in the middle between Data Publishers (like Twitter and...Whassap) and Data Consumers like you and me (and other social aggregator services like FriendFeed, for example).

They accept data however it is produced and allow users to consume data however they want. Big benefits for both producers (reduced load) and consumers (reduced/removed throttling, flexible access methods).

Image source:

<http://www.gnipcentral.com/>

Gnip



gnip

[follow](#)

[message](#)

Name: Gnip Admin

Email: developers@gnipcentral.com

Blog: <http://blog.gnipcentral.com>

Company: Gnip

Location: Boulder, CO

Followers: 4

Public Repos: 6

gnip-dotnet

.Net library for utilizing Gnip services.

Created July 07, 2008

[all commits](#) [commits by owner](#)

52 week participation

gnip-java

Java library for utilizing Gnip services.

Created July 11, 2008

[all commits](#) [commits by owner](#)

52 week participation

gnip-perl

Perl library for utilizing Gnip services.

Created July 11, 2008

[all commits](#) [commits by owner](#)

52 week participation

gnip-php

PHP library for utilizing Gnip services.

Created July 11, 2008

[all commits](#) [commits by owner](#)

52 week participation

gnip-python

Python library for utilizing Gnip services.

Created July 09, 2008

[all commits](#) [commits by owner](#)

52 week participation

gnip-ruby

Ruby library for utilizing Gnip services.

Created July 03, 2008

[all commits](#) [commits by owner](#)


52 week participation

Gnip libs on GitHub.

Remember

XML Messages

Persistent Connection



HTTP
+
XMPP



Thanks

The bigger picture:

- the web is evolving, so developer demands are changing
- HTTP is no longer the end-all and be-all for creating webapps today

Add XMPP to your bag of tricks and you'll be far more ready to tackle some of today's hardest problems.

Image sources:

<http://code.google.com/apis/socialgraph/>

<http://inbabble.com/2008/01/27/interview-what-is-the-greatest-challenge-in-mobile-today/>

Brian Dainton

brian.dainton@gmail.com

brian@fiveruns.com

twitter.com/bdainton