

Search: the right tool ...
... but what is the job?



Jетро Coenradie
jco@trifork.nl
[@gridshore](#)

TRIFORK.

The right tool ...



... but what is the job?



The right tool

Search

and the job is

Two people to help me

Meet the customer

Marc



Meet our
software
developer

John



What is a Search Solution?



Search = finding stuff

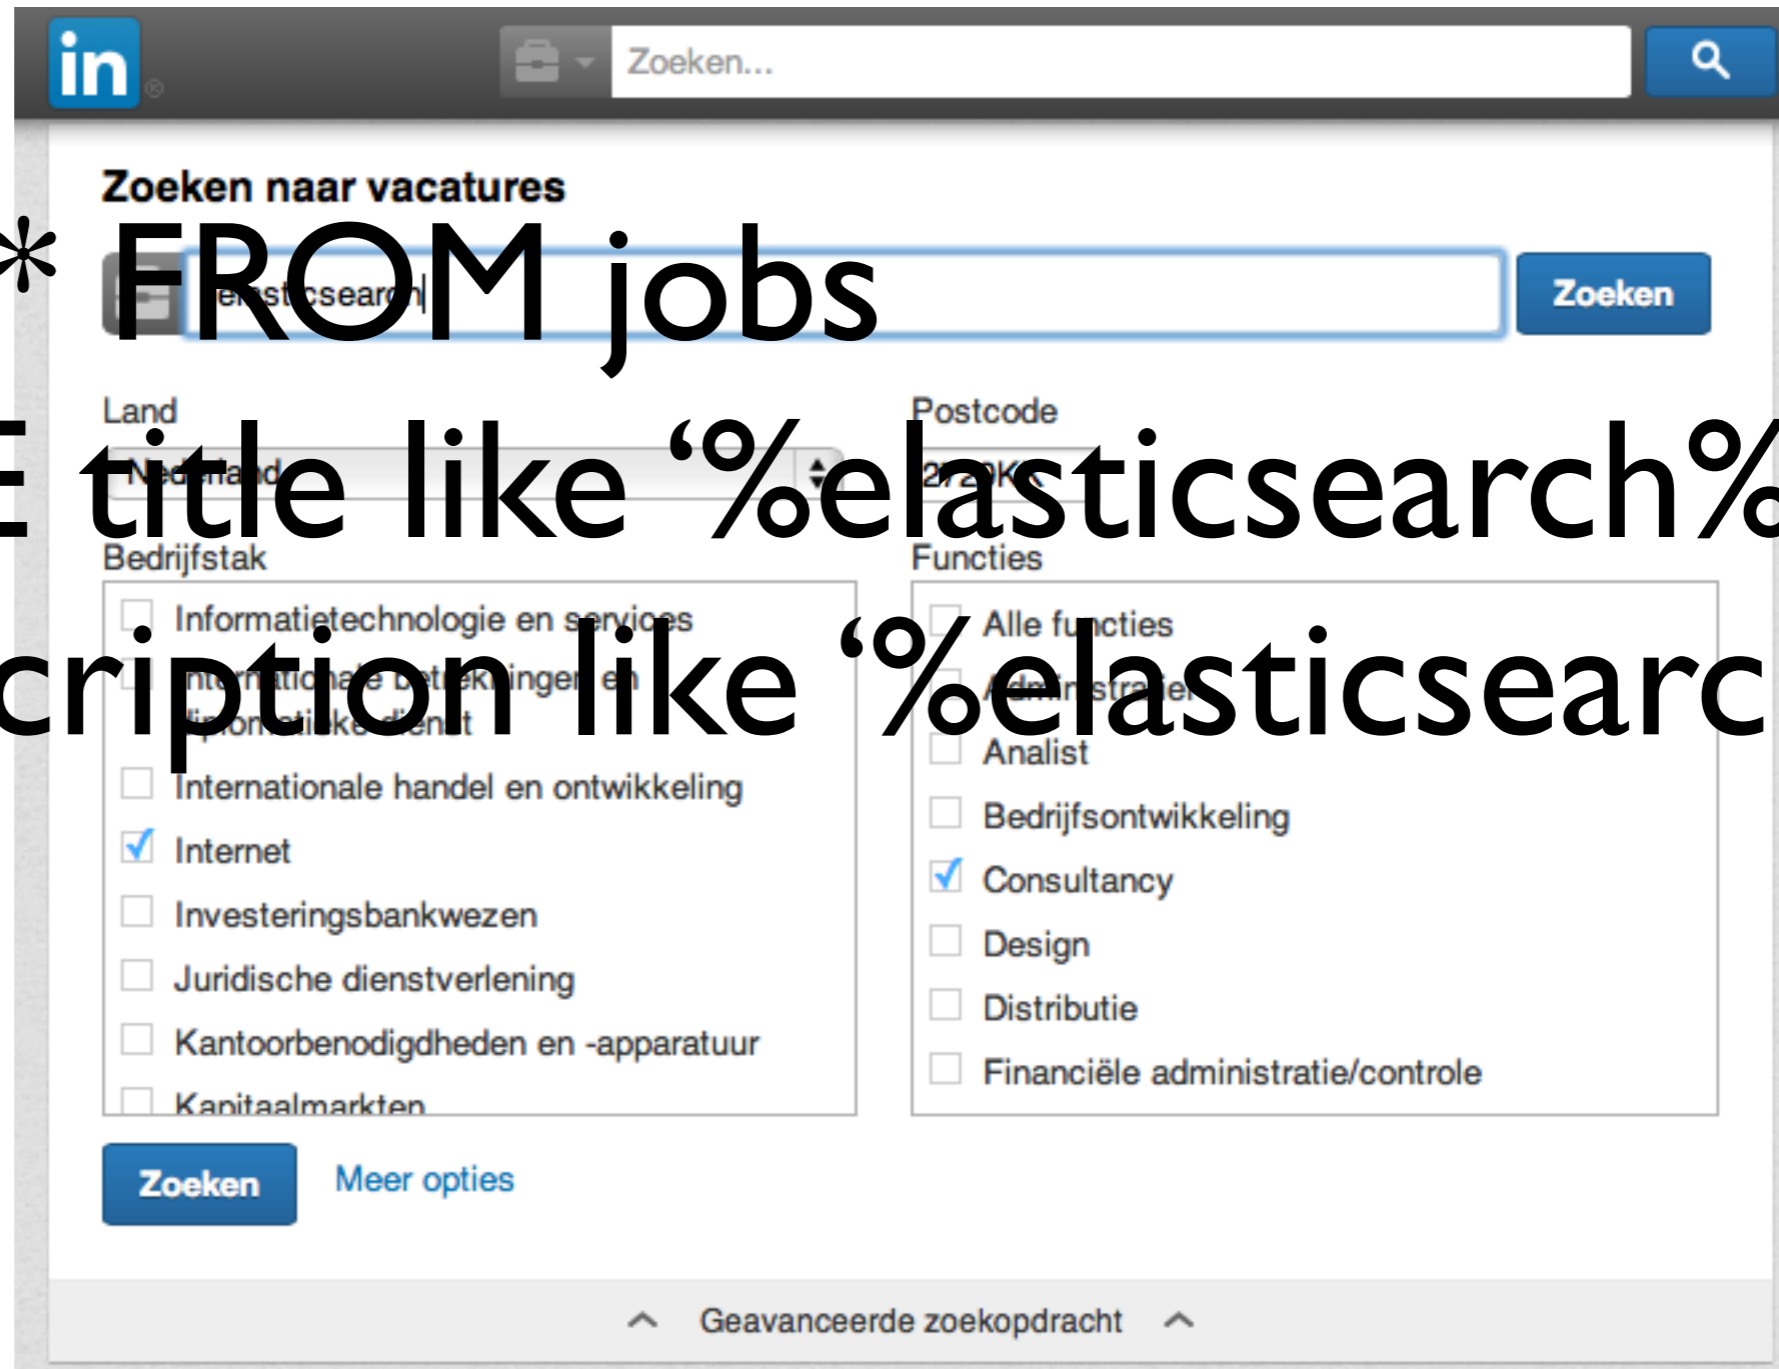


Search Solution =

Services for finding stuff



SELECT * FROM jobs
WHERE title like '%elasticsearch%'
OR description like '%elasticsearch%'
OR ...



The image shows a screenshot of the LinkedIn job search interface. At the top, there is a search bar with the text "Zoeken..." and a magnifying glass icon. Below this, the heading "Zoeken naar vacatures" is visible. A search input field contains the text "elasticsearch" and is highlighted with a blue border. To the right of this field is a blue button labeled "Zoeken". Below the search bar, there are several filter sections: "Land" (with "Nederland" selected), "Postcode" (with "2720K" selected), "Bedrijfstak" (a list of industry categories with "Internet" checked), and "Functies" (a list of job functions with "Consultancy" checked). At the bottom of the filters, there is a blue "Zoeken" button and a link "Meer opties". At the very bottom of the page, there is a link "Geavanceerde zoekopdracht" with up and down arrows.



**Not optimized for
free text search**

Inverted index

For this job you need thorough knowledge of elasticsearch.

We are looking for Elasticsearch trainees

Terms

job	1
knowledge	1
elasticsearch	1,2
trainees	2



Analysis

- Tokenizing
- Filters
- Analyzers



Searching/Querying

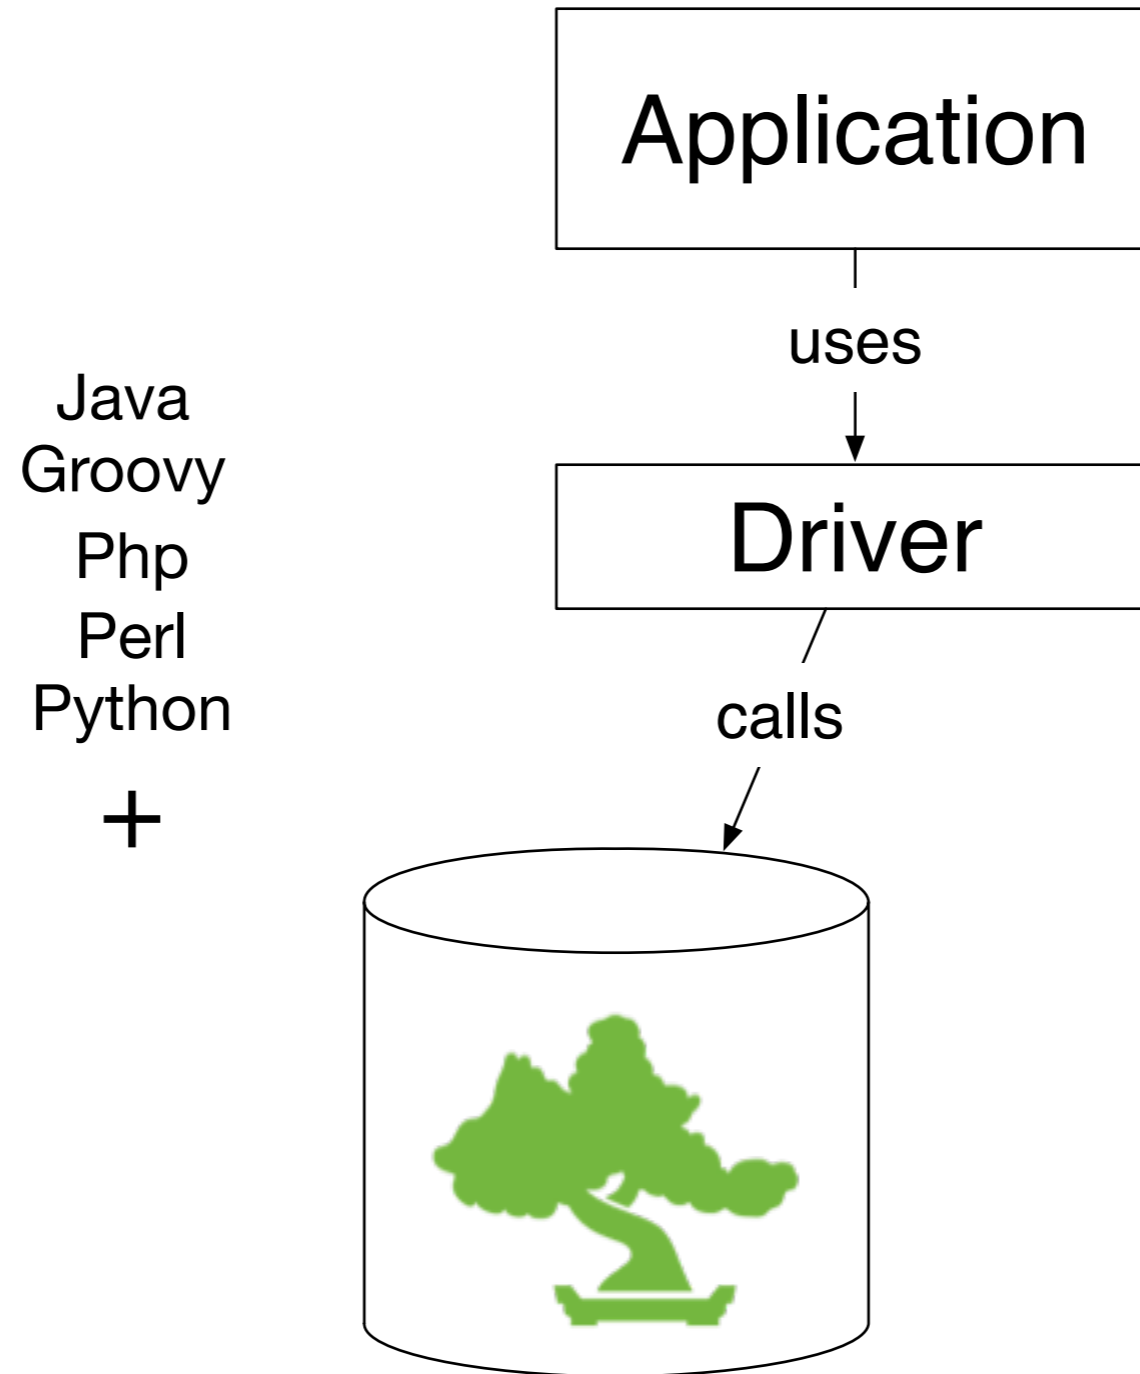
- Stemming
- Synonyms
- Sounds like



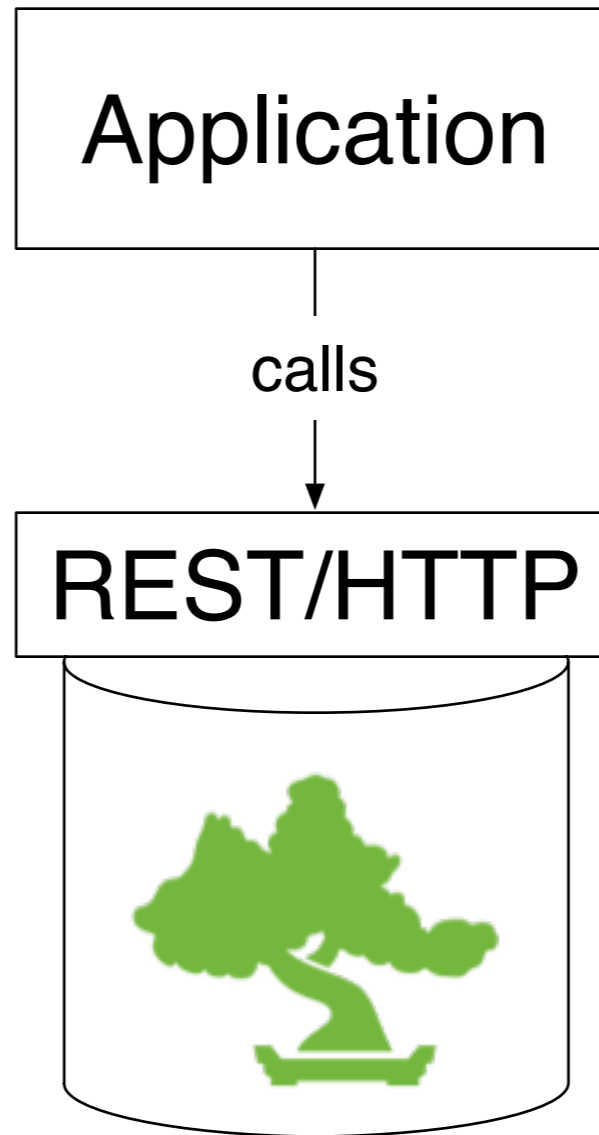
Open source Search solutions



Design overview



Design overview



What can you tell about Security?



There is no security

- Do not expose your search solution to the public
- Use a proxy (guides are available)



Becoming successful

A lot of vacancies

A lot of job seekers




Being successful

Scalability


Availability




Storing more content







Storing more content



Prevent data loss



Keep copies



The right tool ...



GitHub

The screenshot shows a web browser window displaying the GitHub search results for the query 'elasticsearch'. The browser's address bar shows the URL 'github.com/search?q=elasticsearch&ref=cmdform'. The page header includes the GitHub logo, navigation links for 'Explore', 'Gist', 'Blog', and 'Help', and the user profile 'jetro'. The search bar contains the text 'elasticsearch' and a 'Search' button. Below the search bar, the page indicates 'We've found 1,367 repository results' and provides a 'Sort: Best match' dropdown menu. On the left side, there are two filters: 'Repositories' (1,367) and 'Languages'. The 'Languages' filter shows a list of programming languages with their respective counts: Java (346), Ruby (194), JavaScript (157), Python (123), PHP (81), Shell (54), Perl (43), Puppet (42), Scala (18), and C# (15). The main content area displays a list of repository results, each with a repository name, a brief description, the programming language, the number of stars, the number of forks, and a commit history bar. The top results are: 1. 'elasticsearch/elasticsearch' (Java, 5,010 stars, 1,215 forks, last updated 6 hours ago). 2. 'mobz/elasticsearch-head' (JavaScript, 651 stars, 159 forks, last updated 20 hours ago). 3. 'nervetattoo/elasticsearch' (PHP, 180 stars, 46 forks, last updated a month ago). 4. 'elasticsearch/cookbook-elasticsearch' (Ruby, 294 stars, 189 forks, last updated a day ago). 5. 'elasticsearch/elasticsearch.github.com' (CSS, 238 stars, 232 forks, last updated 25 days ago). At the bottom left, there are links for 'Advanced Search' and 'Cheat Sheet'.

Search · elasticsearch

github.com/search?q=elasticsearch&ref=cmdform — Add page to Reading List

Explore Gist Blog Help

jetro

Search

elasticsearch Search

We've found 1,367 repository results Sort: Best match

Repositories 1,367

- Code 74,203
- Issues 6,553
- Users 6

Languages

- Java 346
- Ruby 194
- JavaScript 157
- Python 123
- PHP 81
- Shell 54
- Perl 43
- Puppet 42
- Scala 18
- C# 15

elasticsearch/elasticsearch Java ★ 5,010 🍴 1,215
Open Source, Distributed, RESTful Search Engine
Last updated 6 hours ago

mobz/elasticsearch-head JavaScript ★ 651 🍴 159
A web front end for an elastic search cluster
Last updated 20 hours ago





nervetattoo/elasticsearch PHP ★ 180 🍴 46
Simple PHP client for ElasticSearch
Last updated a month ago

elasticsearch/cookbook-elasticsearch Ruby ★ 294 🍴 189
Chef cookbook for Elasticsearch
Last updated a day ago

elasticsearch/elasticsearch.github.com CSS ★ 238 🍴 232
Last updated 25 days ago

Advanced Search Cheat Sheet

Facets

 Repositories	1,367
 Code	74,203
 Issues	6,553
 Users	6

Languages

Java	346
Ruby	194
JavaScript	157
Python	123
PHP	81
Shell	54
Perl	43
Puppet	42
Scala	18
C#	15


Highlighting

 **elasticsearch/elasticsearch** Java ★ 5,010 🍷 1,215

Open Source, Distributed, RESTful Search Engine


Last updated 6 hours ago




 **karmi/retire** Ruby ★ 1,523 🍷 467

A rich Ruby API and DSL for the **Elasticsearch** search engine

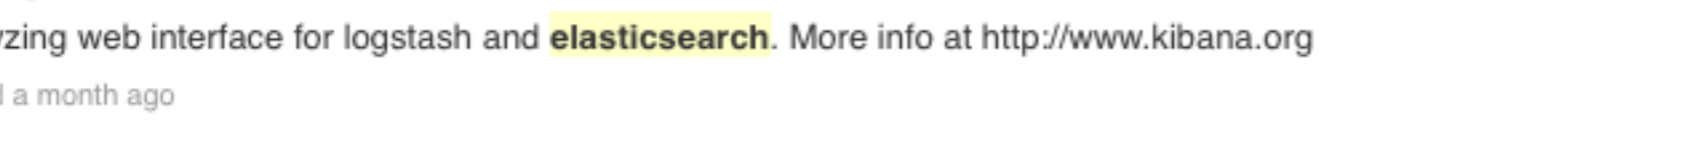
Last updated 14 days ago



 **rashidkpc/Kibana** JavaScript ★ 1,169 🍷 282

A log analyzing web interface for logstash and **elasticsearch**. More info at <http://www.kibana.org>

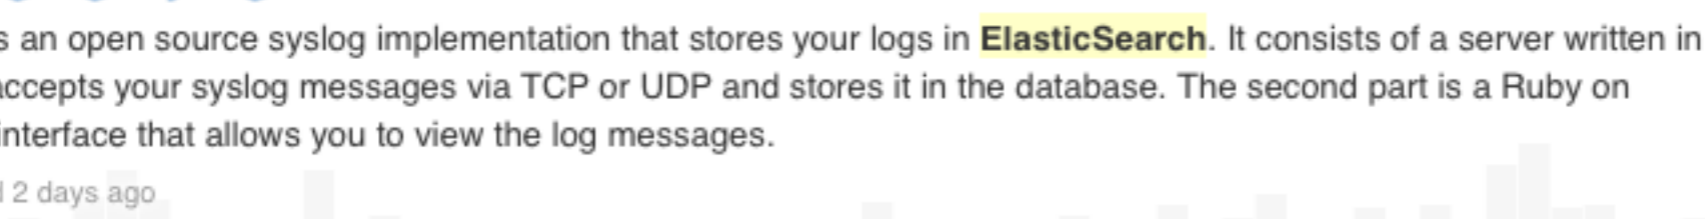
Last updated a month ago



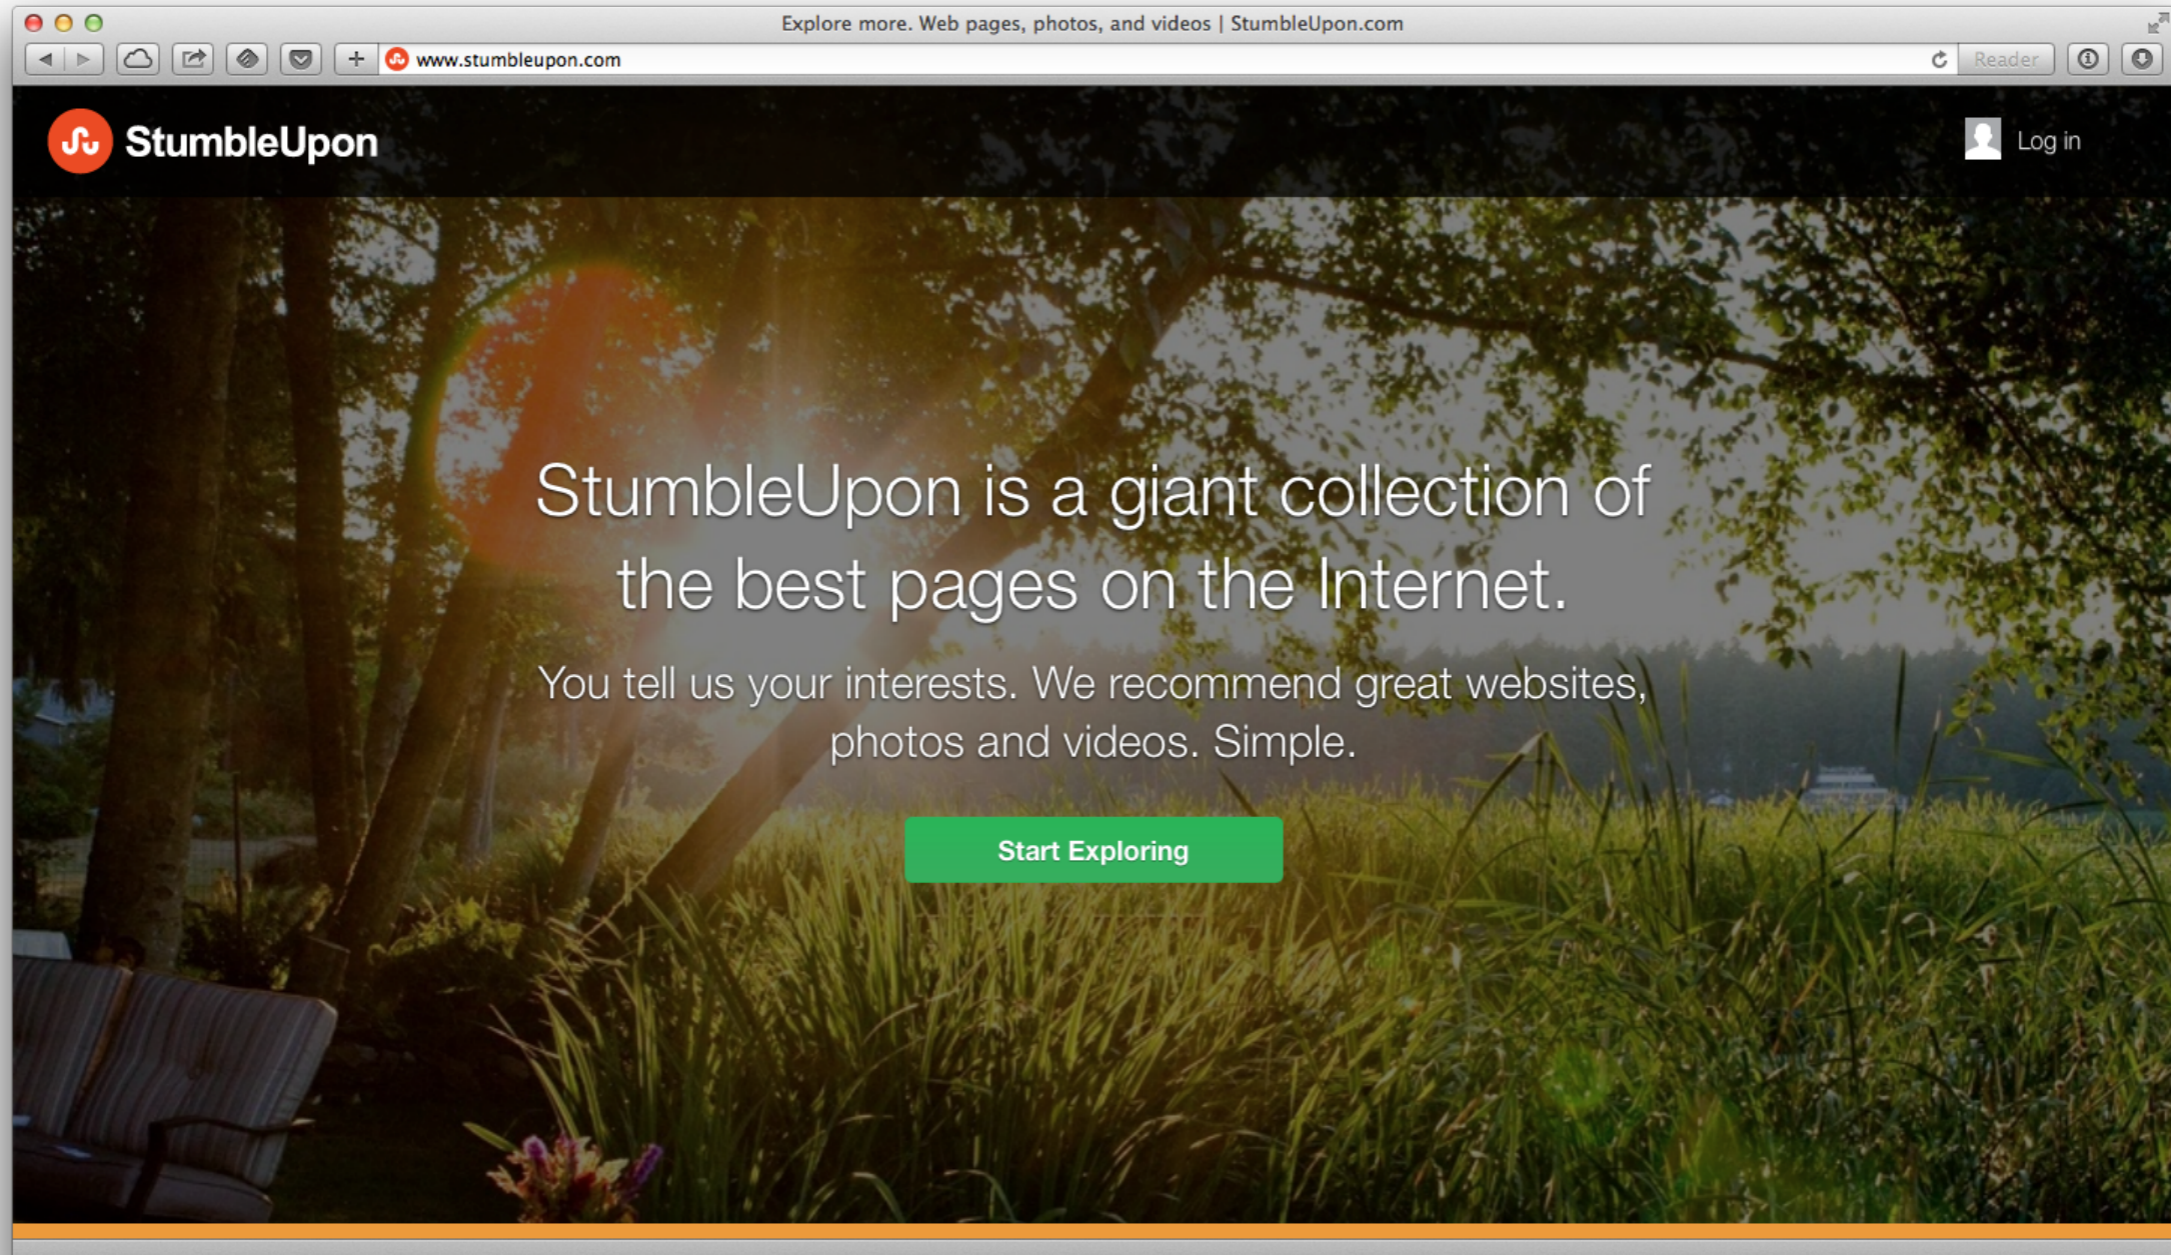
 **Graylog2/graylog2-server** Java ★ 859 🍷 143

Graylog2 is an open source syslog implementation that stores your logs in **ElasticSearch**. It consists of a server written in Java that accepts your syslog messages via TCP or UDP and stores it in the database. The second part is a Ruby on Rails web interface that allows you to view the log messages.

Last updated 2 days ago



Recommendation



Geospatial search

The screenshot displays the Foursquare website interface for a geospatial search. The browser address bar shows the URL: `foursquare.com/explore?mode=url&near=Amsterdam%2C%20Netherlands&q=Pub`. The search bar contains the text "Pub".

Suggestions for Pub
Or try: [nightlife](#), [dinner](#), [publar](#), [beer](#), [bar](#), [drinks](#), [restaurant](#), [lunch](#), [ужин](#), [popular on sundays](#)
Show me: [Specials](#) [Haven't Been](#) [Friends](#) [Price](#) [Open Now](#) [Saved](#)

Pub Recommendations:

- Molly Malone's Irish Pub** (7.8 rating, Oudezijds Kolk 9)
 - "Best Irish **stew** in Amsterdam." (2 tips)
 - "... the huge **St. Nicolas Church** in front of Central..." (1 tip)
- BeerTemple** (9.0 rating, Nieuwezijds Voorburgwal 250)
 - "... choice of **American beers**, very nice place, and..." (3 tips)
 - "Have the **milk stout** and white thal! Delicious!" (2 tips)
 - "There are 25 **American Craft Beers** on tap here,..." (2 tips)
- O'Reilly's** (8.1 rating, Paleisstraat 103)
 - "Really good **Pub in Amsterdam**, Amazing choice of..." (3 tips)
 - "**Great football venue**. Plenty of..." (2 tips)
 - "**Guinness** here is delicious!" (4 tips)

The map shows Amsterdam with numerous numbered location pins (1-30) indicating various pub locations. The city name "Amsterdam" is prominently displayed on the map.

Log analysis



Data analysis



Log analysis



Ralph Meijer

using elasticsearch and logstash to
serve billions of searchable events

<http://www.elasticsearch.org/blog/using-elasticsearch-and-logstash-to-serve-billions-of-searchable-events-for-customers/>

Percolator



Get email alerts for new **Elasticsearch Jobs**

Your Email:

You can cancel anytime



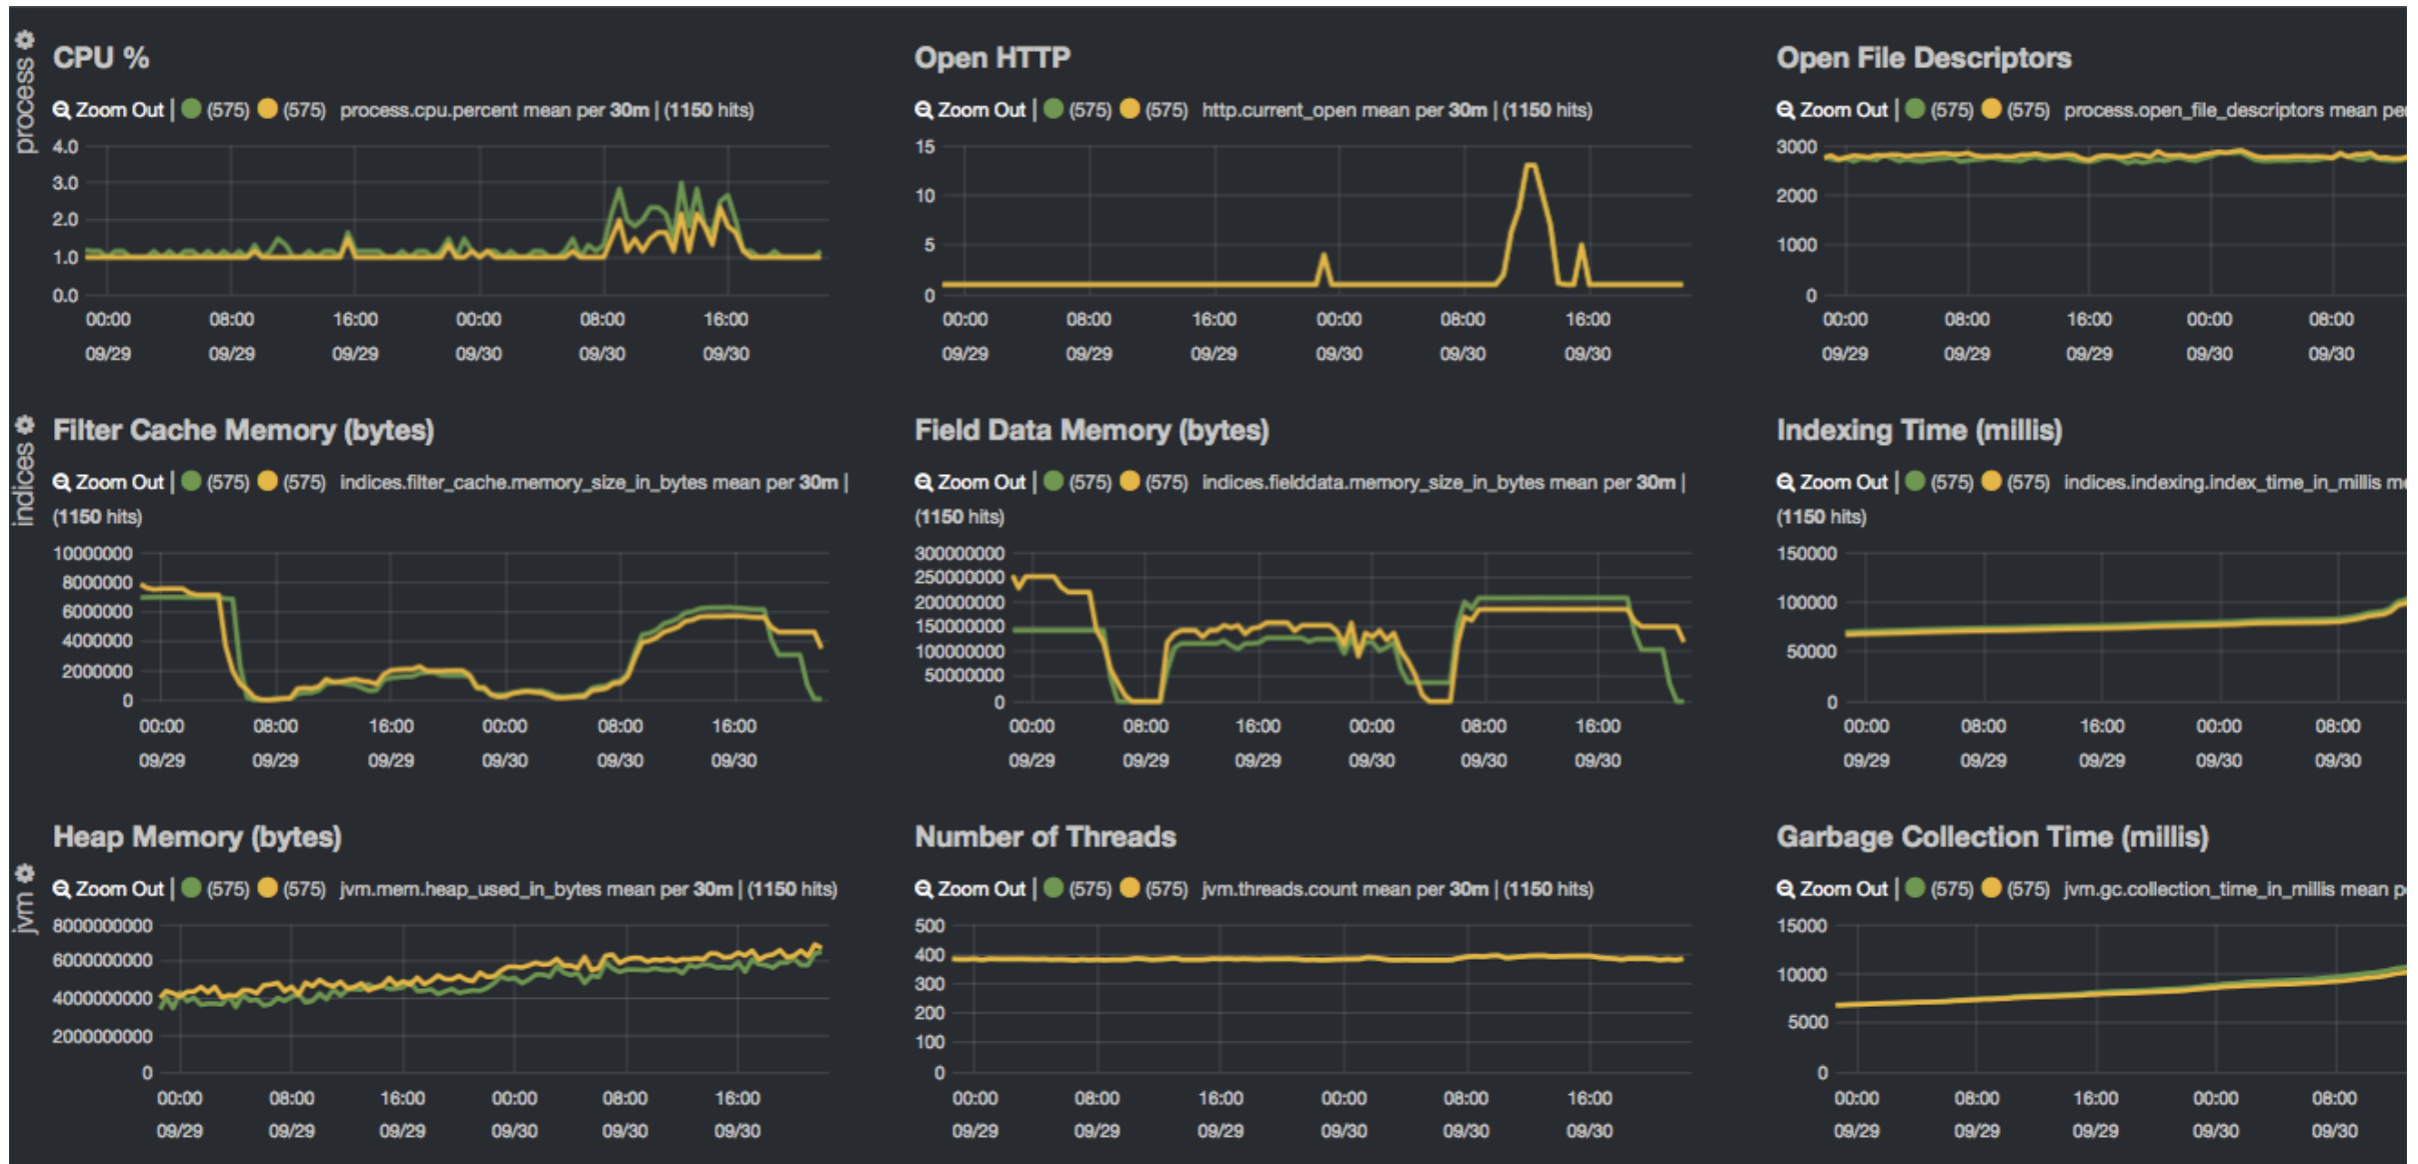
Daily



Weekly

Save

Percolator



<https://github.com/abronner/elasticsearch-monitoring>

Percolator

The screenshot shows a web browser window displaying the news page of the University of Amsterdam. The browser's address bar shows the URL www.uva.nl/nieuws-agenda/nieuws/uva-nieuws/nieuws.html. The page header includes the university logo and navigation links for 'Studenten', 'Medewerkers', 'Extranet', 'Alumni', 'Bibliotheek', 'Contact', and 'English home'. A search bar is located in the top right corner.

The main navigation bar contains links for 'Onderwijs', 'Onderzoek', 'Nieuws & Agenda', 'Over de Uva', and 'Disciplines'. Below this, a breadcrumb trail reads 'Home / Nieuws & Agenda / Nieuws / UvA Nieuws'. A large banner image shows a newspaper rack with various titles like 'EL PAIS' and 'Nieuwsblad'.

The 'Nieuws' section features a sub-header with links: 'UvA Nieuws', 'Hoogleraarsbenoemingen', 'In gesprek met ...', 'Folia', 'Kijk de Uva', and 'Amsterdamlezingen'. A 'Filter op:' section on the left includes a search bar and a list of categories: 'Onderwijs (121)' and 'Onderzoek (357)'. The main content area displays a list of news items, with the top article titled 'UvA, TU Delft en STW tekenen licentieovereenkomst met chemieconcern BASF' dated 29 okt 2013. A second article, 'Enzym biedt inspiratie voor succesvolle katalyse', is partially visible below.

On the right side of the page, there is a small vertical widget with a thumbs-up and thumbs-down icon.

NoSQL

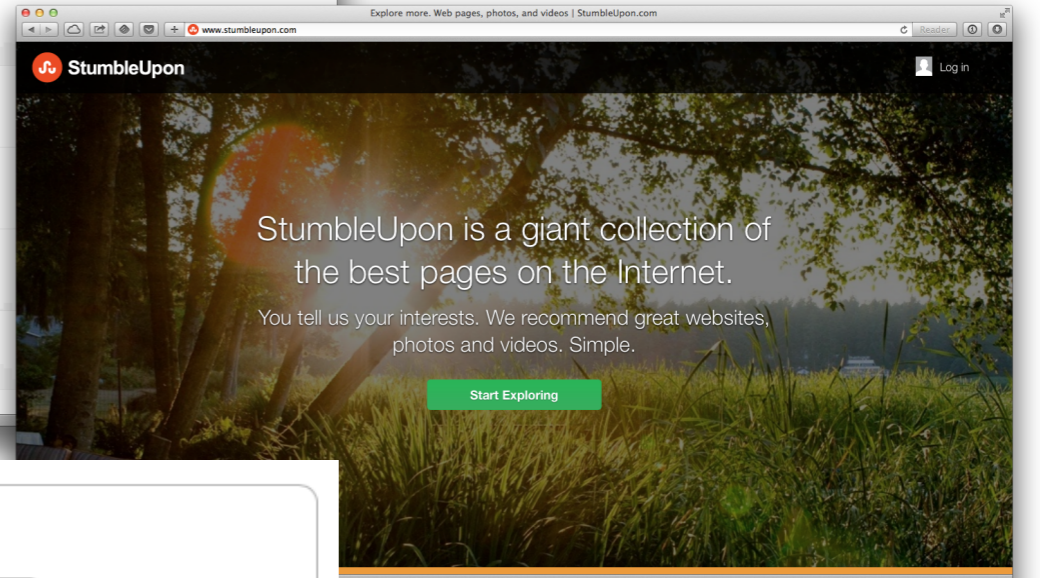
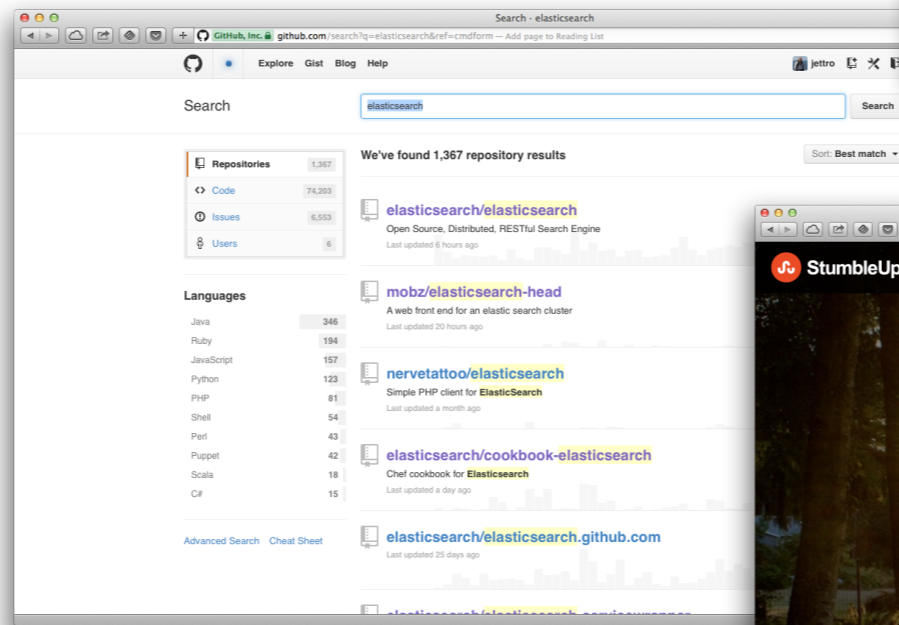
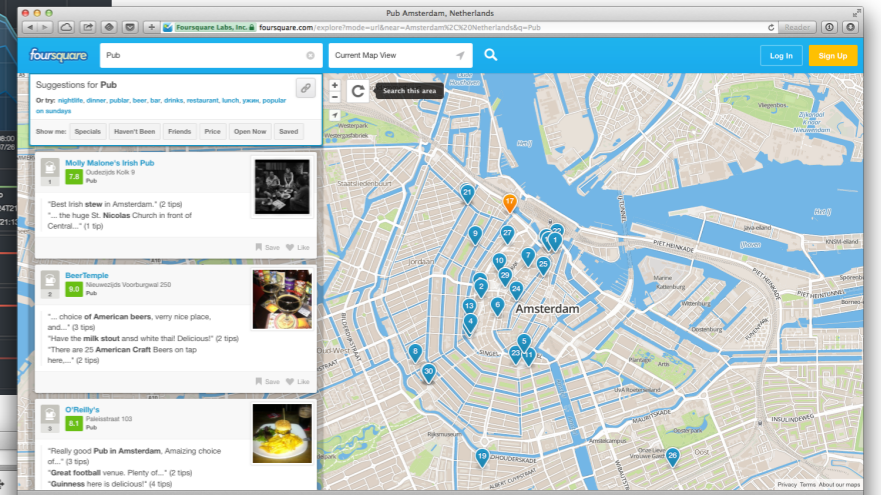
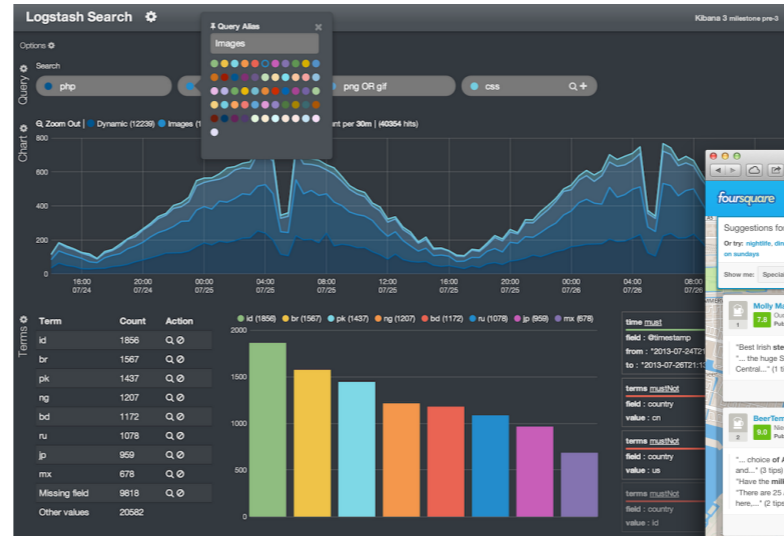


Beeld en Geluid

Store the meta-data of all image, audio and video material

The right tool

The job



✉ **Get email alerts for new Elasticsearch Jobs**

Your Email:

You can cancel anytime Daily Weekly

Questions?

