



Brandwatch

# Detecting events on the Web in real-time with **Java, Kafka & ZooKeeper**

Dr. James Stanier | [brandwatch.com](https://brandwatch.com) | [james@brandwatch.com](mailto:james@brandwatch.com)

# Coming Up/

- Me, Brandwatch and new problems
- Moving to Kafka
- Processing data
- Distributing work
- Finding meaning



Who?



**Dr. James Stanier**

VP Engineering,  
Product (Backend)

Brandwatch

@jstanier | james@brandwatch.com

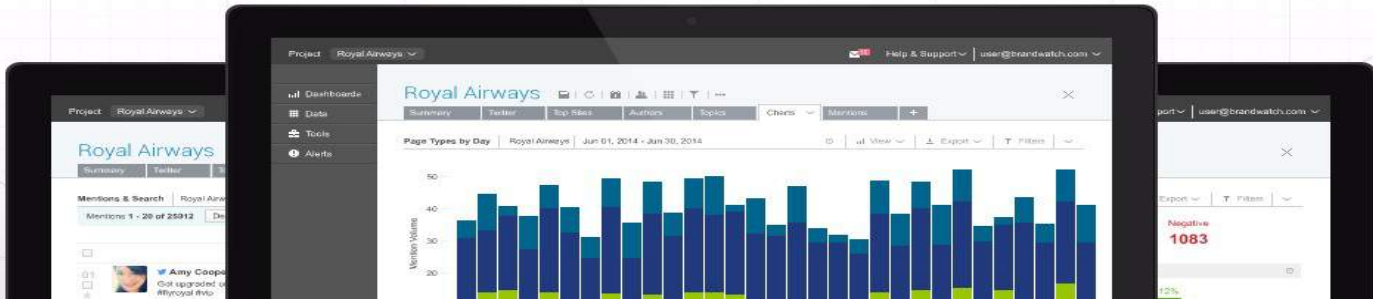


FREE DEMO >

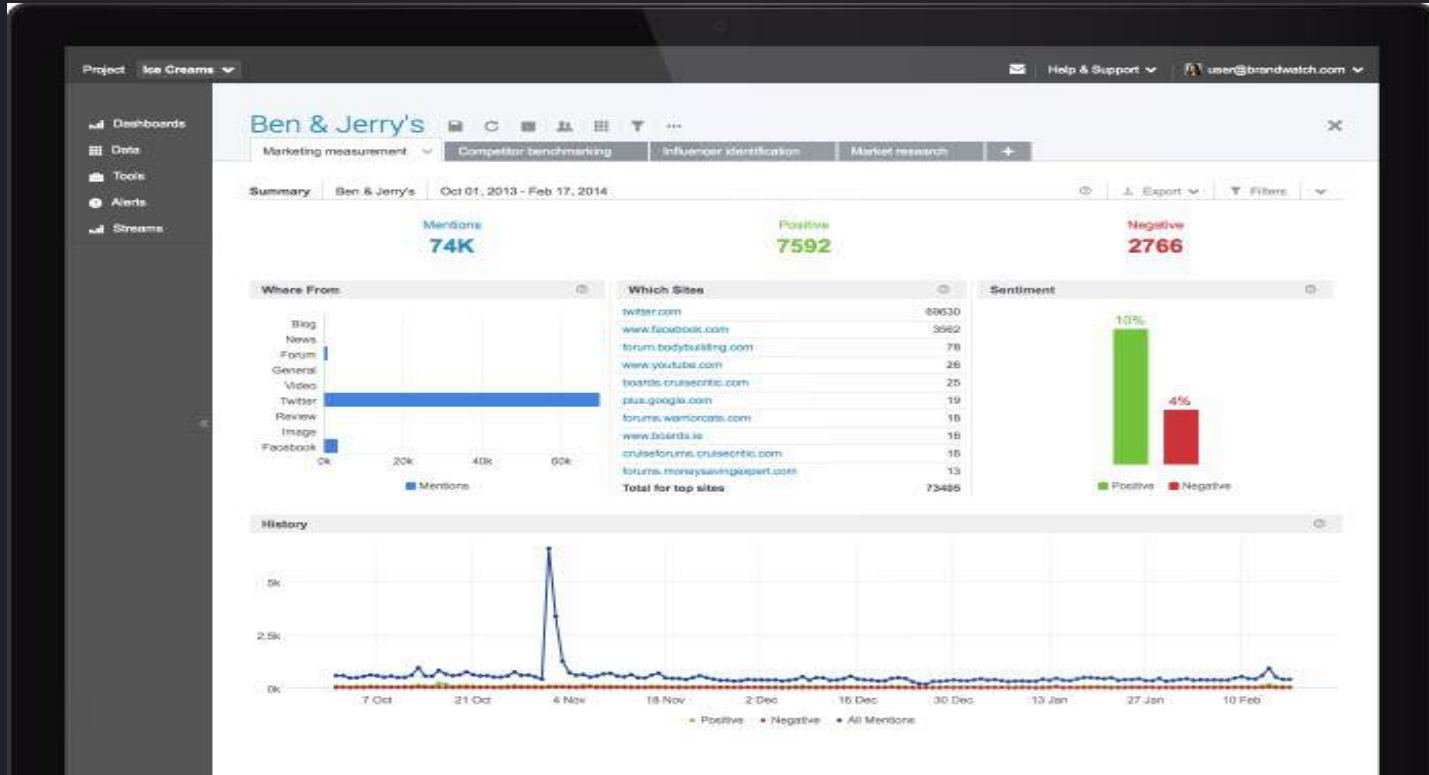
# Use social listening to power any decision

Find meaning in the conversations that matter.  
Act with confidence.

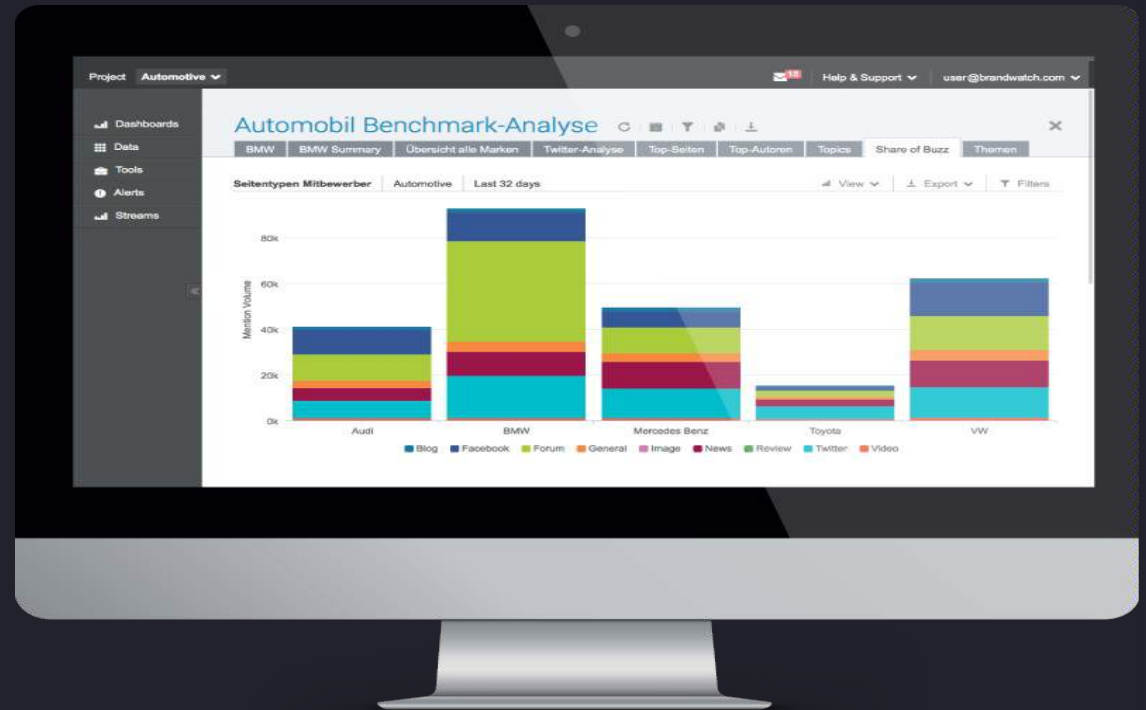
START LISTENING >



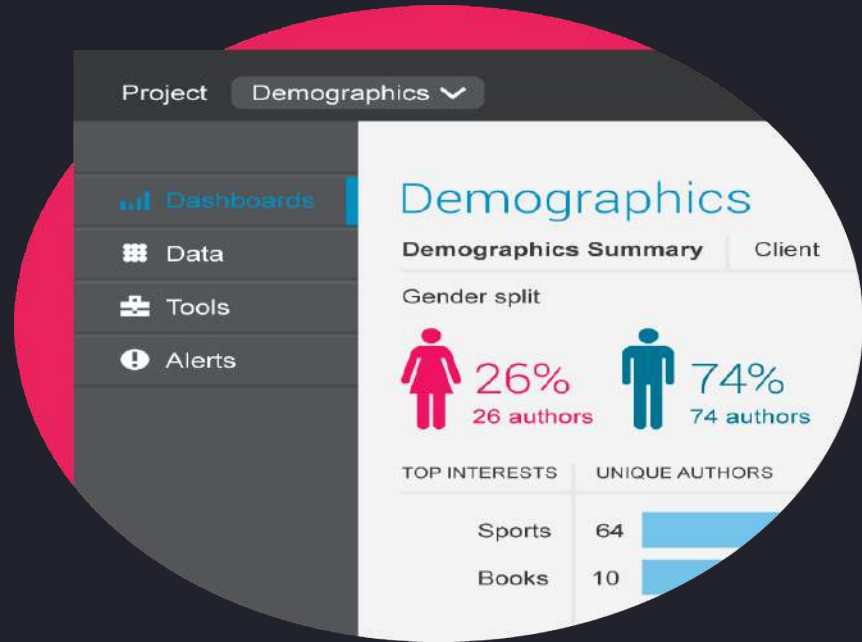
# Data/ Presentation



# Data/ Aggregation



# Data/ Classification





# Data/ Not just top level metrics

The screenshot shows a social media monitoring dashboard for the project "Ice Creams". The main content area displays "Mentions & Search" for "Ben & Jerry's" from June 01, 2013, to June 30, 2013. It shows 28 of 25912 mentions. The dashboard includes a table of mentions with columns for ID, Title, Sentiment, and Page. The table lists three mentions: one negative mention from Jennifer? (@jennywilson) about vodka prices, one neutral mention from Lynne McFarlane (@lynnemcfarlane1) about watching friends with a boyfriend, and one positive mention from Lyndsey Leigh Moran on Facebook about Ben & Jerry's ice cream. The right sidebar contains filters for "Data" (Group), "Date" (Jun 01, 2013 - Jun 30, 2013), and "Filters" (Sentiment & Page Type, Author, Impact, Forum thread, Blog).




Project: Ice Creams

Ben & Jerry's

Summary | Twitter | Top Sites | Authors | Topics | Charts | Mentions

Mentions & Search | Ben & Jerry's | Jun 01, 2013 - Jun 30, 2013

Mentions 1 - 28 of 25912

	TITLE	SENTIMENT	PAGE
01	 <b>Jennifer?</b> @jennywilson when a double vodka is cheaper than ben and jerrys... we know there's an issue Retweets: 0   Replies: 0   Impressions: 328	Negative	2011 Jun 2013-06-30
02	 <b>Lynne McFarlane</b> @lynnemcfarlane1 ben and jerrys & watching friends with benefits, who needs a boyfriend when you have movies and ice cream Retweets: 0   Replies: 1   Impressions: 117	Neutral	2011 Jun 2013-06-30
03	 <b>Lyndsey Leigh Moran</b> - facebook.com I'm goin to start of by sayin i want deep fried oreos followed by I'm goin to start of by sayin i want deep fried oreos followed by ben and jerrys awesome some choc icecream with pudding brownie chunks and almonds the i want to add a order of wings and then some pickets until i get it ima be a very unpleasant woman	Positive	30 Jun 2013-06-30

Filters

- Sentiment & Page Type
- Author
- Impact
- Forum thread
- Blog

## Data/ The numbers

- 50+ Java Web Crawlers
- 10+ Historical crawlers for new queries
- Twitter via GNIP (now Twitter), Weibo, Disqus and more
- 80M+ query matches per day





### TWITTER Mentions

Globe

**Twitter Mention 1**  
RT @brandwatch: @Globe what "Globe" does? "Globe" is a brand name of a company that provides services around the world. (See the link below)

**Twitter Mention 2**  
@brandwatch: @Globe what "Globe" does? "Globe" is a brand name of a company that provides services around the world. (See the link below)

**Twitter Mention 3**  
@brandwatch: @Globe what "Globe" does? "Globe" is a brand name of a company that provides services around the world. (See the link below)

**Twitter Mention 4**  
@brandwatch: @Globe what "Globe" does? "Globe" is a brand name of a company that provides services around the world. (See the link below)

**Twitter Mention 5**  
@brandwatch: @Globe what "Globe" does? "Globe" is a brand name of a company that provides services around the world. (See the link below)

**Twitter Mention 6**  
@brandwatch: @Globe what "Globe" does? "Globe" is a brand name of a company that provides services around the world. (See the link below)

**Twitter Mention 7**  
@brandwatch: @Globe what "Globe" does? "Globe" is a brand name of a company that provides services around the world. (See the link below)

### Watch and Smartwatch Reviews

September 28 11:27am



Device	Score	Reviews
Apple Watch	8133	164
Sony Smartwatch 3	7271	220

### Demo project

Daily 1

Daily 2

### Demo project

Daily 1

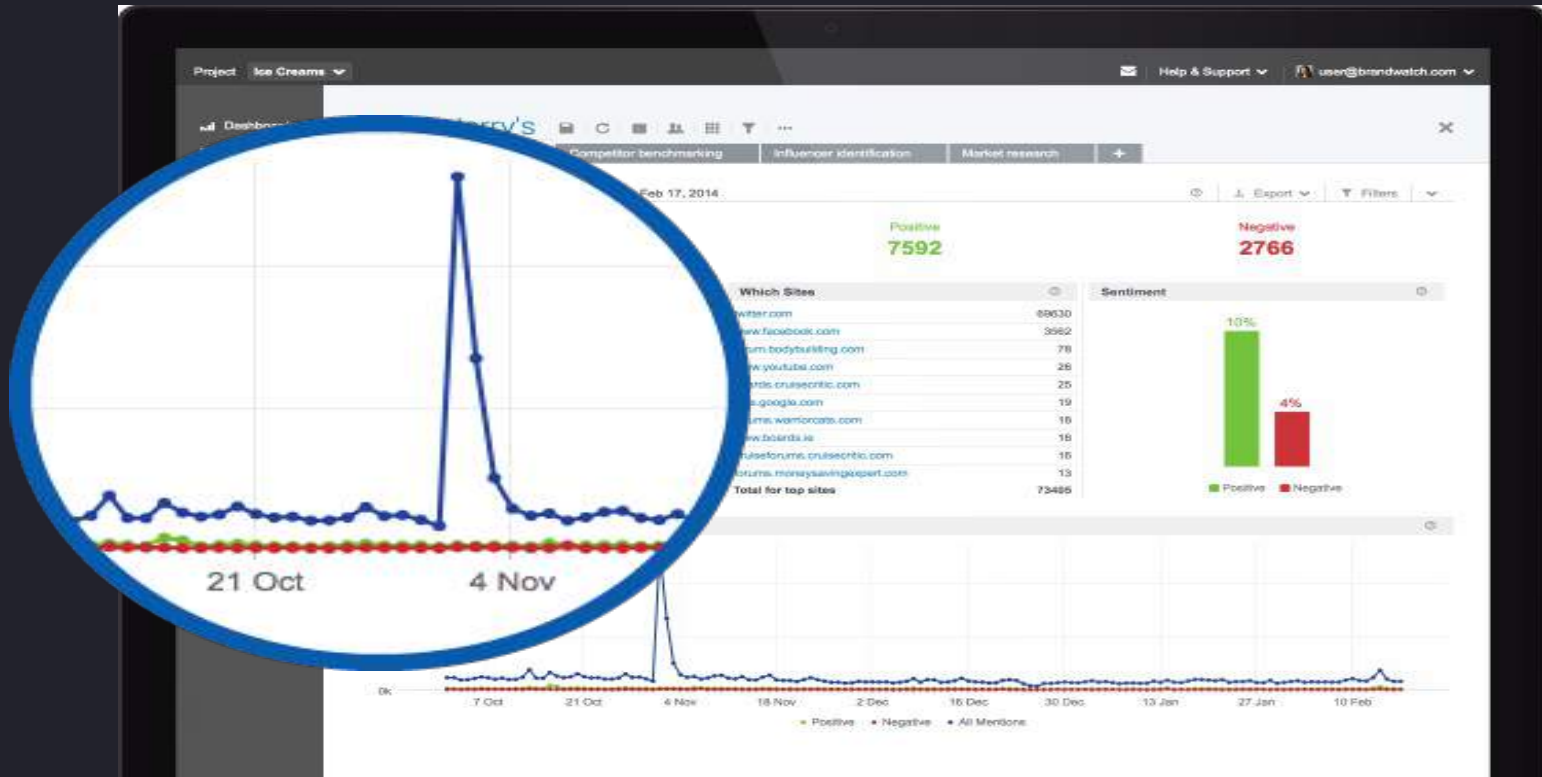
Daily 2





A new challenge

# The challenge/ The signal from the noise

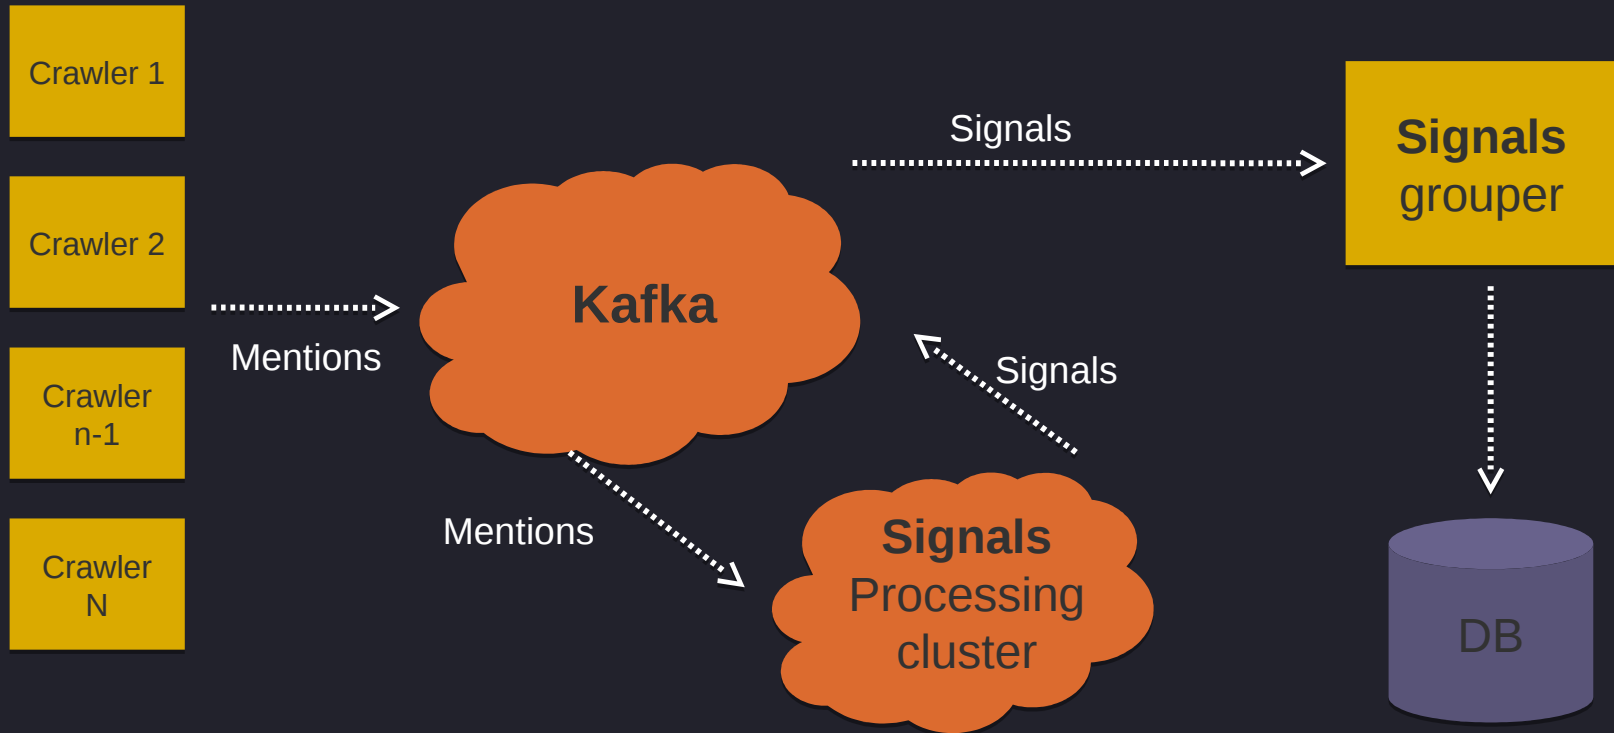


## The challenge/ at scale

- 130K+ user queries
- 80M+ mentions per day
- Polling the data stores for mentions for all queries takes 8hrs for one pass



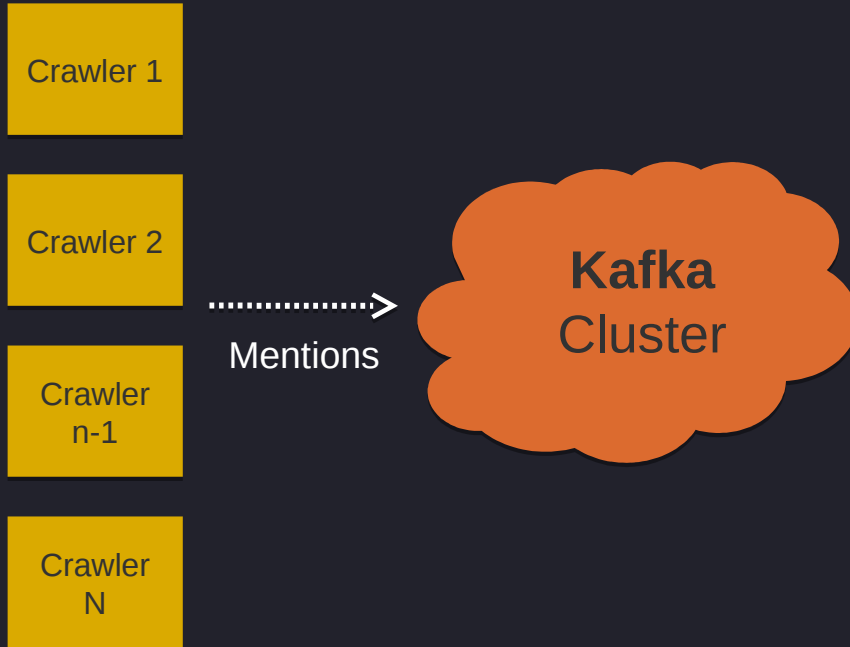
# The Problem/ How we handled it...





Kafka

# Step 1/ Kafka



# Kafka/ What is it?

- **Apache Kafka** is a publish-subscribe messaging system rethought as a distributed commit log
- Apache top level project **November 2013**
- Started at **LinkedIn**


## Kafka/ is...

- **Fast:** hundreds of MBs read/write per second from thousands of clients
- **Scalable:** clustered, partitioned over many machines, expanded without downtime
- **Durable:** messages persisted to disk and replicated in cluster

# Kafka/ Written to disk?

FIGURE 3

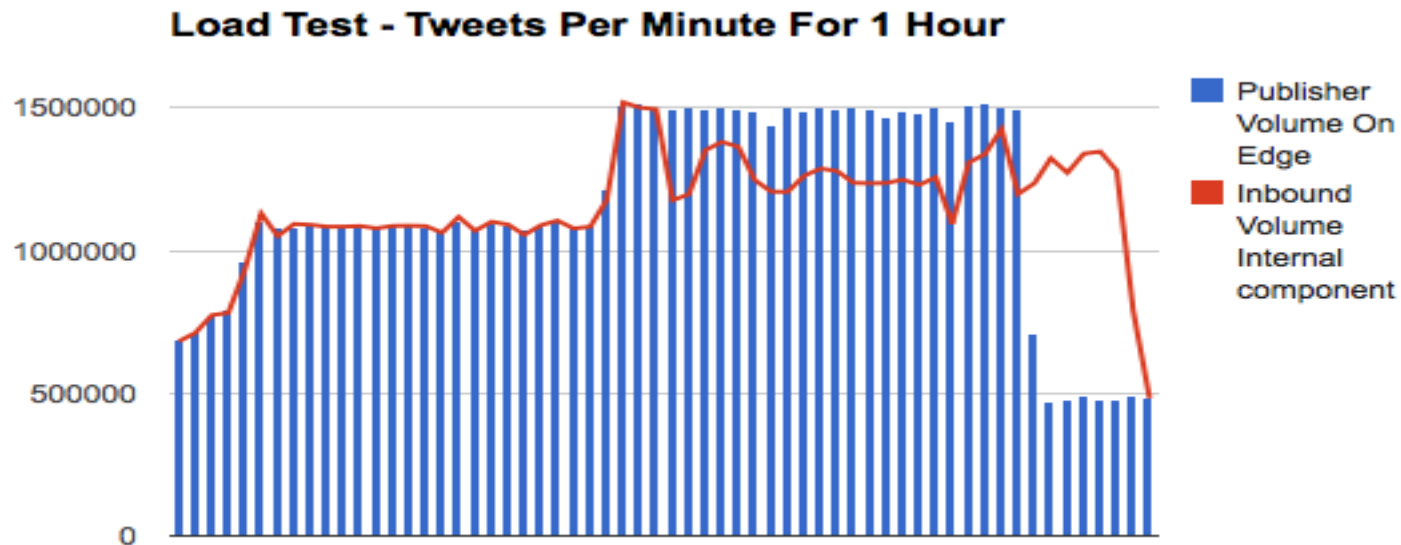
Comparing Random and Sequential Access in Disk and Memory



Note: Disk tests were carried out on a freshly booted machine (a Windows 2003 server with 64-GB RAM and eight 15,000-RPM SAS disks in RAID5 configuration) to eliminate the effect of operating-system disk caching. SSD test used a latest-generation Intel high-performance SATA SSD.

<http://q.acm.org/detail.cfm?id=1563874>

# Kafka/ Bending, not breaking



<http://engineering.gnip.com/tag/kafka/>

# Kafka/ Sending from the crawlers

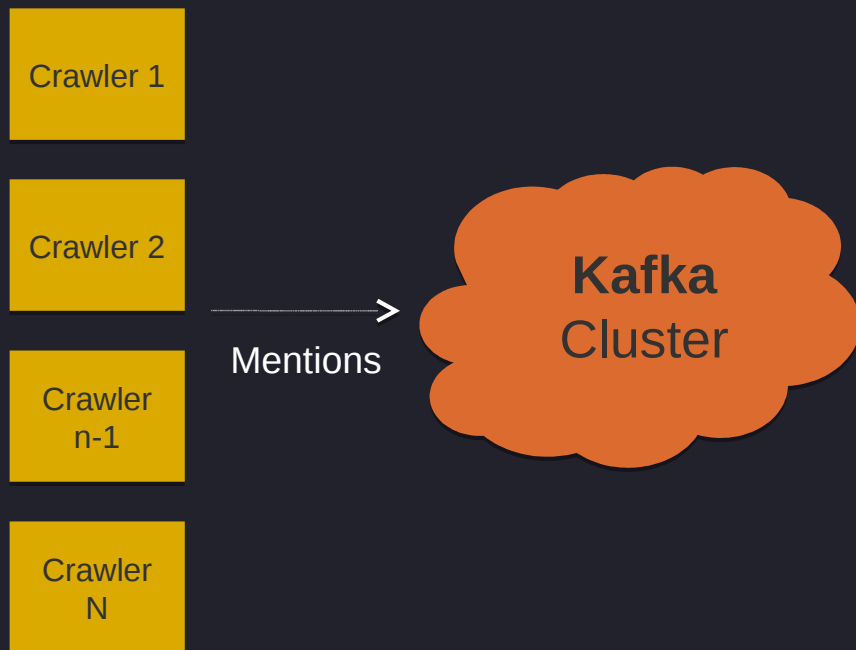
```
String message = toJson(...);
```

```
KeyedMessage<String, String> message = new  
    KeyedMessage<String, String>("query.mentions",  
    queryId, message);
```

```
producer.send(message);
```

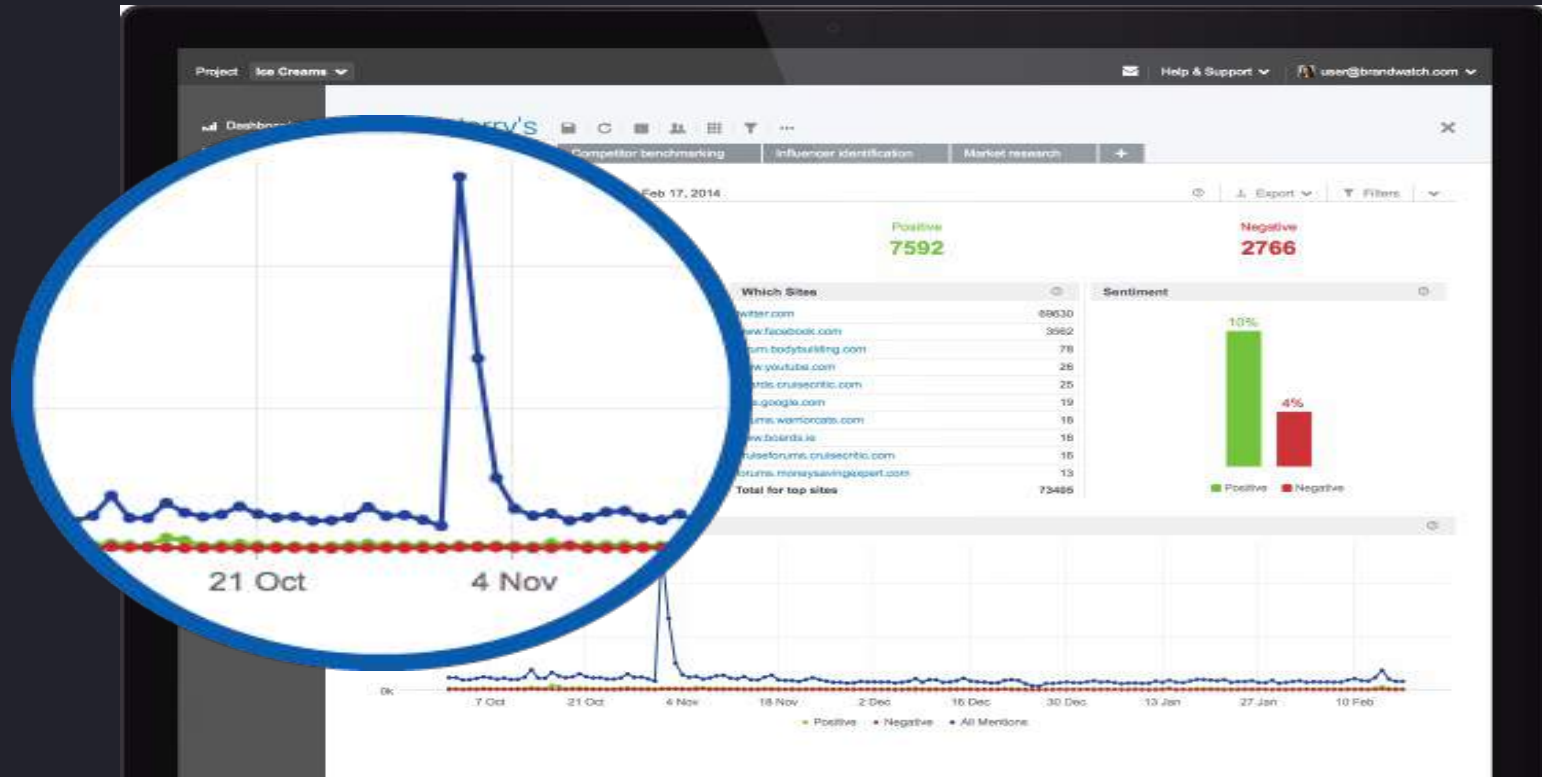


# Step 1/ Done

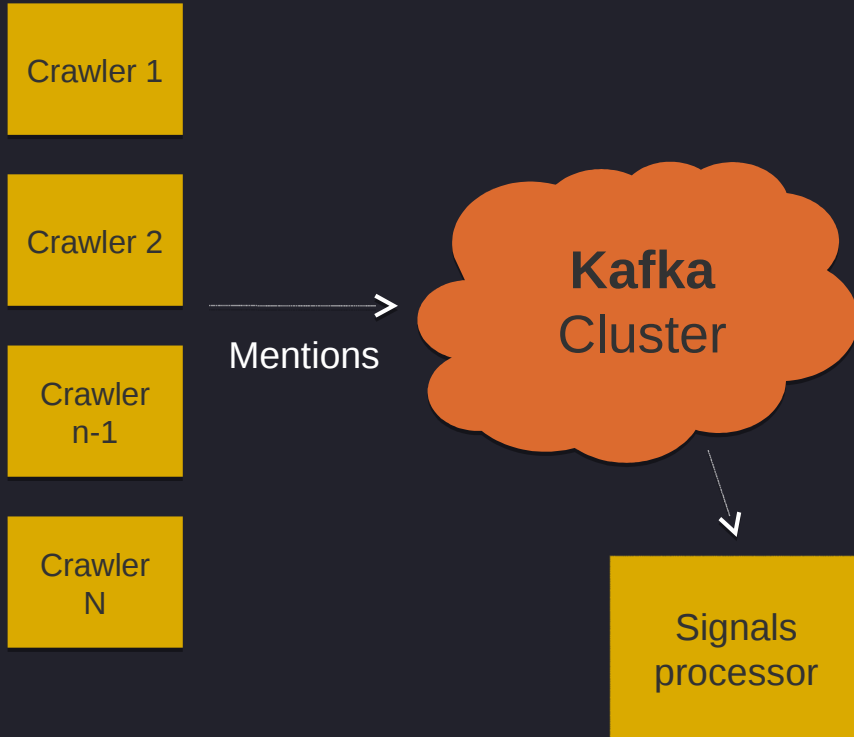


# Processing

# Processing/ What's happening now?



# Step 2.1/ One processing JVM



# Processing/ A wild tweet appears!

Mention

**date:** 01/06/2015 16:05

**pageType:** twitter

**author:** @berlinperson

**hashtags:** [#berlinbuzzwords, #amazingtalk,  
#greatshoes]

**mentionedTweeters:** [@jstanier]

**text:** “@jstanier is at #berlinbuzzwords  
#amazingtalk #greatshoes”

# Processing/ Storing hashtags

```
Map<Date, Multiset<String>>
```

Initialise with the last 24 hours

# Processing/ Storing hashtags

Map<Date, Multiset<String>>

Mention

date: 01/06/2015 4:10PM

hashtags: [#berlinbuzzwords, #amazingtalk,  
#greatshoes]

# Processing/ Storing hashtags

Map<Date, Multiset<String>>

Mention

date: 01/06/2015 4:10PM

hashtags: [#berlinbuzzwords, #amazingtalk,  
#greatshoes]

```
add("#berlinbuzzwords")  
add("#amazingtalk")  
add("#greatshoes")
```



# Processing/ Cycling the buckets

```
@Scheduled(cron = "0 0 * * * *")  
  
public void cycleBuckets() {  
    Date oldest = buckets.lastKey();  
    removeBucket(oldest);  
  
    DateTime newest = new  
        DateTime(buckets.firstKey());  
    addBucket(newest.plusHours(1).toDate());  
  
}
```

# Processing/ Detecting spikes

- At scheduled intervals
- For each #
  - Convert to a timeseries [5, .... 1002, 5499]
  - Compare previous hour to history
  - Give a score to it
- If score  $>$  threshold, it's interesting
- Send it on a new Kafka topic

# Processing/ What we just did

#hashtag  
data  
model

# Processing/ But we also track...

#hashtag data model	country data model	author data model	page type data model	sentiment data model	volume data model	link share data model
---------------------------	--------------------------	-------------------------	----------------------------	----------------------------	-------------------------	-----------------------------

# Processing/ ...for one query

“Berlin Buzzwords” query

```
graph TD; A["Berlin Buzzwords query"] --- B["#hashtag data model"]; A --- C["country data model"]; A --- D["author data model"]; A --- E["page type data model"]; A --- F["sentiment data model"]; A --- G["volume data model"]; A --- H["link share data model"];
```

#hashtag  
data  
model

country  
data  
model

author  
data  
model


page type  
data  
model

sentiment  
data  
model

volume  
data  
model

link share  
data  
model

# Processing/ 100K+ queries and rising



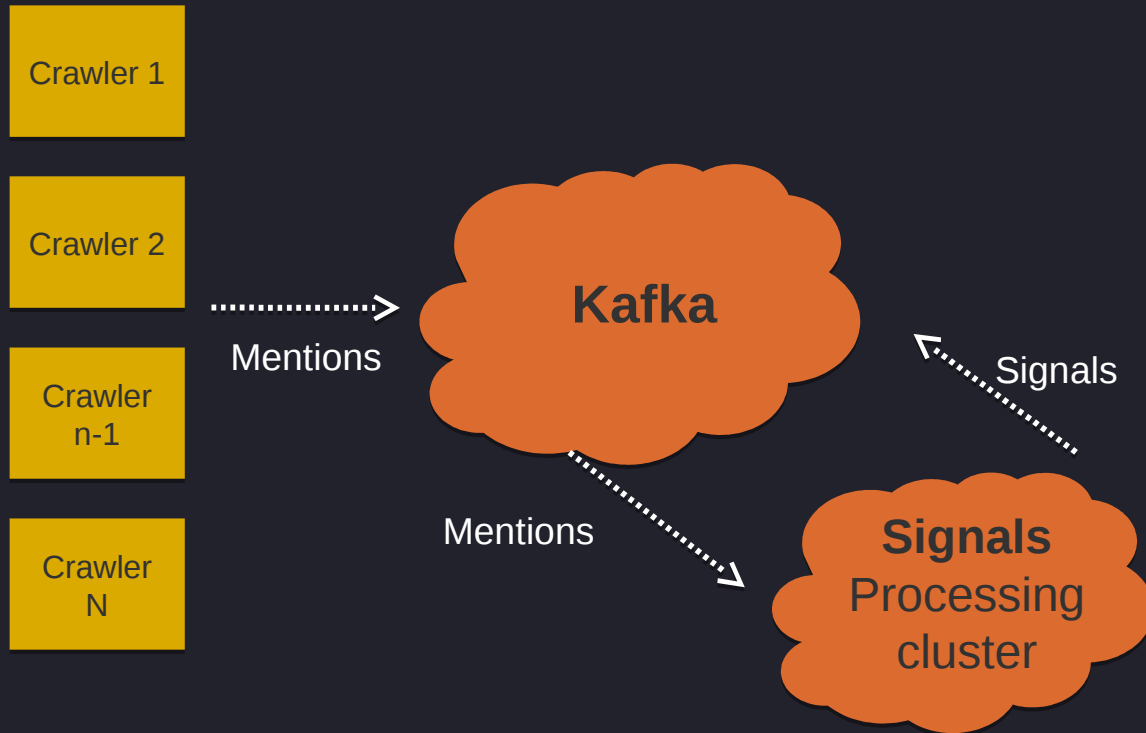
Processing/ We need more JVMs

But how do we **share** the workload?

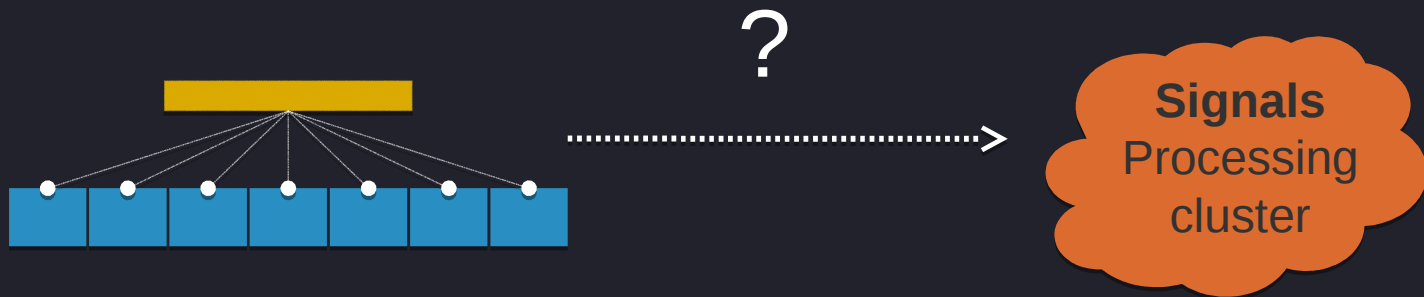
# Distribution of work



## Step 2.2/ A cluster of processing JVMs



# Distribution/ An atomic unit of work



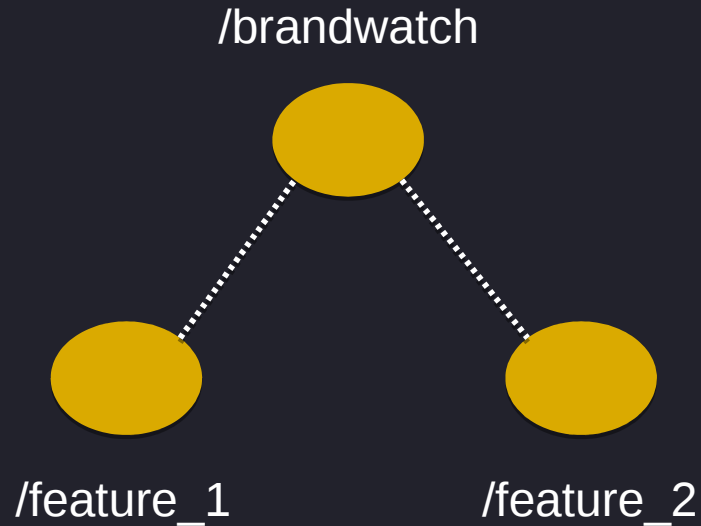
# Distribution/ Leader election

A way of deciding who is the **leader** for a task in a group of **distributed nodes**

# Distribution/ Zookeeper

A way of **coordinating** and **managing** distributed applications

# Zookeeper/ It's like a file system



# Zookeeper/ At the command line

```
[zk: localhost:12181(CONNECTED) 6] ls /  
[zookeeper, admin, consumers, brandwatch, controller, brokers, controller_epoch]  
[zk: localhost:12181(CONNECTED) 7] ls /brokers  
[topics, ids]  
[zk: localhost:12181(CONNECTED) 8] █
```

# Distribution/ Recipes



Last Published: 2014-07-05



Apache ZooKeeper

## APACHE CURATOR

- About
- Getting Started
- Examples
- Recipes
- Framework
- Utilities
- Client

## DETAILS

- Error Handling
- Logging and Tracing
- Tech Notes
- Exhibitor Integration
- Source Code
- Project Team
- Project Information
- Javadoc

Wiki

Releases

## EXTENSIONS

## Recipes

Curator implements all of the recipes listed on the ZooKeeper recipes doc (except two phase commit). Click on the recipe name below for detailed documentation.

### Elections

**Leader Latch** - In distributed computing, leader election is the process of designating a single process as the organizer of some task distributed among several computers (nodes). Before the task is begun, all network nodes are unaware which node will serve as the "leader," or coordinator, of the task. After a leader election algorithm has been run, however, each node throughout the network recognizes a particular, unique node as the task leader.

**Leader Election** - Initial Curator leader election recipe.

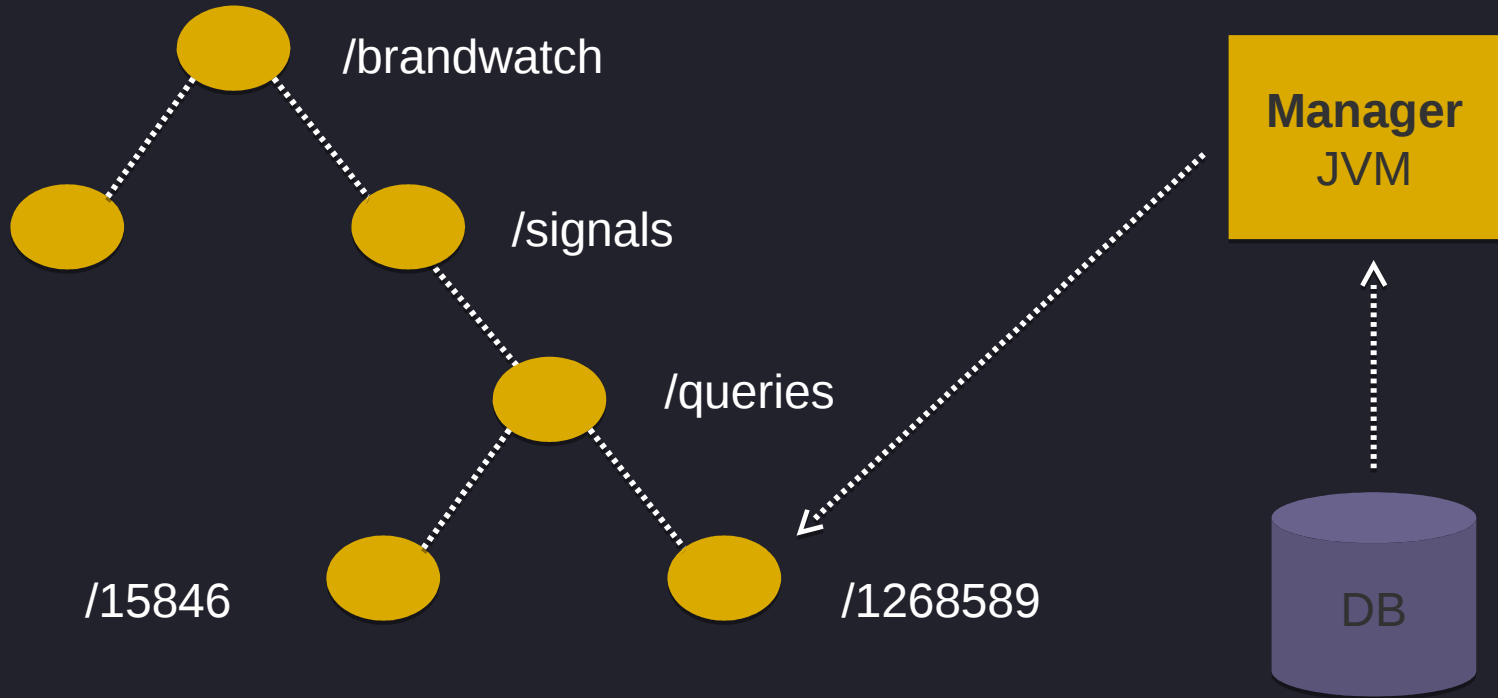
### Locks

**Shared Reentrant Lock** - Fully distributed locks that are globally synchronous, meaning at any snapshot in time no two clients think they hold the same lock.

**Shared Lock** - Similar to Shared Reentrant Lock but not reentrant.

**Shared Reentrant Read Write Lock** - A re-entrant read/write mutex that works across JVMs. A read write lock maintains a pair of associated locks, one for read-only operations and one for writing. The read lock may be held simultaneously by multiple reader processes, so long as there are no writers. The write lock is exclusive.

# Distribution/ Offering jobs





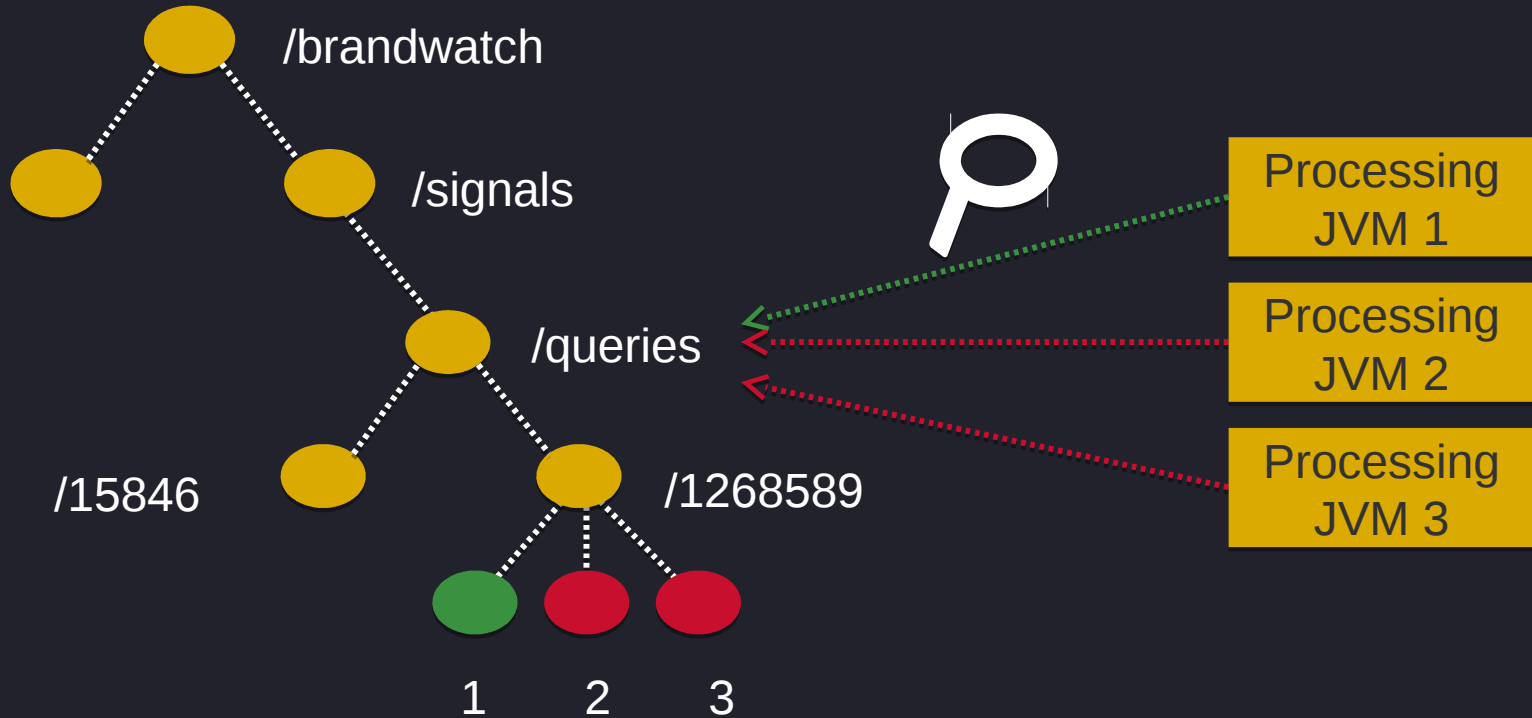
# Distribution/ PGQ

The screenshot shows the GitHub interface for the repository 'BrandwatchLtd / pgq-consumer'. At the top, it displays the repository name, a timestamp 'Feb 22, 2015, 9:46 PM GMT', and interaction buttons for 'Unwatch' (4), 'Unstar' (6), and 'Fork' (6). Below this, the repository title 'PGQ consumer for Java. — Edit' is shown, along with statistics: '38 commits', '1 branch', '2 releases', and '2 contributors'. A navigation bar includes a refresh icon, a dropdown for 'branch: master', and the repository name 'pgq-consumer / +'. A prominent blue banner indicates a 'Merge pull request #11 from jstanier/increase-unit-test-coverage'. Below the banner, a commit by 'jstanier' is detailed, listing files such as 'src', '.gitignore', '.travis.yml', 'LICENCE', 'README.md', and 'pom.xml' with their respective commit dates. The 'README.md' section is partially visible, showing the repository name 'pgq-consumer' with a 'build passing' status and a description: 'A PGQ consumer written in Java, using Spring's JdbcTemplate for database access.' On the right side, a sidebar contains links for 'Code', 'Issues' (0), 'Pull requests' (0), 'Wiki', 'Pulse', 'Graphs', and 'Settings'. At the bottom of the sidebar, it provides the 'HTTPS clone URL' as 'https://github.com/Brandw' and buttons for 'Clone in Desktop' and 'Download ZIP'.

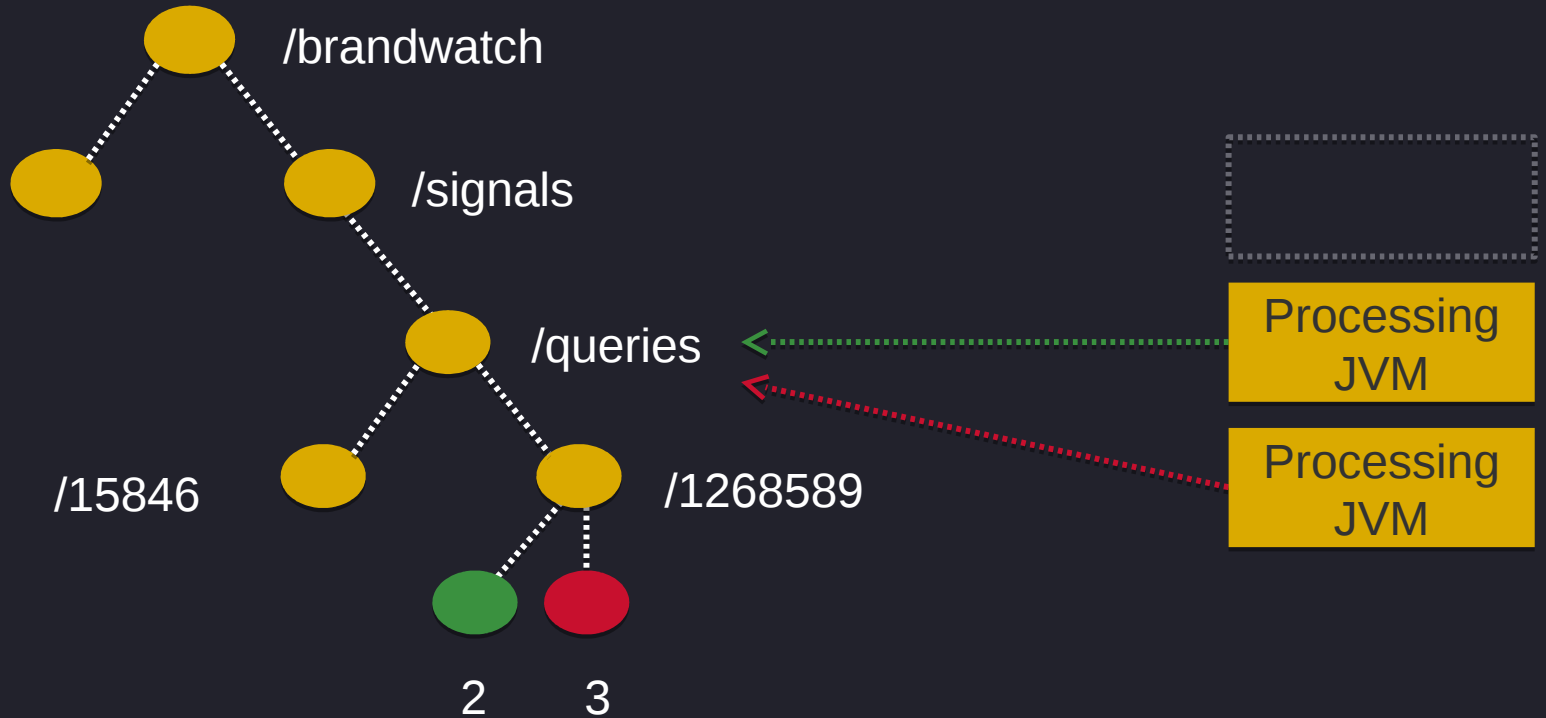
# Distribution/ Leader election 101



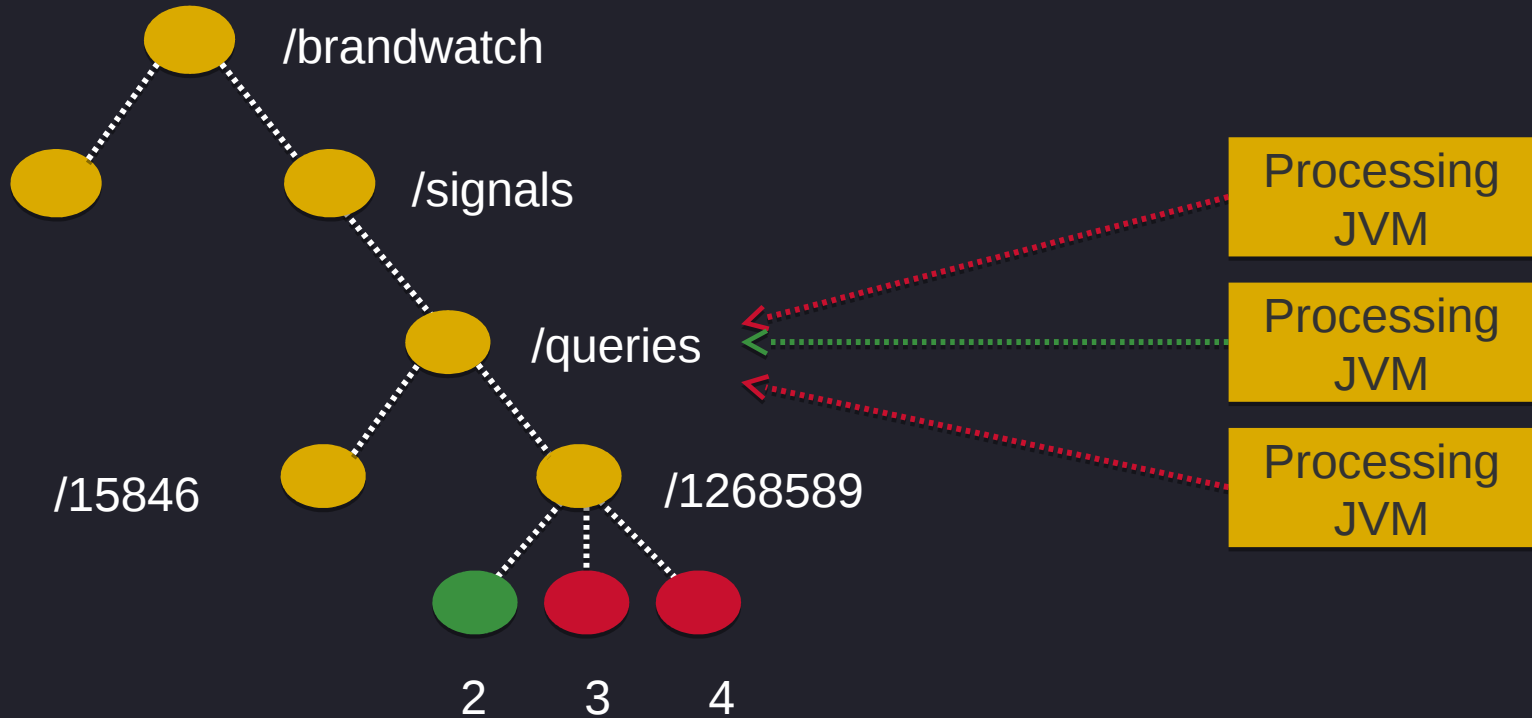
# Distribution/ Leader election 101



# Distribution/ The leader dies



# Distribution/ The dead rises again



# Distribution/ That's how we do it

Each time someone turns on the feature, we  
leader elect for processing

# Distribution/ Almost there?

We are processing **long running** jobs

What about workers getting **overloaded**?

# Distribution/ After leader election

1. Take leadership

2. Hit max queries?

- a. No – go to 3

- b. Yes – give up leadership, try again


3. Start working



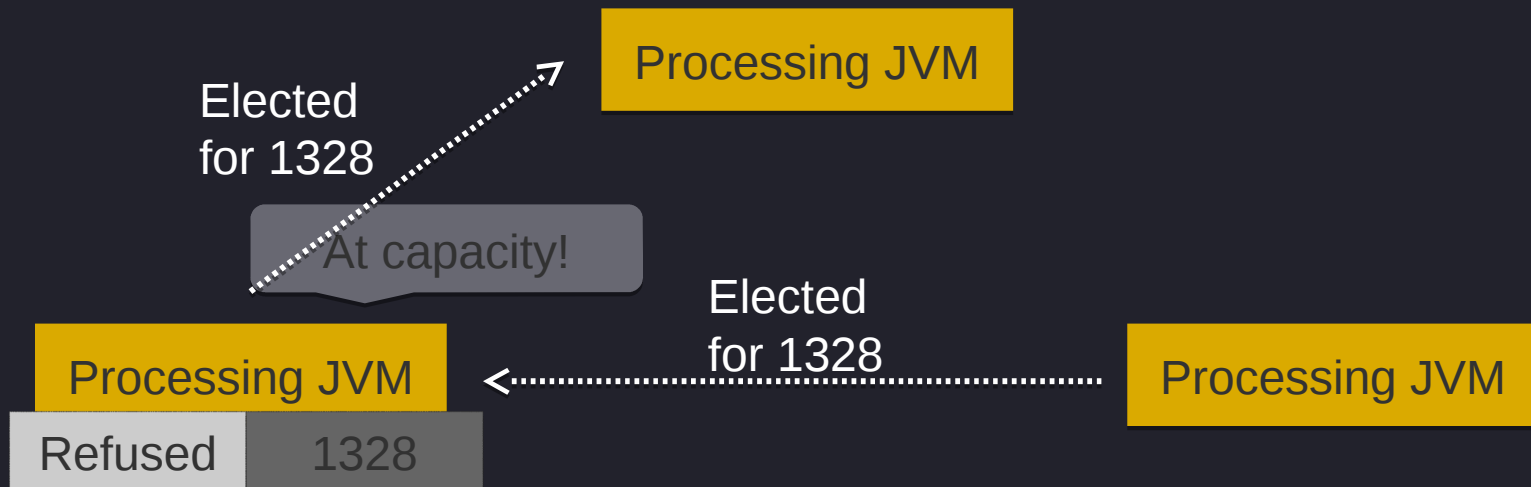
**Distribution/** Now we're almost there?

Actually, **no**...

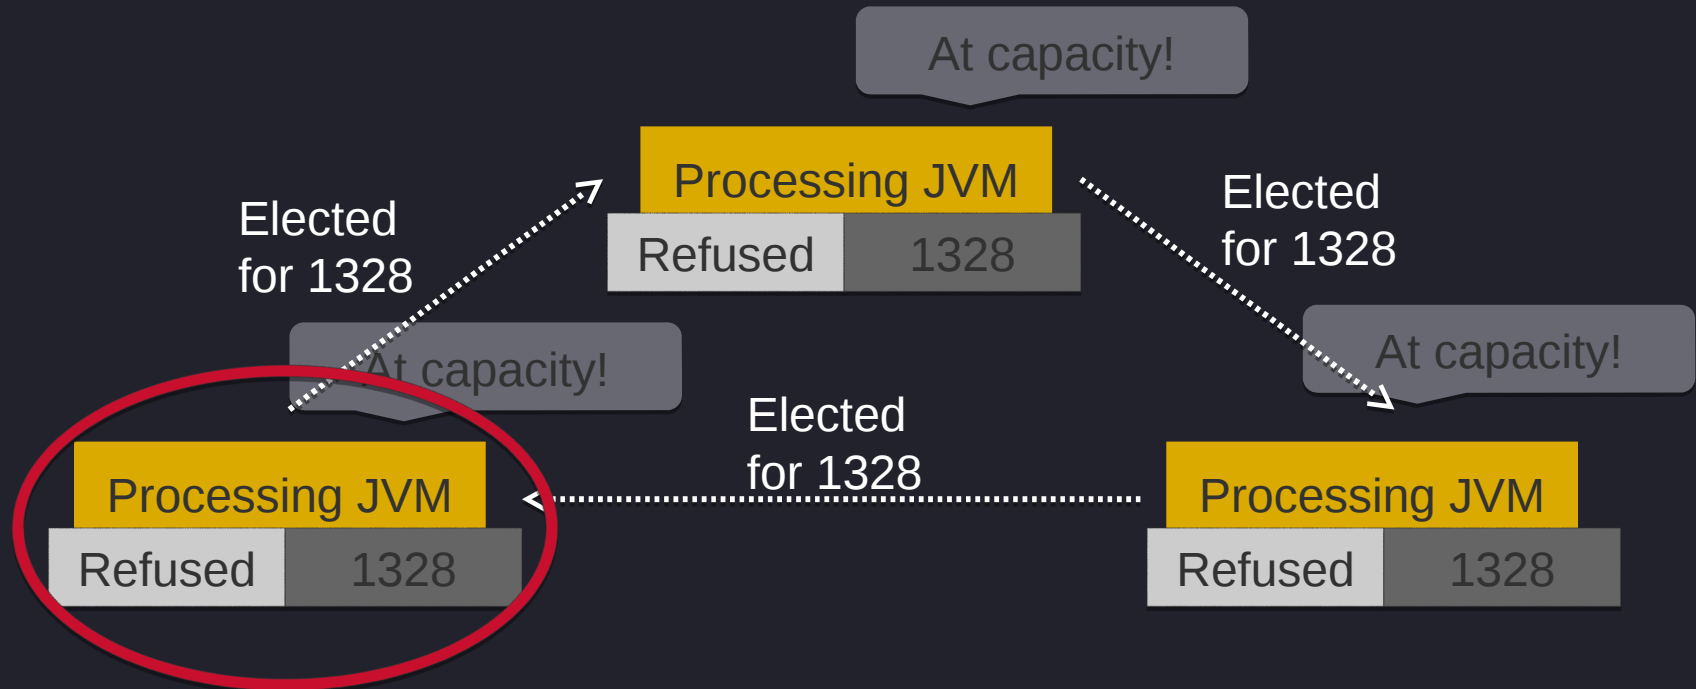
# Distribution/ Infinite election



# Distribution/ Solution



# Distribution/ Solution





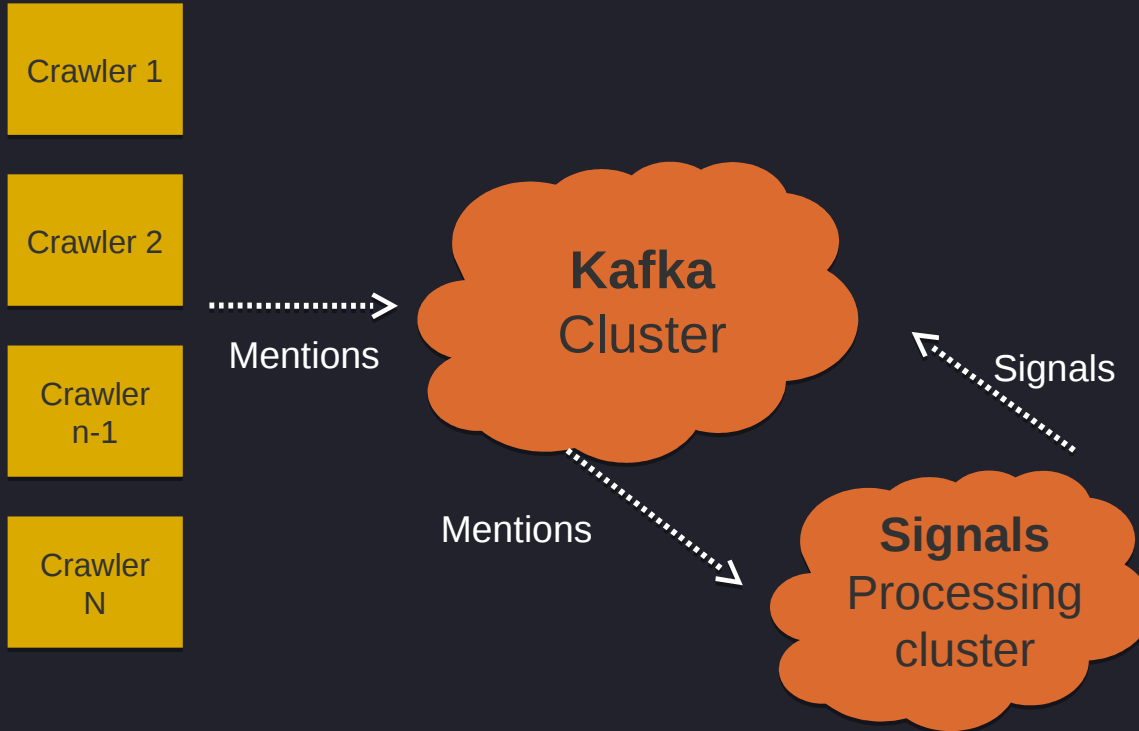
State


# State/ Snapshotting of worker data

If one worker dies, we want the other to pick up where it left off

Regular snapshotting to **HBase**

# Step 2.2/ Done!

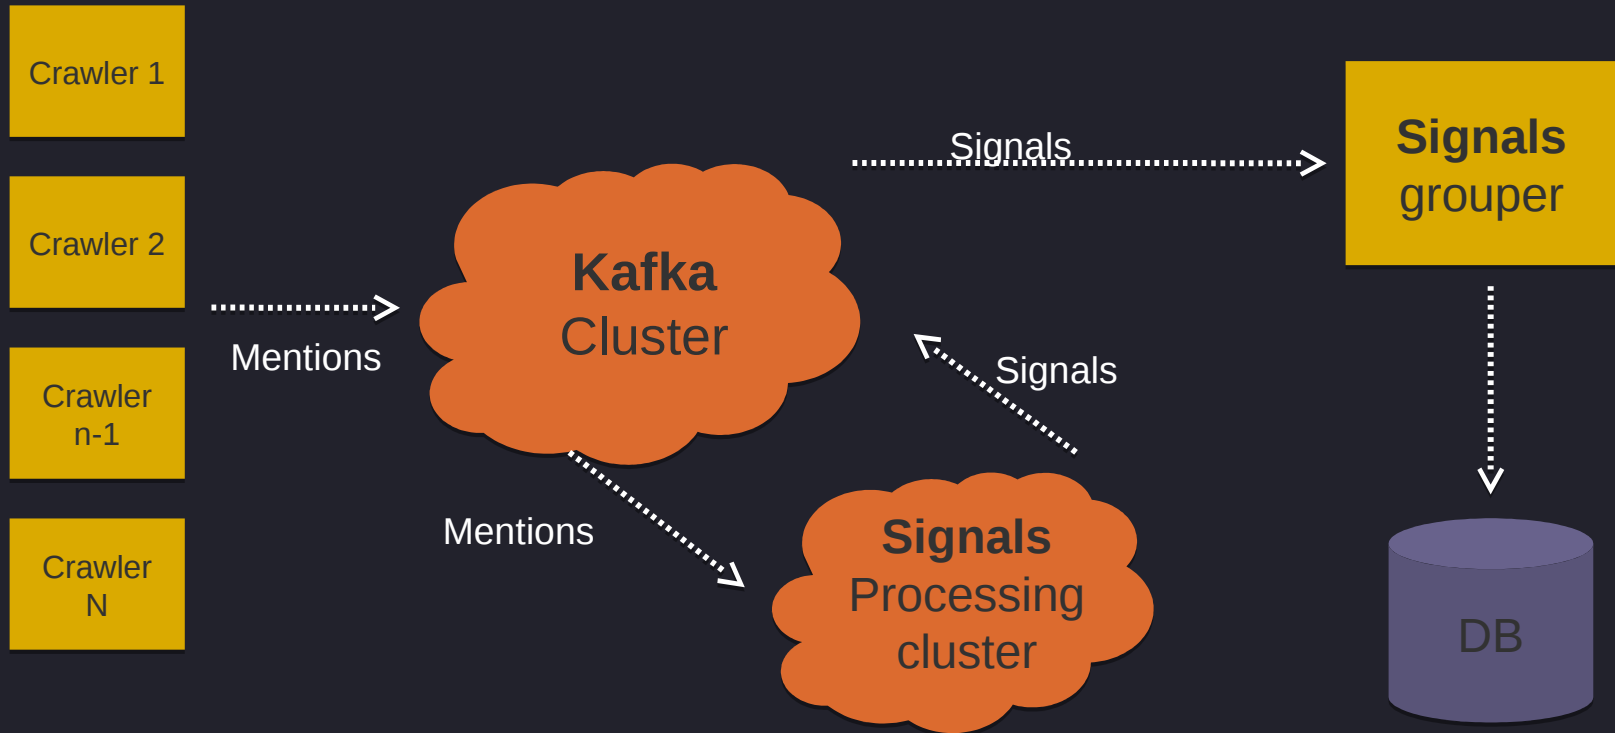




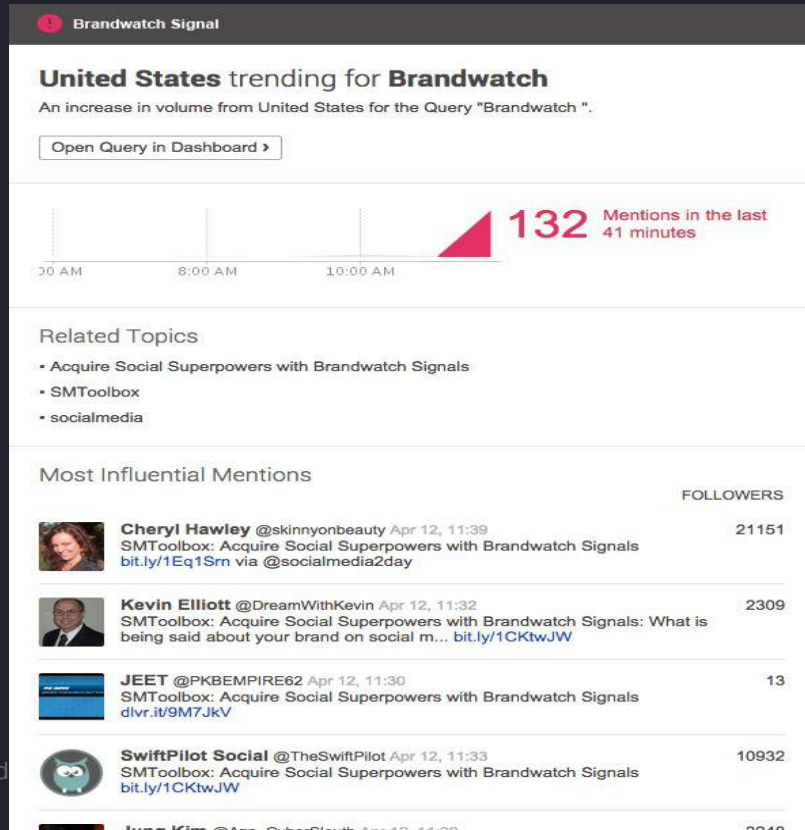
Finding  
meaning



# Step 3/ Meaning



# Meaning/ Desired outcome



# Meaning/ 1. Topics



{ “Hey @jstanier Your talk sucks! #bbuzz”,  
“I love #bbuzz”,  
“Where’s the Club Mate? #bbuzz”,  
... ,  
“I have consumed ALL THE FREE COFFEE!  
#bbuzz” }

# Meaning/ 1. Topics

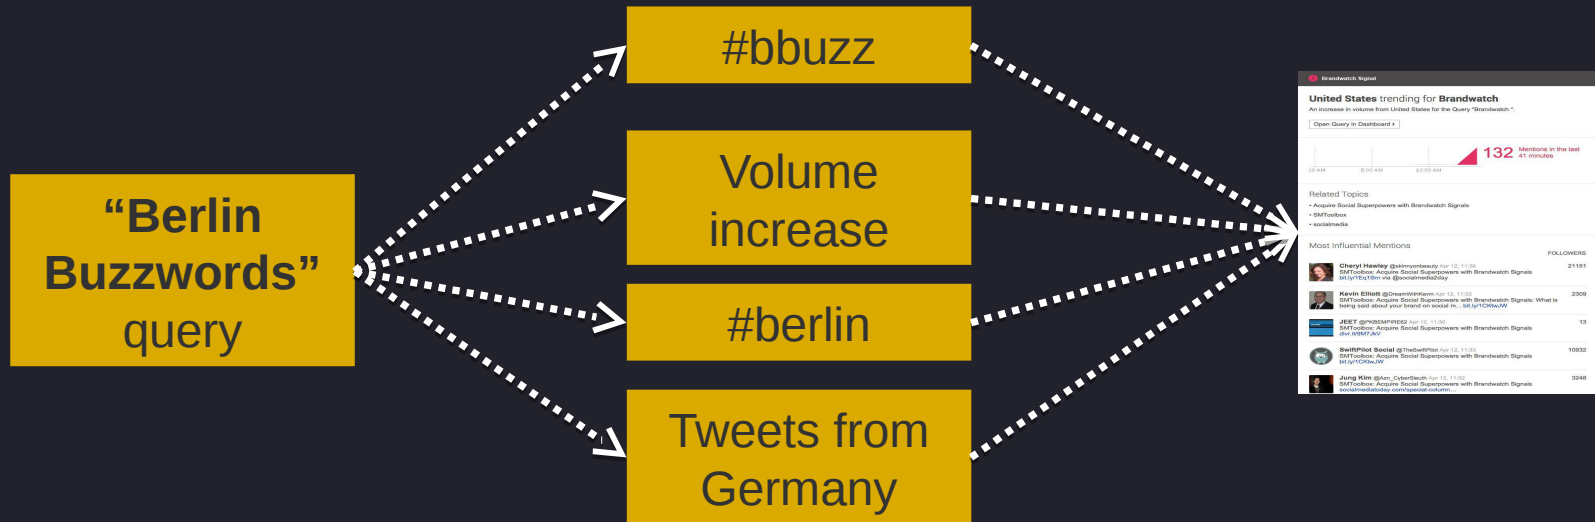


{

Berlin Buzzwords  
ElasticSearch  
Scaling

}

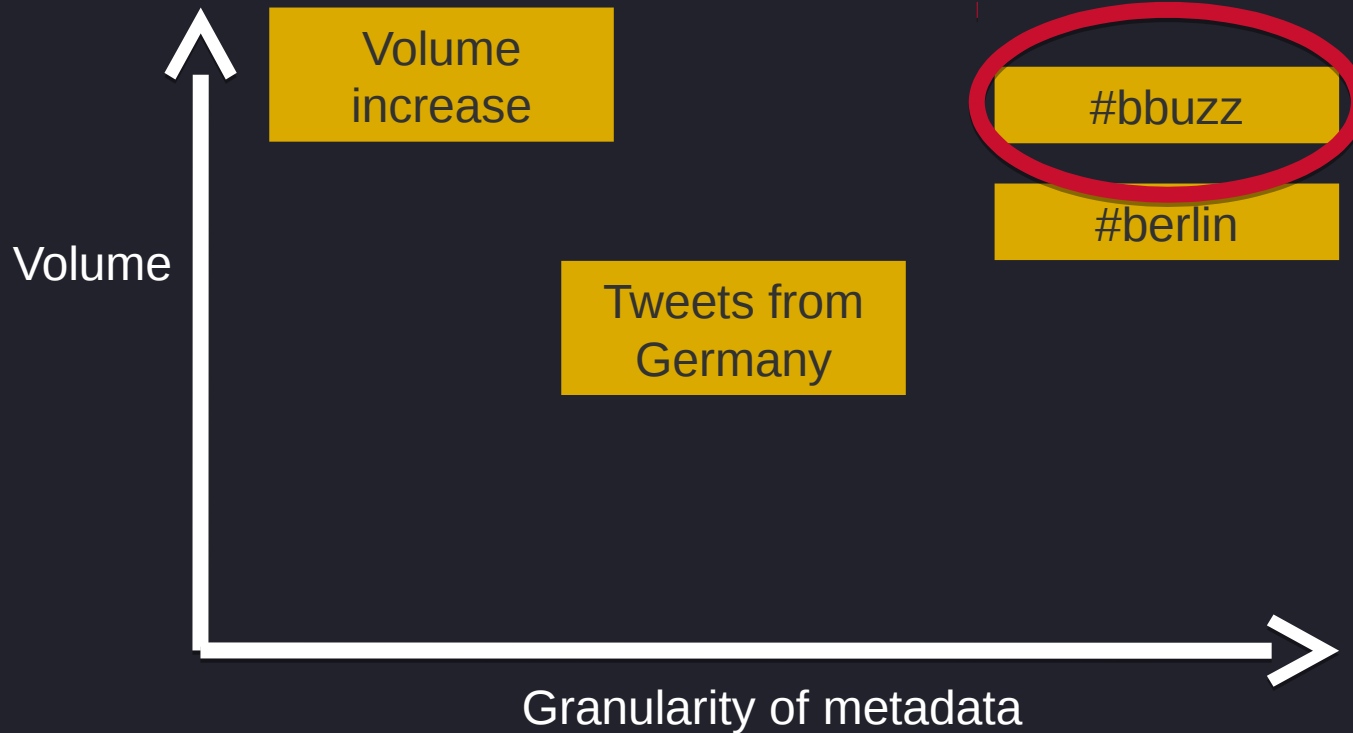
# Meaning/ 2. Grouping events




## Meaning/ 2. Grouping events

- Granularity
- Text similarity
- Shape of volume sparkline

# Meaning/ 2. Grouping events



# Meaning/ An example


 Brandwatch Signal

---

Trending **Tweet** for **Brandwatch**

We've detected a significant increase in people retweeting [the following tweet](#) for the Query "Brandwatch".

---

 **TheDrum** The Drum  
Nigel Farage bombs on social media during #leadersdebate says @brandwatch [ow.ly/LJh6c](http://ow.ly/LJh6c) [pic.twitter.com/KTCmSQOmZw](http://pic.twitter.com/KTCmSQOmZw)


---

Related Changes

An increase in the usage of the hashtag #leadersdebate

An increase in mentions linking to <http://www.thedrum.com/news/2015/04/17/nigel-farage-bombs-social-media-nationalist-leaders-pick-mostly-positive-reception> & <http://twitter.com/TheDrum/status/588980340624117760/photo/1>

---

now you know 



# Closing remarks

Say hello/

@jstanier



Q&A