



Elasticsearch.

Le moteur de recherche
élastique pour tous

Qui ?

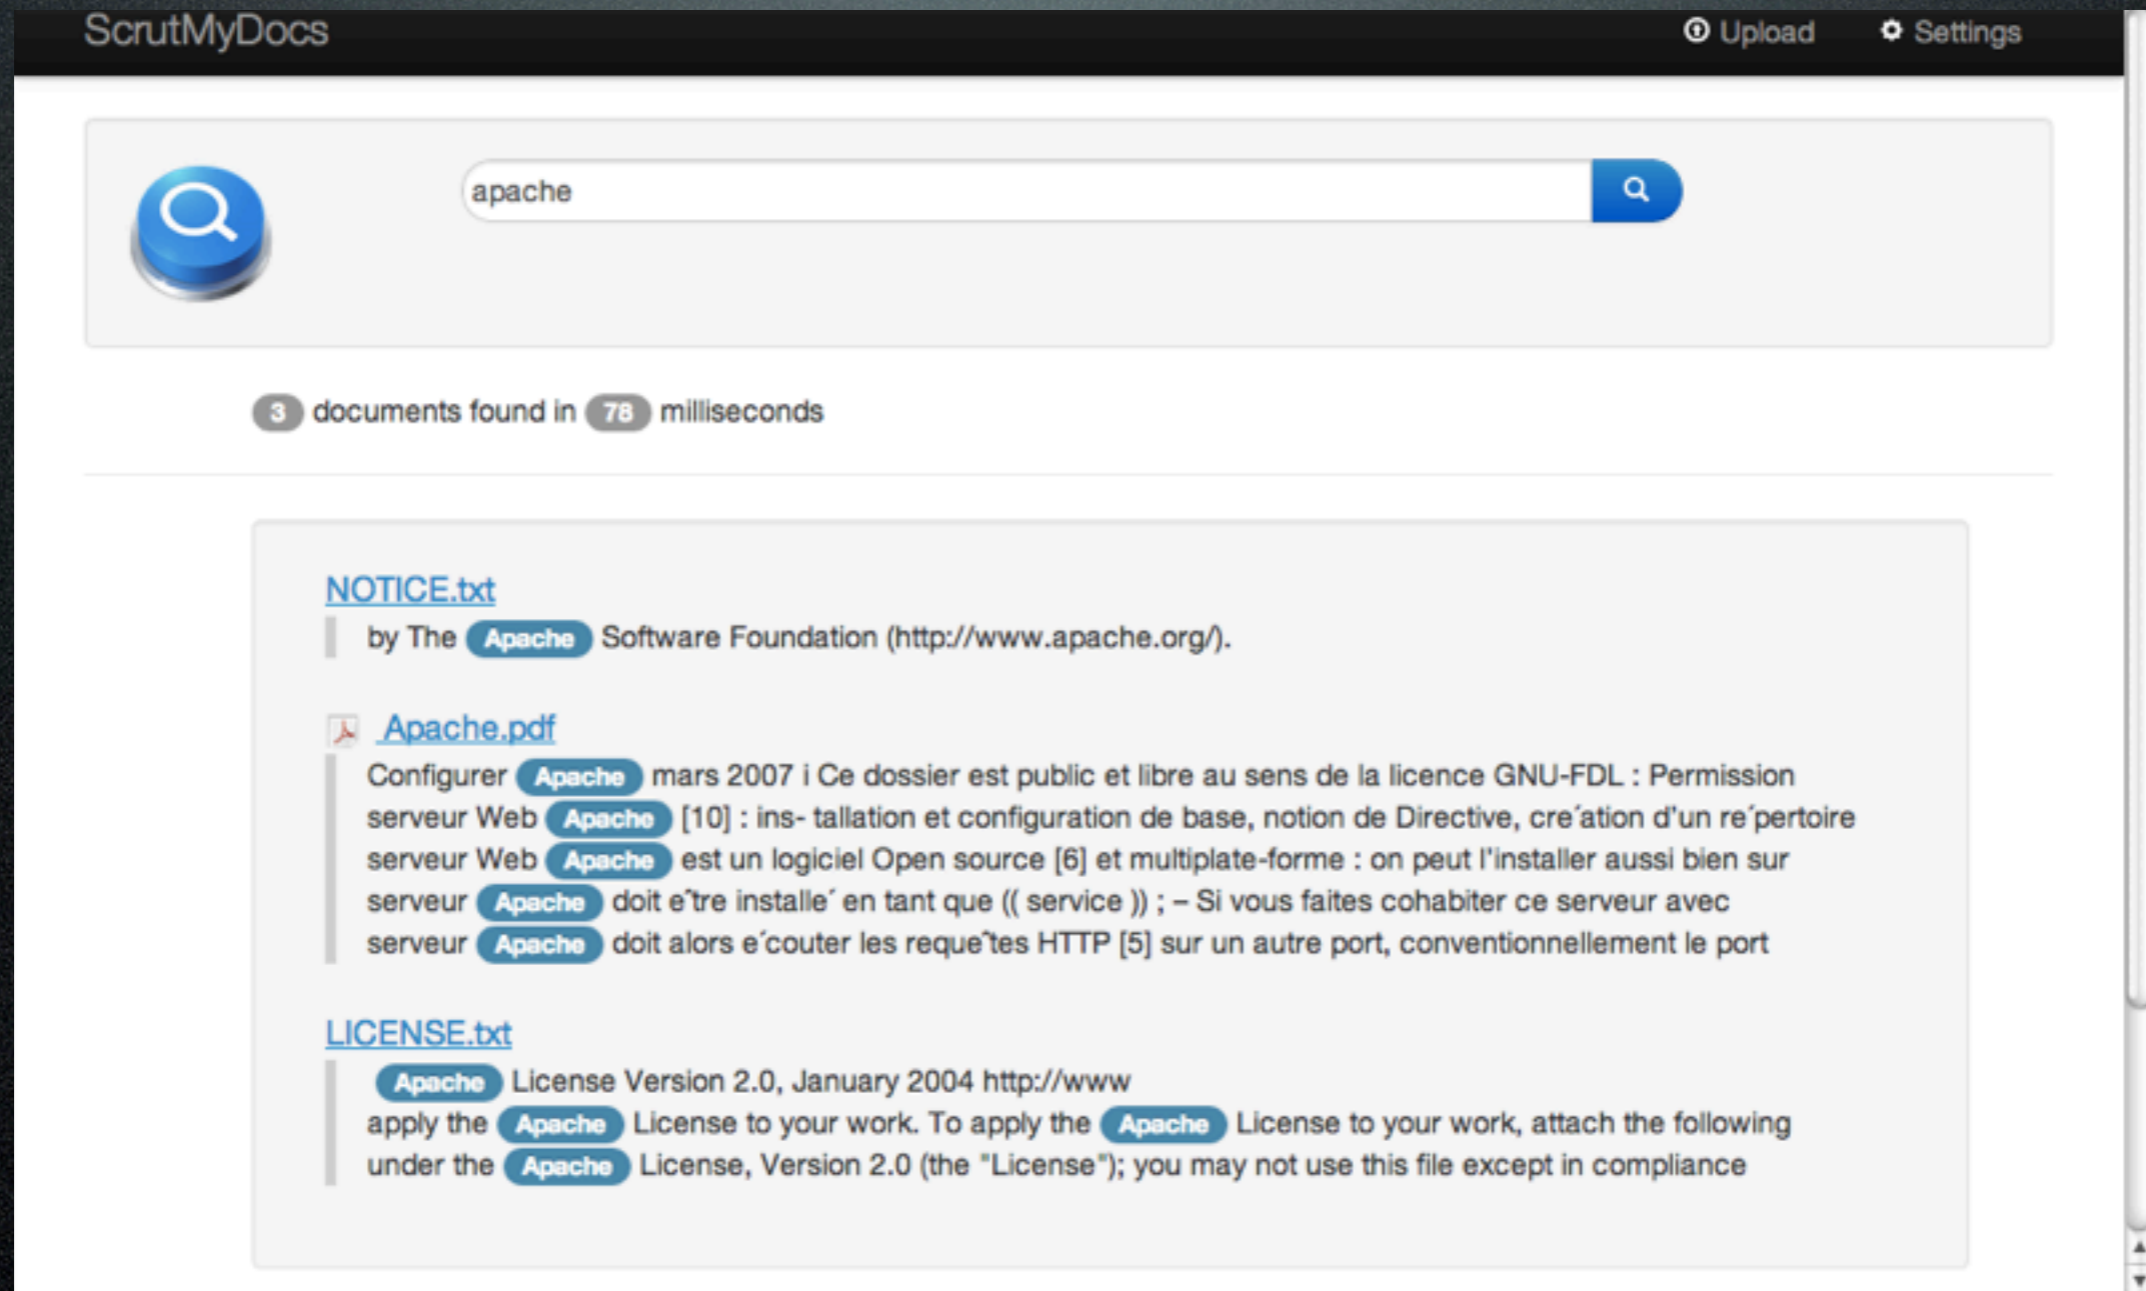
```
$ curl http://localhost:9200/talk/speaker/dpilato
{
  "nom" : "David Pilato",
  "jobs" : [
    { "boite" : "SRA Europe (SSII)", "mission" : "bon à tout faire", "date" : "1995" },
    { "boite" : "SFR", "mission" : "touche à tout", "date" : "1997" },
    { "boite" : "e-Brands / Vivendi", "mission" : "chef de projets", "date": "2000" },
    { "boite" : "DGDDI (douane)", "mission" : "mouton à 5 pattes", "date" : "2005" },
    { "boite" : "IDEO Technologies", "mission" : "directeur technique", "date" : "2012" },
    { "boite" : "Elasticsearch.com", "mission" : "technical advocate", "date" : "2013" } ],
  "passions" : [ "famille", "job", "deejay" ],
  "blog" : "http://dev.david.pilato.fr/",
  "twitter" : [ "@dadoonet", "@elasticsearchfr", "@scrutmydocs" ],
  "email" : "david@pilato.fr"
}
```

Qui ?



```
$ curl http://localhost:9200/talk/speaker/dpilato
{
  "nom" : "David Pilato",
  "jobs" : [
    { "boite" : "SRA Europe (SSII)", "mission" : "bon à tout faire", "date" : "1995" },
    { "boite" : "SFR", "mission" : "touche à tout", "date" : "1997" },
    { "boite" : "e-Brands / Vivendi", "mission" : "chef de projets", "date": "2000" },
    { "boite" : "DGDDI (douane)", "mission" : "mouton à 5 pattes", "date" : "2005" },
    { "boite" : "IDEO Technologies", "mission" : "directeur technique", "date" : "2012" },
    { "boite" : "Elasticsearch.com", "mission" : "technical advocate", "date" : "2013" } ],
  "passions" : [ "famille", "job", "deejay" ],
  "blog" : "http://dev.david.pilato.fr/",
  "twitter" : [ "@dadoonet", "@elasticsearchfr", "@scrutmydocs" ],
  "email" : "david@pilato.fr"
}
```



ScrutMyDocs.org




ScrutMyDocs Upload Settings

3 documents found in 78 milliseconds

[NOTICE.txt](#)
by The **Apache** Software Foundation (<http://www.apache.org/>).

 [Apache.pdf](#)
Configurer **Apache** mars 2007 | Ce dossier est public et libre au sens de la licence GNU-FDL : Permission serveur Web **Apache** [10] : installation et configuration de base, notion de Directive, création d'un répertoire serveur Web **Apache** est un logiciel Open source [6] et multiplate-forme : on peut l'installer aussi bien sur serveur **Apache** doit être installé en tant que ((service)) ; - Si vous faites cohabiter ce serveur avec serveur **Apache** doit alors écouter les requêtes HTTP [5] sur un autre port, conventionnellement le port

[LICENSE.txt](#)
Apache License Version 2.0, January 2004 <http://www>
apply the **Apache** License to your work. To apply the **Apache** License to your work, attach the following under the **Apache** License, Version 2.0 (the "License"); you may not use this file except in compliance

SQL Classique

Cherche moi un document
de **décembre 2011** portant sur la **france**
et contenant **produit** et **david**

En SQL :



SQL Classique

Cherche moi un document
de **décembre 2011** portant sur la **france**
et contenant **produit** et **david**










En SQL :

```
SELECT
  doc.*, pays.*
FROM
  doc, pays
WHERE
  doc.pays_code = pays.code AND
  doc.date_doc > to_date('2011-12', 'yyyy-mm') AND
  doc.date_doc < to_date('2012-01', 'yyyy-mm') AND
  lower(pays.libelle) = 'france' AND
  lower(doc.commentaire) LIKE '%produit%' AND
  lower(doc.commentaire) LIKE '%david%';
```



Au final, on obtient

Power Search:

ID Number	<input type="text"/>
Web Title	<input type="text"/>
Url	<input type="text"/>
Category	Select <input type="button" value="v"/>
Web Description	<input type="text"/>
Keywords	<input type="text"/>
Contact Name	<input type="text"/>
Contact Email	<input type="text"/>
Featured Links 	Select <input type="button" value="v"/>
Cool Links 	Select <input type="button" value="v"/>
Bold Links	Select <input type="button" value="v"/>
Icon	<input type="radio"/>  <input type="radio"/>  <input type="radio"/>  <input type="radio"/>  <input type="radio"/>  <input type="radio"/> 
Rating Average 	Select <input type="button" value="v"/>
Number of Votes	between <input type="text"/> and <input type="text"/>
Total Hits	between <input type="text"/> and <input type="text"/>
Hits Today	between <input type="text"/> and <input type="text"/>
IP Address	<input type="text"/>
Submission Software Name	<input type="text"/>




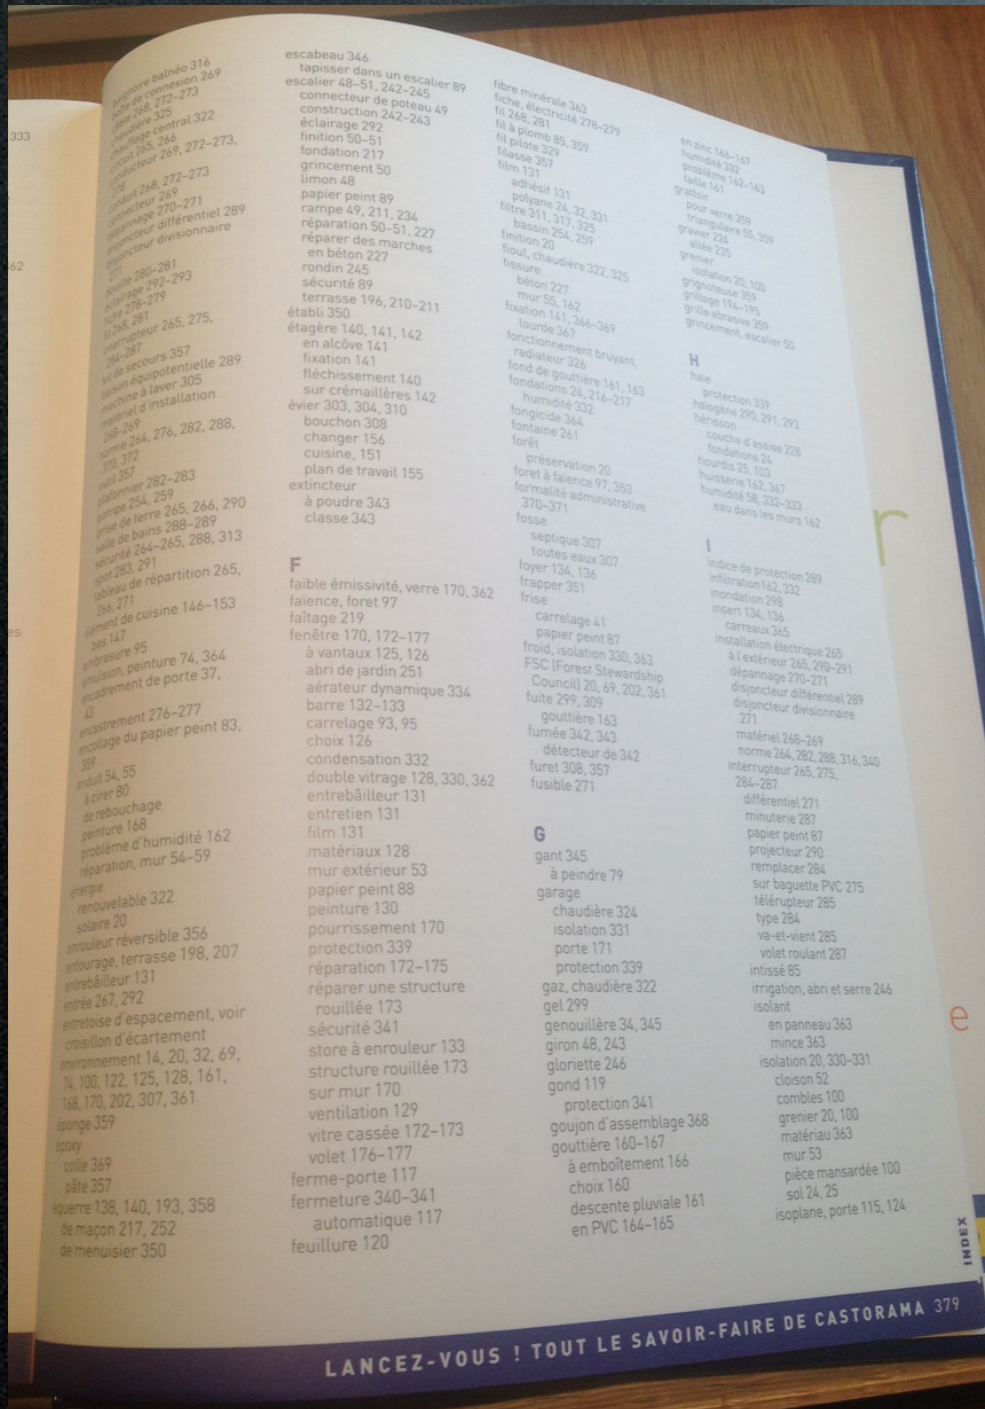
Un moteur de recherche

- Un moteur de recherche est composé de :
 - un moteur d'**indexation de documents**
 - un moteur de **recherche sur les index**
- De fait, un moteur de recherche est énormément plus rapide qu'une base de données pour faire des recherches :

c'est son métier !



L'indexation, c'est quoi en fait ?



Elasticsearch





Elasticsearch

- Moteur de recherche pour la génération NoSQL
- Basé sur le standard Apache Lucene
- Masque la complexité Java/Lucene à l'aide de services standards HTTP / RESTful / JSON
- Utilisable à partir de n'importe quelle technologie
- Ajoute la couche cloud manquante à Lucene
- C'est un **moteur**, pas une **interface graphique** !



Indexer

```
$ curl -XPUT localhost:9200/twitter/tweet/1 -d '{
  "text": "Bienvenue à la conférence #elasticsearch pour #JUG",
  "created_at": "2012-04-06T20:45:36.000Z",
  "source": "Twitter for iPad",
  "truncated": false,
  "retweet_count": 0,
  "hashtag": [ { "text": "elasticsearch", "start": 27, "end": 40 },
               { "text": "JUG", "start": 47, "end": 55 } ],
  "user": { "id": 51172224, "name": "David Pilato",
            "screen_name": "dadoonet", "location": "France",
            "description": "Soft Architect, Project Manager, Senior Developer.\r\nAt this time, enjoying NoSQL world : CouchDB, ElasticSearch.\r\nDeeJay 4 times a year, just for fun !" }
}'
```



Indexer

```
$ curl -XPUT localhost:9200/twitter/tweet/1 -d '{
  "text": "Bienvenue à la conférence #elasticsearch pour #JUG",
  "created_at": "2012-04-06T20:45:36.000Z",
  "source": "Twitter for iPad",
  "truncated": false,
  "retweet_count": 0,
  "hashtag": [ { "text": "elasticsearch", "start": 27, "end": 40 },
               { "text": "JUG", "start": 47, "end": 55 } ],
  "user": { "id": 51172224, "name": "David Pilato",
            "screen_name": "dadoonet", "location": "France",
            "description": "Soft Architect, Project Manager, Senior Developer.\r\n\r\nAt this time, enjoying NoSQL world : CouchDB, ElasticSearch.\r\n\r\nDeeJay 4 times a year, just for fun !" }
}'
```



```
{
  "ok":true,
  "_index":"twitter",
  "_type":"tweet",
  "_id":"1"
}
```



Chercher

```
$ curl localhost:9200/twitter/tweet/_search?q=elasticsearch
```



Chercher

```
$ curl localhost:9200/twitter/tweet/_search?q=elasticsearch
```



```
{
  "took" : 24,
  "timed_out" : false,
  "_shards" : { "total" : 5, "successful" : 5, "failed" : 0 },
  "hits" : {
    "total" : 1,
    "max_score" : 0.227,
    "hits" : [ {
      "_index" : "twitter",
      "_type" : "tweet",
      "_id" : "1",
      "_score" : 0.227, "_source" : {
        "text": "Bienvenue à la conférence #elasticsearch pour #JUG",
        "created_at": "2012-04-06T20:45:36.000Z",
        "source": "Twitter for iPad",
        [...]
      }
    } ]
  }
}
```



Chercher

```
$ curl localhost:9200/twitter/tweet/_search?q=elasticsearch
```



```
{
  "took" : 24,
  "timed_out" : false,
  "_shards" : { "total" : 5, "successful" : 5, "failed" : 0 },
  "hits" : {
    "total" : 1,
    "max_score" : 0.227,
    "hits" : [ {
      "_index" : "twitter",
      "_type" : "tweet",
      "_id" : "1",
      "_score" : 0.227, "_source" : {
        "text": "Bienvenue à la conférence #elasticsearch pour #JUG",
        "created_at": "2012-04-06T20:45:36.000Z",
        "source": "Twitter for iPad",
        [...]
      }
    } ]
  }
}
```

Nb de documents



Chercher

```
$ curl localhost:9200/twitter/tweet/_search?q=elasticsearch
```



```
{
  "took" : 24,
  "timed_out" : false,
  "_shards" : { "total" : 5, "successful" : 5, "failed" : 0 },
  "hits" : {
    "total" : 1,
    "max_score" : 0.227,
    "hits" : [ {
      "_index" : "twitter",
      "_type" : "tweet",
      "_id" : "1",
      "_score" : 0.227, "_source" : {
        "text": "Bienvenue à la conférence #elasticsearch pour #JUG",
        "created_at": "2012-04-06T20:45:36.000Z",
        "source": "Twitter for iPad",
        [...]
      }
    } ]
  }
}
```

Coordonnées



Chercher

```
$ curl localhost:9200/twitter/tweet/_search?q=elasticsearch
```



```
{
  "took" : 24,
  "timed_out" : false,
  "_shards" : { "total" : 5, "successful" : 5, "failed" : 0 },
  "hits" : {
    "total" : 1,
    "max_score" : 0.227,
    "hits" : [ {
      "_index" : "twitter",
      "_type" : "tweet",
      "_id" : "1",
      "_score" : 0.227, "_source" : {
        "text": "Bienvenue à la conférence #elasticsearch pour #JUG",
        "created_at": "2012-04-06T20:45:36.000Z",
        "source": "Twitter for iPad",
        [...]
      }
    } ]
  }
}
```

Pertinence



Chercher

```
$ curl localhost:9200/twitter/tweet/_search?q=elasticsearch
```



```
{
  "took" : 24,
  "timed_out" : false,
  "_shards" : { "total" : 5, "successful" : 5, "failed" : 0 },
  "hits" : {
    "total" : 1,
    "max_score" : 0.227,
    "hits" : [ {
      "_index" : "twitter",
      "_type" : "tweet",
      "_id" : "1",
      "_score" : 0.227, "_source" : {
        "text": "Bienvenue à la conférence #elasticsearch pour #JUG",
        "created_at": "2012-04-06T20:45:36.000Z",
        "source": "Twitter for iPad",
        [...]
      }
    } ]
  }
}
```

Document
source



Query DSL

- Requêtes précises : plutôt que de chercher « à la google », on peut utiliser des critères précis :



Query DSL

- Requêtes précises : plutôt que de chercher « à la google », on peut utiliser des critères précis :

```
$ curl -XPOST localhost:9200/twitter/tweet/_search -d '{
  "bool" : {
    "must" : {
      "term" : { "user" : "kimchy" }
    },
    "must_not" : {
      "range" : {
        "age" : { "from" : 10, "to" : 20 }
      }
    },
    "should" : [
      {
        "term" : { "tag" : "wow" }
      }, {
        "match" : { "tag" : "elasticsearch is cool" }
      }
    ]
  }
}'
```



Analyser

La puissance des facettes !

Faites parler vos données en les regardant sous différentes facettes !
(Et en temps quasi réel, s'il vous plait !)



Des tweets

ID	Username	Date	Hashtag
1	dadoonet	2012-04-18	1
2	talk	2012-04-18	5
3	elasticsearch	2012-04-18	2
4	dadoonet	2012-04-18	2
5	talk	2012-04-18	6
6	elasticsearch	2012-04-19	3
7	dadoonet	2012-04-19	3
8	talk	2012-04-19	7
9	elasticsearch	2012-04-20	4



Term Facet


Username	Date	Hashtag
dadoonet	2012-04-18	1
talktalk	2012-04-18	5
elasticsearch	2012-04-18	2
dadoonet	2012-04-18	2
talk	2012-04-18	6
elasticsearch	2012-04-19	3
dadoonet	2012-04-19	3
talk	2012-04-19	7
elasticsearch	2012-04-20	4



Term Facet

Username	Date	Hashtag
dadoonet	2012-04-18	1
talktalk	2012-04-18	5
elasticsearch	2012-04-18	3
dadoonet	2012-04-19	3
talk	2012-04-19	7
elasticsearch	2012-04-20	4
dadoonet	2012-04-19	3
talk	2012-04-19	3
elasticsearch	2012-04-19	3

Username	Count
dadoonet	3
talk	3
elasticsearch	3



Term Facet

Username	Date	Hashtag
dadoonet	2012-04-18	1
talk	2012-04-18	5
elasticsearch	2012-04-18	2
dadoonet	2012-04-18	2
talk	2012-04-18	6
elasticsearch	2012-04-19	3
dadoonet	2012-04-19	3
talk	2012-04-19	7
elasticsearch	2012-04-20	4



Term Facet


Username	Date	Hashtag
dadoonet	2012-04-18	2
talk	2012-04-18	6
elasticsearch	2012-04-18	2
dadoonet	2012-04-18	3
talk	2012-04-19	7
elasticsearch	2012-04-19	3
dadoonet	2012-04-19	3
talk	2012-04-19	7
elasticsearch	2012-04-20	4

```
"facets" : {  
  "users" : { "terms" : { "field" : "username" } }  
}
```



Term Facet

Username	Date	Hashtag
dadoonet	2012-04-20	
talk	2012-04-20	
elasticsearch	2012-04-20	
dadoonet	2012-04-20	
talk	2012-04-20	
elasticsearch	2012-04-20	
dadoonet	2012-04-20	
talk	2012-04-20	
elasticsearch	2012-04-20	4



```
"facets" : {  
  "users" : { "terms" : { "field" : "username" } }  
}  
  
"facets" : {  
  "users" : {  
    "_type" : "terms",  
    "missing" : 0,  
    "total": 9,  
    "other": 0,  
    "terms" : [  
      { "term" : "dadoonet", "count" : 3 },  
      { "term" : "talk", "count" : 3 },  
      { "term" : "elasticsearch", "count" : 3 }  
    ]  
  }  
}
```



Date Histogram Facet

Date	Hashtag
2012-04-18	1
2012-04-18	5
2012-04-18	2
2012-04-18	2
2012-04-18	6
2012-04-19	3
2012-04-19	3
2012-04-19	7
2012-04-20	4



Date Histogram Facet

Date	Hashtag
2012-04-18	1
2012-04-18	5
2012-04-18	2
2012-04-18	2
2012-04-18	6
2012-04-19	3
2012-04-19	3
2012-04-19	7
2012-04-20	4


Par mois

Date	Count
2012-04	9




Date Histogram Facet

Date	Hashtag	Par mois	
Date		Date	Count
2012-04-18	1	2012-04	9
2012-04-18	5		
2012-04-18	2		
2012-04-18	2		
2012-04-18	6		
2012-04-19	3	Date	Count
2012-04-19	3	2012-04-18	5
2012-04-19	7	2012-04-19	3
2012-04-20	4	2012-04-20	1



Date Histogram Facet


Date	Hashtag
2012-04-18	1
2012-04-18	5
2012-04-18	2
2012-04-18	2
2012-04-18	6
2012-04-19	3
2012-04-19	3
2012-04-19	7
2012-04-20	4



Date Histogram Facet

Date	Hashtag
2012-04-18	1
2012-04-18	5
2012-04-18	2
2012-04-18	2
2012-04-18	6
2012-04-19	3
2012-04-19	3
2012-04-19	7
2012-04-20	4

```
"facets" : {  
  "perday" : {  
    "date_histogram" : {  
      "field" : "date",  
      "interval" : "day"  
    }  
  }  
}
```



Date Histogram Facet

Date	Hashtag
2012-04-18	1
2012-04-18	5
2012-04-18	2
2012-04-18	2
2012-04-18	6
2012-04-19	3
2012-04-19	3
2012-04-19	7
2012-04-20	4



```
"facets" : {  
  "perday" : {  
    "date_histogram" : {  
      "field" : "date",  
      "interval" : "day"  
    }  
  }  
}
```

```
"facets" : {  
  "perday" : {  
    "_type" : "date_histogram",  
    "entries" : [  
      { "time" : 1334700000000, "count" : 5 },  
      { "time" : 1334786400000, "count" : 3 },  
      { "time" : 1334872800000, "count" : 1 }  
    ]  
  }  
}
```



Range Facet

	Hashtag
	1
	5
	2
	2
	6
	3
	3
	7
	4



Range Facet

Hashtag
1
5
2
2
6
3
3
7
4



Hashtag	Count	Min	Max	Moy	Total
$x < 3$	3	1	2	1.667	5
$3 \leq x < 5$	3	3	4	3.333	10
$x \geq 5$	3	5	7	6	18



Range Facet

	Hashtag
	1
	5
	2
	2
	6
	3
	3
	7
	4



Range Facet

```
"facets" : { "hashtags" : {  
  "range" : { "field" : "hashtag",  
  "ranges" : [  
    { "to" : 3 },  
    { "from" : 3, "to" : 5 },  
    { "from" : 5 }  
  ] } } }
```

Hashtag
1
5
2
2
6
3
3
7
4



Range Facet

```
"facets" : { "hashtags" : {  
  "range" : { "field" : "hashtag",  
  "ranges" : [  
    { "to" : 3 },  
    { "from" : 3, "to" : 5 },  
    { "from" : 5 }  
  ] } } }
```

```
"facets" : {  
  "hashtags" : {  
    "_type" : "range",  
    "ranges" : [ {  
      "to": 3,  
      "count": 3,  
      "min": 1, "max": 2,  
      "total": 5, "mean": 1.667  
    }, {  
      "from": 3, "to" : 5,  
      "count": 3,  
      "min": 3, "max": 4,  
      "total": 10, "mean": 3.333  
    }, {  
      "from": 5,  
      "count": 3,  
      "min": 5, "max": 7,  
      "total": 18, "mean": 6  
    } ] } }
```

Hashtag
1
5
2
2
6
3
3
7
4



Site marchand

The screenshot displays the Fotolia website interface. At the top, the Fotolia logo is accompanied by the text "16 525 874 images haute définition pour tous vos besoins professionnels" and the phone number "0970 465 265". Navigation links for "Connexion", "Créer un compte", and "Mon Panier (0)" are visible. Below the header, a search bar contains the word "recherche" and a "GO!" button. The main content area shows a grid of image thumbnails, with one image (#36840357) highlighted in a larger preview window. The sidebar on the left offers advanced search filters under "Recherche avancée", including "Ancienneté" (age), "Collection", "Type", and "Orientation". The bottom of the page shows licensing information: "LICENCES XS - XXL | V | XV" and "CRÉDITS 1" with a price of "2403" and "533".



Site marchand

The screenshot shows the Fotolia website interface. At the top, the Fotolia logo is displayed with the text "16 525 874 images haute définition pour tous vos besoins professionnels" and the phone number "0970 465 265". Navigation links for "Connexion", "Créer un compte", and "Mon Panier" are visible. A search bar contains the word "recherche" and is highlighted with a blue box. Below the search bar, the results are shown as "65 900 résultats" with a "Pertinence" sorting option. A grid of image thumbnails is displayed, each with a unique ID and a "Ajouter au Panier" button. A detailed view of a specific image (ID #36840357) is shown in the foreground, featuring a search interface with various filters and options. The bottom of the page shows licensing information: "LICENCES XS - XXL | V | XV" and "CRÉDITS 1" with a price of "2403" and "533".



Site marchand

The screenshot displays the Fotolia website interface. At the top, the Fotolia logo is accompanied by the text "16 525 874 images haute définition pour tous vos besoins professionnels" and the phone number "0970 465 265". Navigation links for "Connexion", "Créer un compte", and "Mon Panier" are visible. Below the header, there are tabs for "Acheter des images", "Images", "Vecteurs", "Vidéos", and "Logos". A search bar contains the word "recherche" and a "GO!" button. The main content area shows "65 900 résultats" and a "Pertinence" sorting option. A grid of image thumbnails is displayed, with one image (#36840357) highlighted in a larger preview window. The preview window shows a collection of colorful, modern web interface elements. The grid includes various images such as a magnifying glass, a hand pointing at a screen, a maze, a 3D figure with a magnifying glass, a camera lens, and a globe with "WWW".



Site marchand

The screenshot displays the Fotolia website interface. At the top, the Fotolia logo is accompanied by the text "16 525 874 images haute définition pour tous vos besoins professionnels" and the phone number "0970 465 265". Navigation links for "Connexion", "Créer un compte", and "Mon Panier (0)" are visible. Below the header, there are tabs for "Acheter des images", "Images", "Vecteurs", "Vidéos", and "Logos". A search bar contains the word "recherche" and a "GO!" button. The main content area shows "65 900 résultats" with a "Pertinence" sorting option. A left sidebar contains filters for "Ancienneté" (ranging from "≤ 1 semaine" to "≤ 2 ans"), "Collection" (Standard, Infinite, Abonnement, Toutes les collections), "Type" (Photos, Illustrations, Vecteurs, Vidéos, Tous types), and "Orientation" (Paysage, Portrait, Carré). The main grid features a large preview of a web application interface with the ID "#36840357" and a "LICENCES: XS - XXL | V | XV" label. Other thumbnails include a magnifying glass, a hand pointing at a screen, a maze, a 3D figure with a magnifying glass, a camera lens, and a globe with "WWW".



Site marchand

The image shows a search filter sidebar on a marketplace website. The sidebar is divided into several sections, each with a dropdown arrow and a title. The sections are: Ancienneté, Collection, Type, Orientation, Options, and Tailles. Each section contains a list of filter options with radio buttons or checkboxes. Annotations with red arrows point to specific filter options: 'Ranges' points to the '≤ 2 ans (~ 38 600)' option in the 'Ancienneté' section; 'Term' points to the 'Photos (~ 30 000)' option in the 'Type' section; 'Term' points to the 'Paysage (~ 35 400)' option in the 'Orientation' section; and 'Ranges' points to the 'L (~ 64 700)' option in the 'Tailles' section. The background shows a grid of search results with various images and product listings.

Ancienneté

- Tout
- ≤ 2 ans (~ 38 600)
- ≤ 1 an (~ 21 000)
- ≤ 6 mois (~ 10 100)
- ≤ 1 mois (~ 1 900)
- ≤ 1 semaine (~ 140)

Collection

- Collection Standard
- Collection Infinie
- Collection abonnement
- Toutes les collections

Type

- Photos (~ 30 000)
- Illustrations (~ 18 000)
- Vecteurs (~ 16 800)
- Vidéos (~ 950)
- Tous types

Orientation

- Paysage (~ 35 400)
- Portrait (~ 13 500)
- Carré (~ 16 800)
- Tous

Options

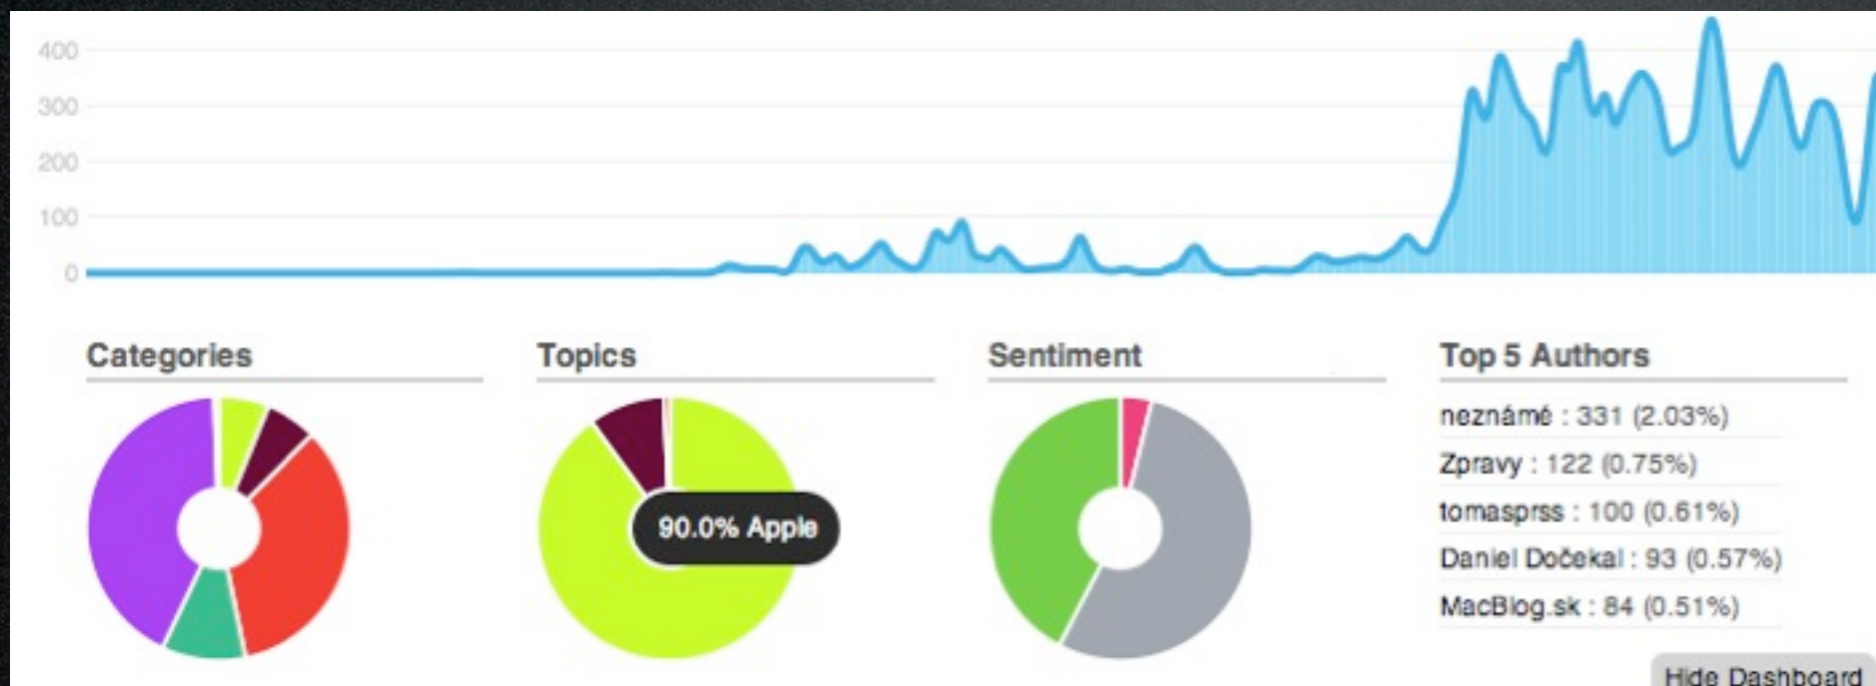
Tailles

- L (~ 64 700)
- XL (~ 56 300)
- XXL (~ 34 800)
- Y (~ 50 000)



Analyse temps-réel

- Faire un matchAll sur l'ensemble des données
- Actualiser toutes les x secondes
- Indexer en même temps les nouvelles données

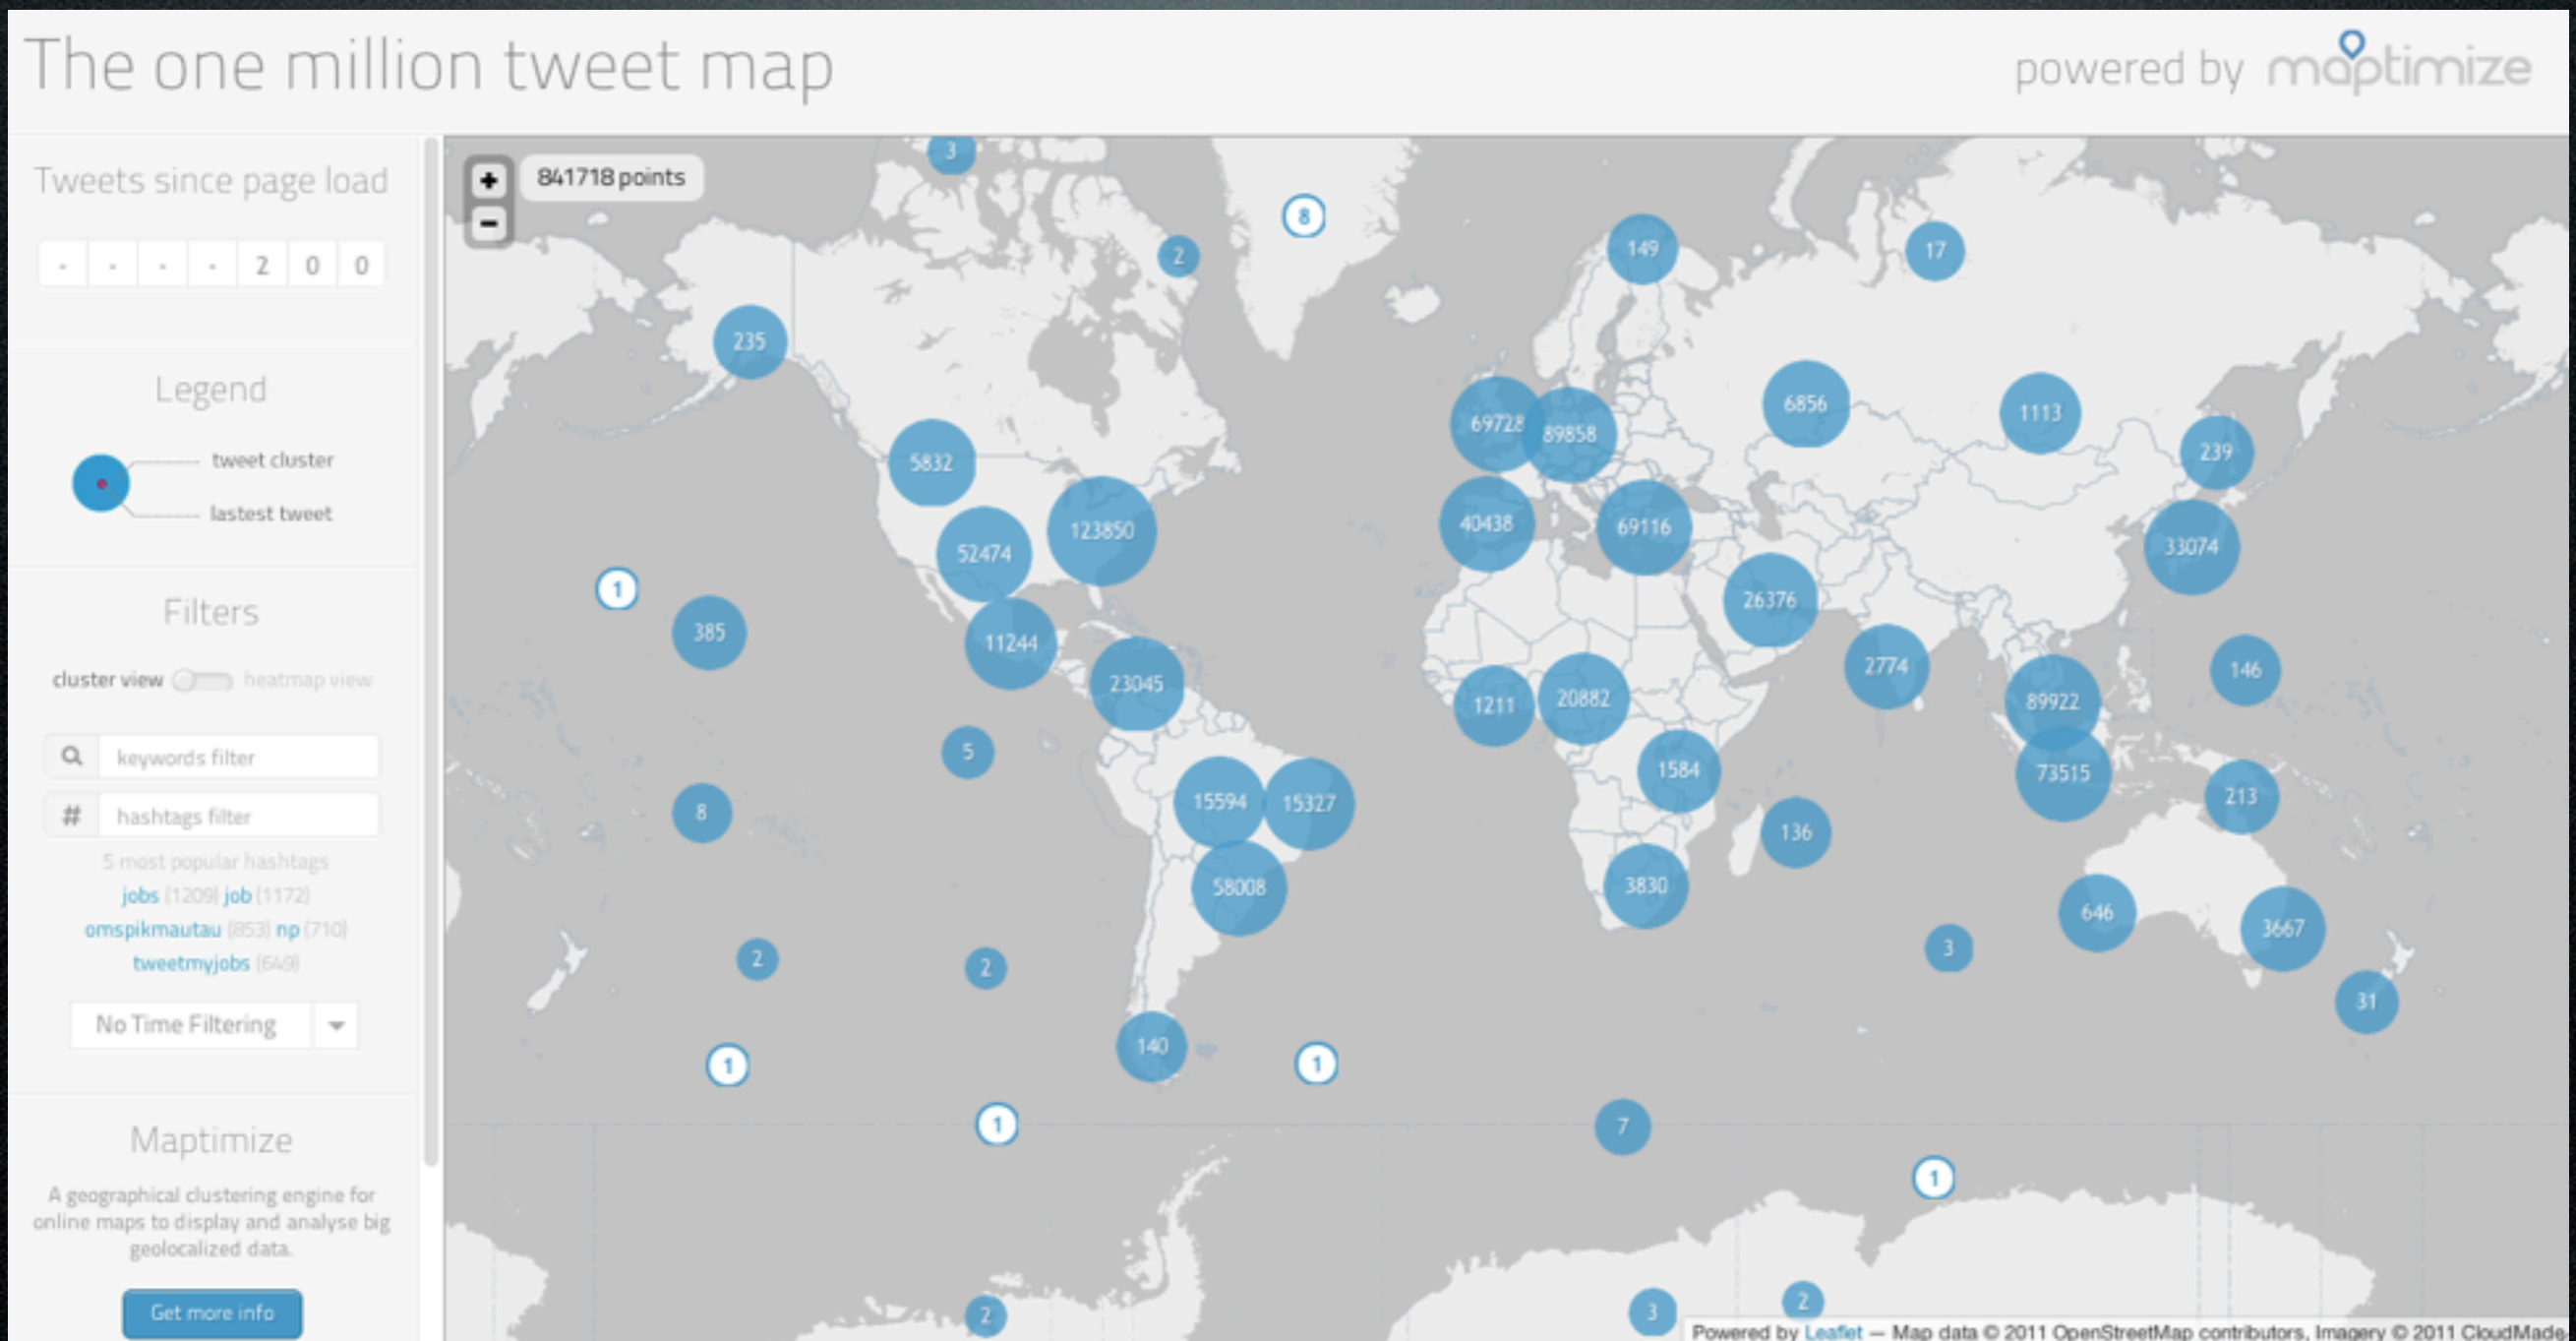


← **Date histogram**

← **Term**











Facettes Cartographiques



Reprenons notre formulaire

Power Search:

ID Number	<input type="text"/>
Web Title	<input type="text"/>
Url	<input type="text"/>
Category	Select <input type="button" value="v"/>
Web Description	<input type="text"/>
Keywords	<input type="text"/>
Contact Name	<input type="text"/>
Contact Email	<input type="text"/>
Featured Links 	Select <input type="button" value="v"/>
Cool Links 	Select <input type="button" value="v"/>
Bold Links	Select <input type="button" value="v"/>
Icon	<input type="radio"/>  <input type="radio"/>  <input type="radio"/>  <input type="radio"/>  <input type="radio"/>  <input type="radio"/> 
Rating Average 	Select <input type="button" value="v"/>
Number of Votes	between <input type="text"/> and <input type="text"/>
Total Hits	between <input type="text"/> and <input type="text"/>
Hits Today	between <input type="text"/> and <input type="text"/>
IP Address	<input type="text"/>
Submission Software Name	<input type="text"/>



Reprenons notre formulaire

Power Search:

ID Number

Web Title

Url

Category

Web Description

Keywords




Contact Name




Contact Email

Featured Links

Cool Links

Bold Links

Icon   

Rating Average

Number of Votes between and

Total Hits between and

Hits Today between and

IP Address

Submission Software Name



Recherche Full Text



Reprenons notre formulaire

Power Search:

ID Number

Web Title

Url

Category

Web Description

Keywords







Contact Name

Contact Email

Featured Links

Cool Links

Bold Links

Icon   
   

Rating Average

Number of Votes between and

Total Hits between and

Hits Today between and

IP Address

Submission Software Name



[ElasticSearch - Big Data Search and Analytics Made Easy](#)
[www.elasticsearch.com/](#) - Traduire cette page
We're the company behind **elasticsearch**, the open source distributed search engine. We develop, support and provide professional services. Discover and ...

[Améliorer les performances des applications web avec Elasticsearch](#)
[www.clever-age.com](#) > Veille > Notre blog de veille
3 juil. 2012 – Le modèle de données relationnel de certaines applications web est parfois si complexe que les performances se dégradent rapidement en ...

[ElasticSearch - Wikipédia](#)
[fr.wikipedia.org/wiki/ElasticSearch](#)
ElasticSearch est un moteur de recherche libre (open source) basé sur Apache Lucene. Il est : distribué (architecture cloud); utilise une base de données noSQL ...

[ElasticSearch - Wikipedia, the free encyclopedia](#)
[en.wikipedia.org/wiki/ElasticSearch](#) - Traduire cette page
ElasticSearch is a distributed, RESTful, free/open source search server based on Apache Lucene. It is developed by Shay Banon and is released under the ...

[ElasticSearch : moteur de recherche taillé pour le Cloud](#)
[soat.developpez.com/articles/elasticsearch/](#)
4 mai 2012 – Cette rétrospective de la présentation de David Pilato lors de Devcoxx France 2012 dont le titre était : « **ElasticSearch** : Moteur de recherche ...

[ElasticSearch - Twitter](#)
[twitter.com/search?q=%23ElasticSearch](#) - Traduire cette page
La description de ce résultat n'est pas accessible à cause du fichier robots.txt de ce site. En savoir plus



Reprenons notre formulaire

Power Search:

ID Number

Web Title

Url

Category

Web Description

Keywords

Contact Name

Contact Email

Featured Links

Cool Links

Bold Links

Icon

Rating Average




Number of Votes and

Total Hits and

Hits Today and

IP Address

Submission Software Name



ElasticSearch - Big Data Search and Analytics Made Easy
www.elasticsearch.com/ - Traduire cette page
We're the company behind **elasticsearch**, the open source distributed search engine. We develop, support and provide professional services. Discover and ...

Améliorer les performances des applications web avec Elasticsearch
www.clever-age.com > Veille > Notre blog de veille
3 jul. 2012 – Le modèle de données relationnel de certaines applications web est parfois si complexe que les performances se dégradent rapidement en ...

ElasticSearch - Wikipédia
fr.wikipedia.org/wiki/ElasticSearch
ElasticSearch est un moteur de recherche libre (open source) basé sur Apache Lucene. Il est : distribué (architecture cloud); utilise une base de données noSQL ...

ElasticSearch - Wikipedia, the free encyclopedia
en.wikipedia.org/wiki/ElasticSearch - Traduire cette page
ElasticSearch is a distributed, RESTful, free/open source search server based on Apache Lucene. It is developed by Shay Banon and is released under the ...

ElasticSearch : moteur de recherche taillé pour le Cloud
soat.developpez.com/articles/elasticsearch/
4 mai 2012 – Cette rétrospective de la présentation de David Pilato lors de Devcoxx France 2012 dont le titre était : « **ElasticSearch** : Moteur de recherche ...

ElasticSearch - Twitter
twitter.com/search?q=%23ElasticSearch - Traduire cette page
La description de ce résultat n'est pas accessible à cause du fichier robots.txt de ce site. En savoir plus



Démonstration

<http://onemilliontweetmap.com/>

<http://www.francetvinfo.fr/en-direct/>

http://localhost:9200/_plugin/demo/



Architecture

Un peu plus de technique : partitions / répliquions / scalabilité



Lexique



Lexique

- **Nœud (node)** : Une instance d'Elasticsearch (~ machine ?)



Lexique

- **Nœud (node)** : Une instance d'Elasticsearch (~ machine ?)
- **Cluster** : Un ensemble de nœuds



Lexique

- **Nœud (node)** : Une instance d'Elasticsearch (~ machine ?)
- **Cluster** : Un ensemble de nœuds
- **Partition (shard)** : permet de découper un index en plusieurs parties pour y distribuer les documents

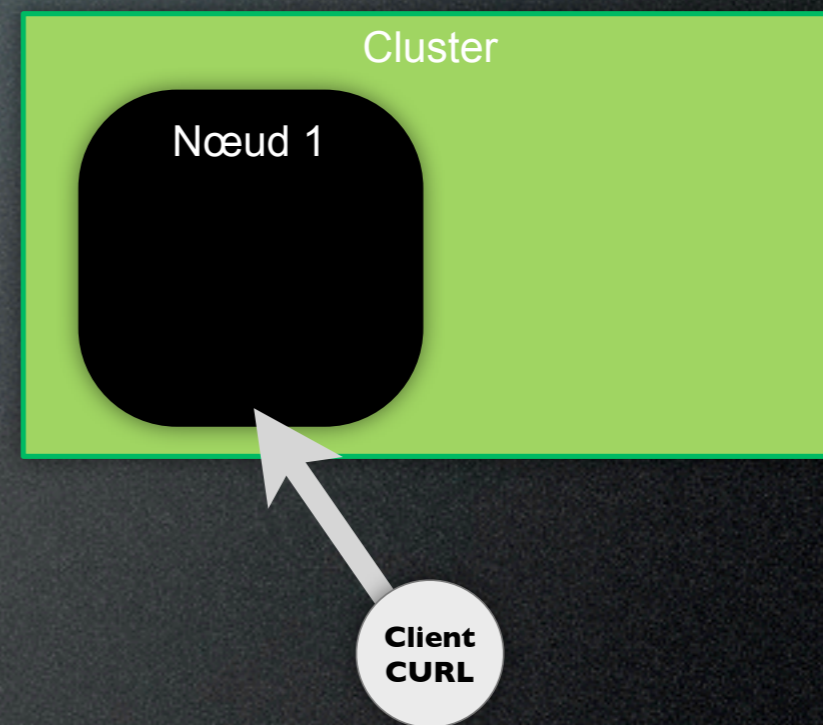


Lexique

- **Nœud (node)** : Une instance d'Elasticsearch (~ machine ?)
- **Cluster** : Un ensemble de nœuds
- **Partition (shard)** : permet de découper un index en plusieurs parties pour y distribuer les documents
- **Réplication (replica)** : recopie d'une partition en une ou plusieurs copies dans l'ensemble du cluster

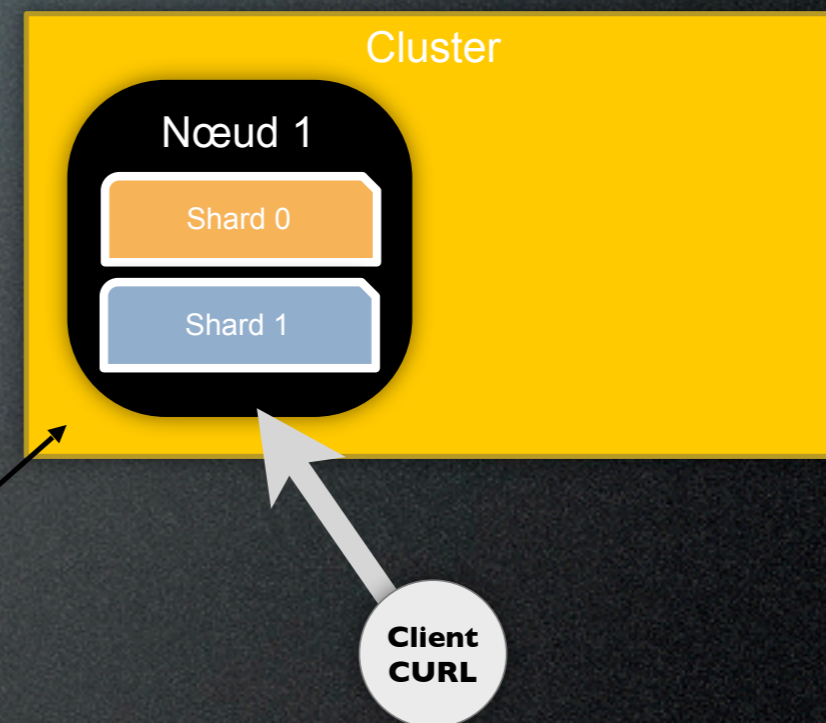


Créons un index



Créons un index

```
$ curl -XPUT localhost:9200/twitter -d '{  
  "index" : {  
    "number_of_shards" : 2,  
    "number_of_replicas" : 1  
  }  
'
```

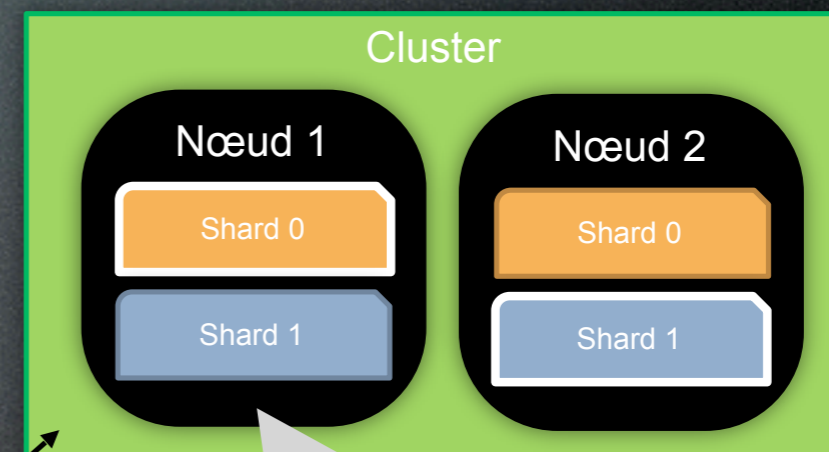


réplication non respectée



Créons un index

```
$ curl -XPUT localhost:9200/twitter -d '{  
  "index" : {  
    "number_of_shards" : 2,  
    "number_of_replicas" : 1  
  }  
'
```




réplication respectée

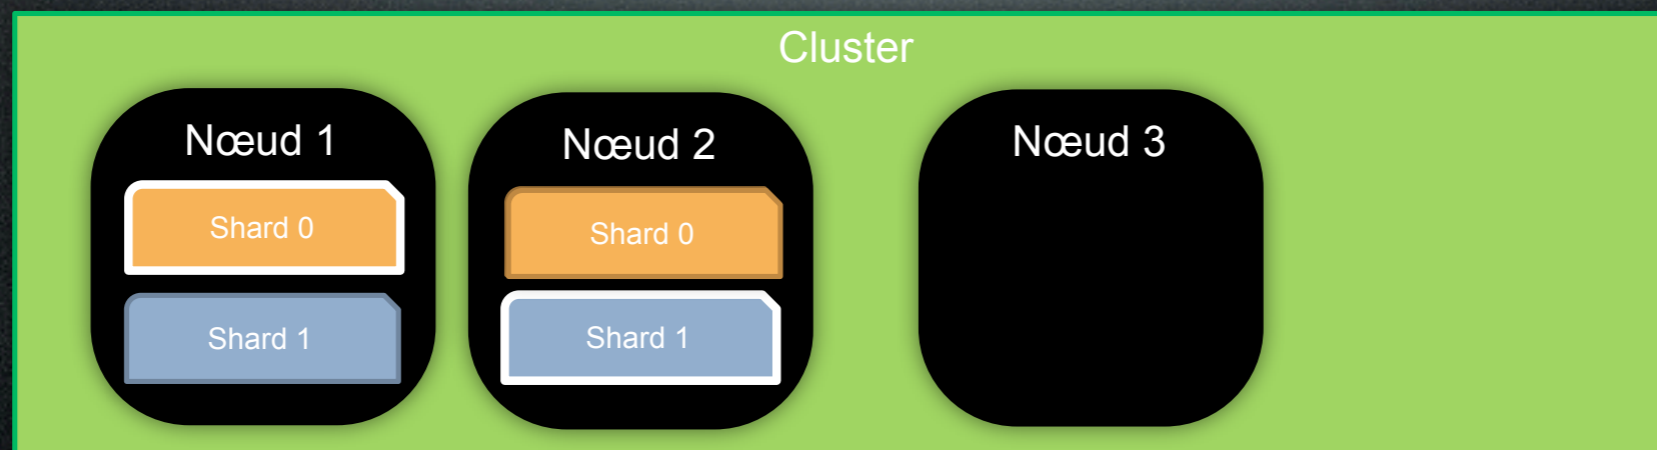
Client
CURL



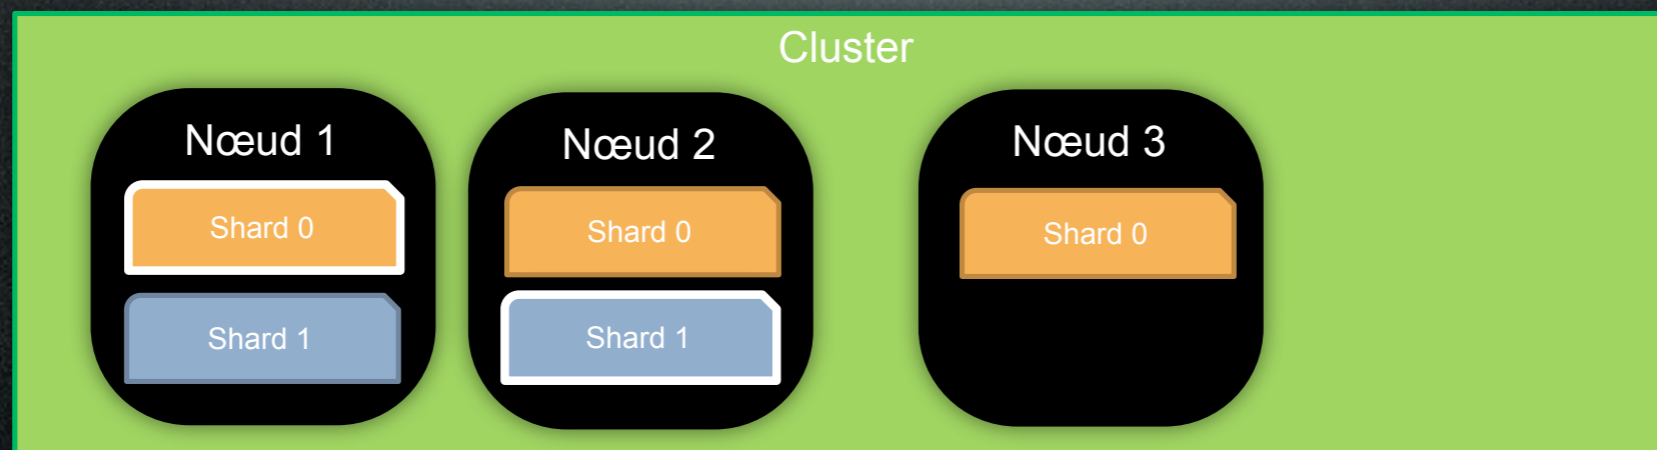
Réallocation dynamique



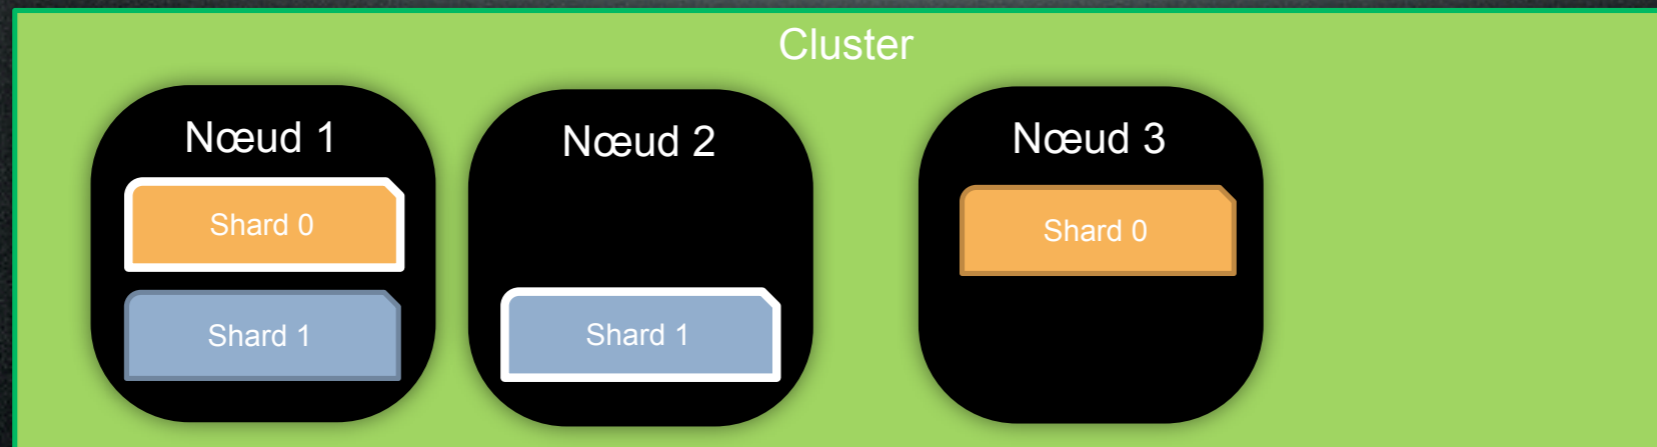
Réallocation dynamique



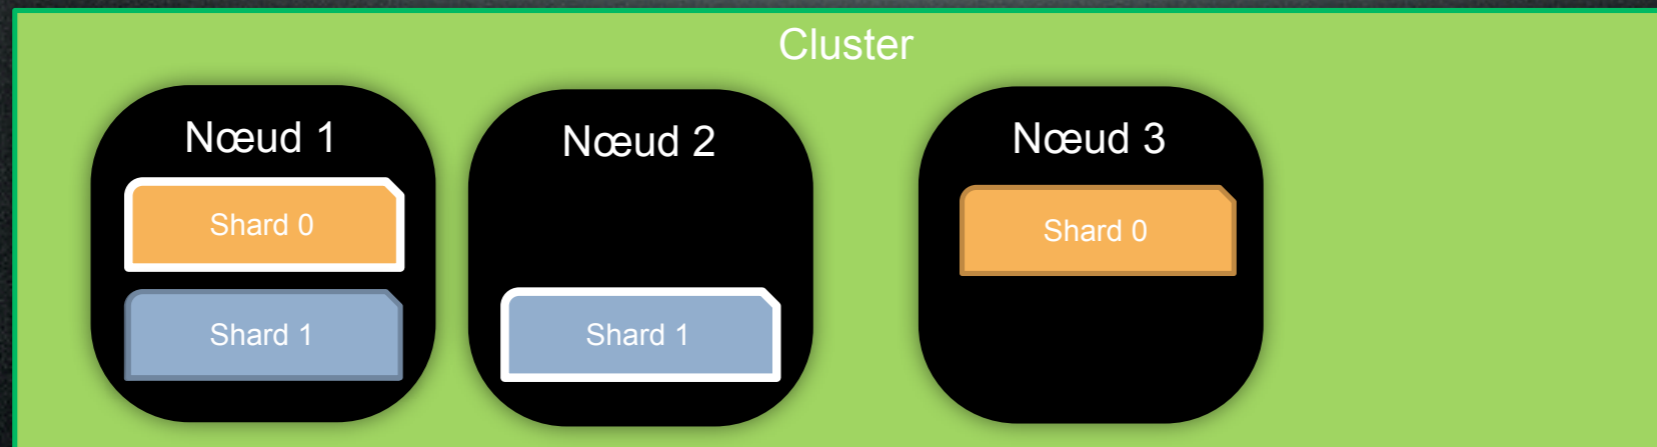
Réallocation dynamique



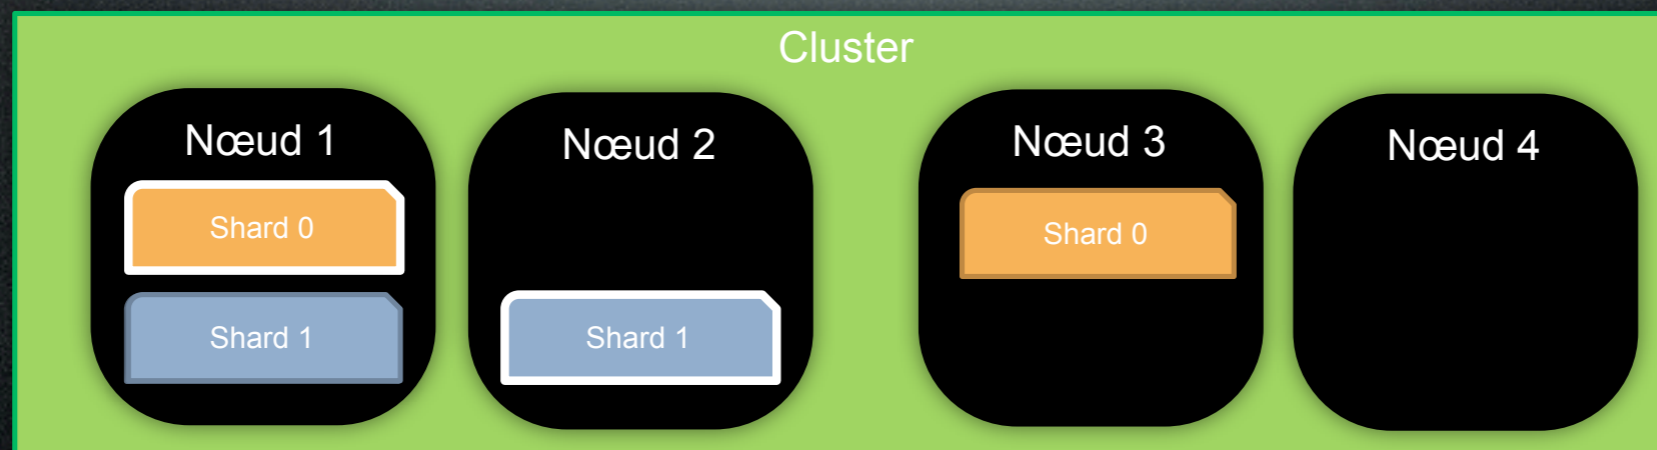
Réallocation dynamique



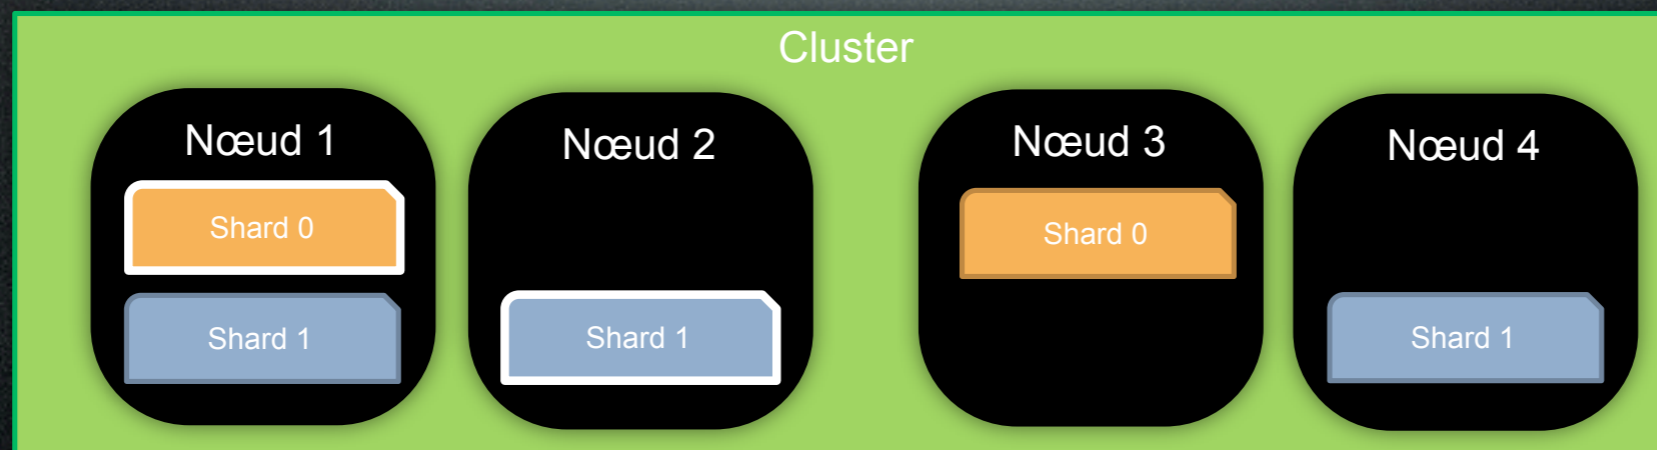
Réallocation dynamique



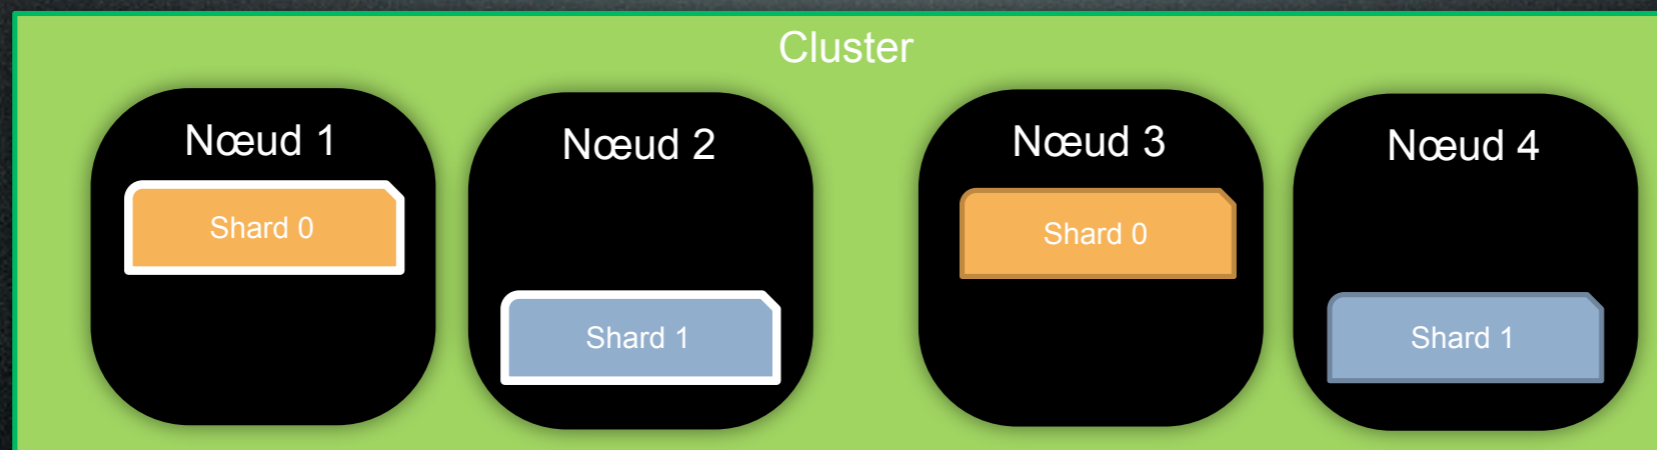
Réallocation dynamique



Réallocation dynamique



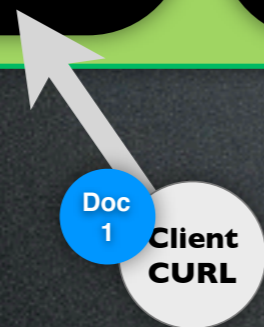
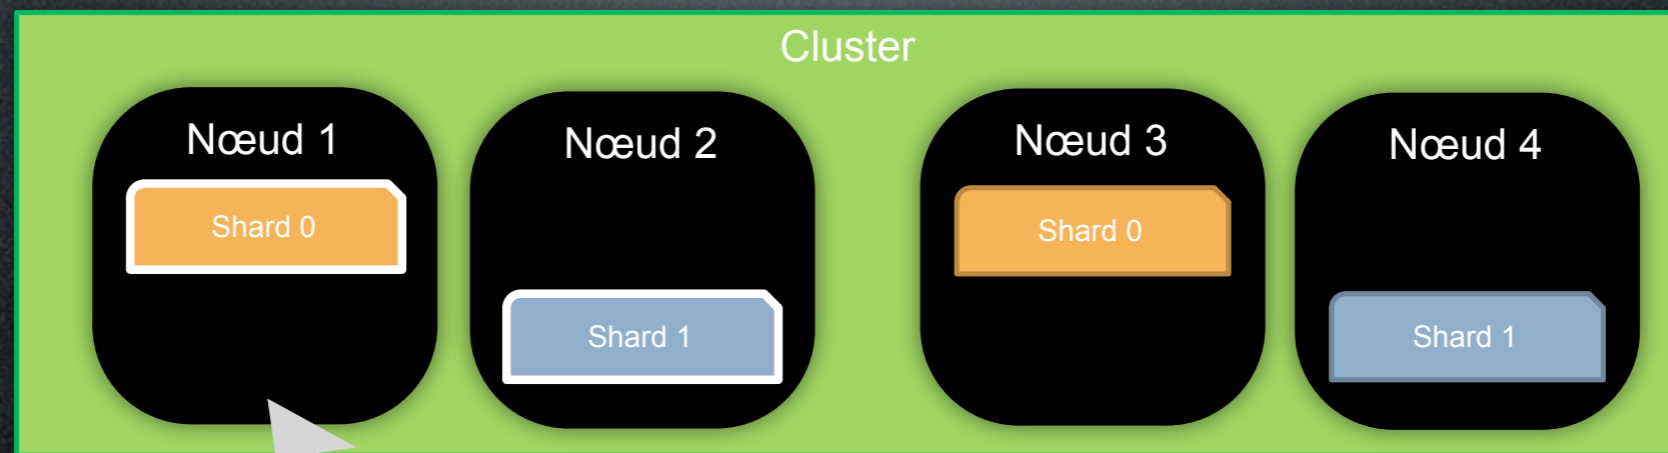
Réallocation dynamique



Le **tuning**, c'est trouver le bon équilibre entre le nombre de **nodes**, **shards** et **replicas** !



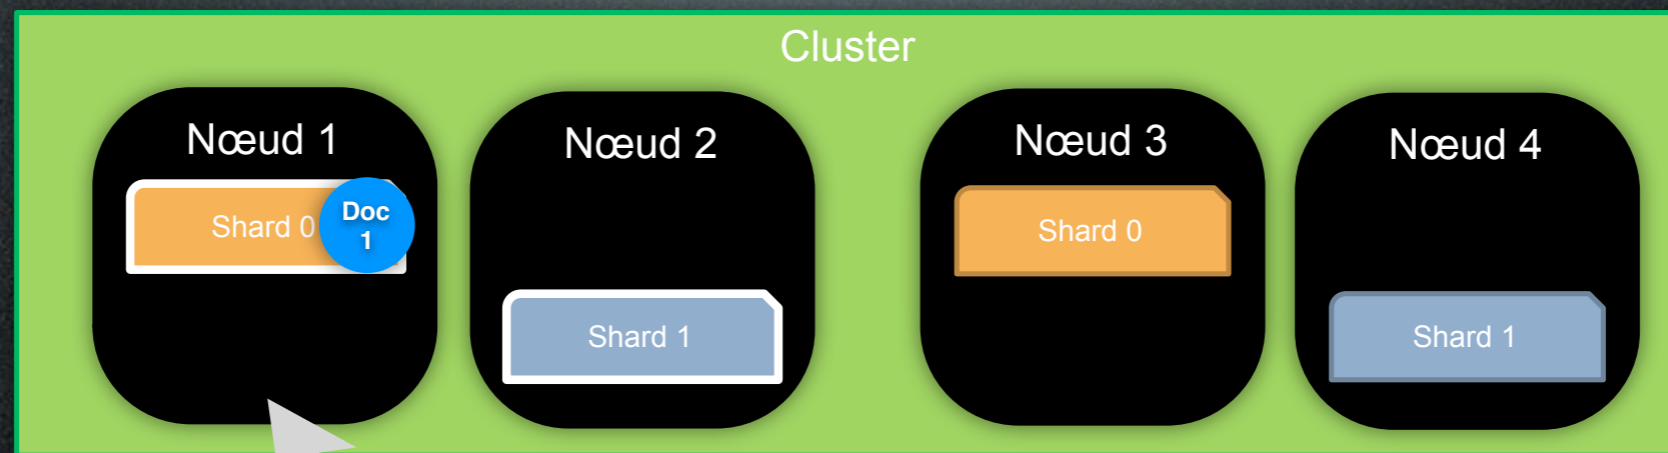
Indexons un document



```
$ curl -XPUT localhost:9200/twitter/tweet/1 -d '{
  "text": "Bienvenue à la conférence #elasticsearch pour #JUG",
  "created_at": "2012-04-06T20:45:36.000Z",
  "source": "Twitter for iPad",
  ...
}'
```



Indexons un document

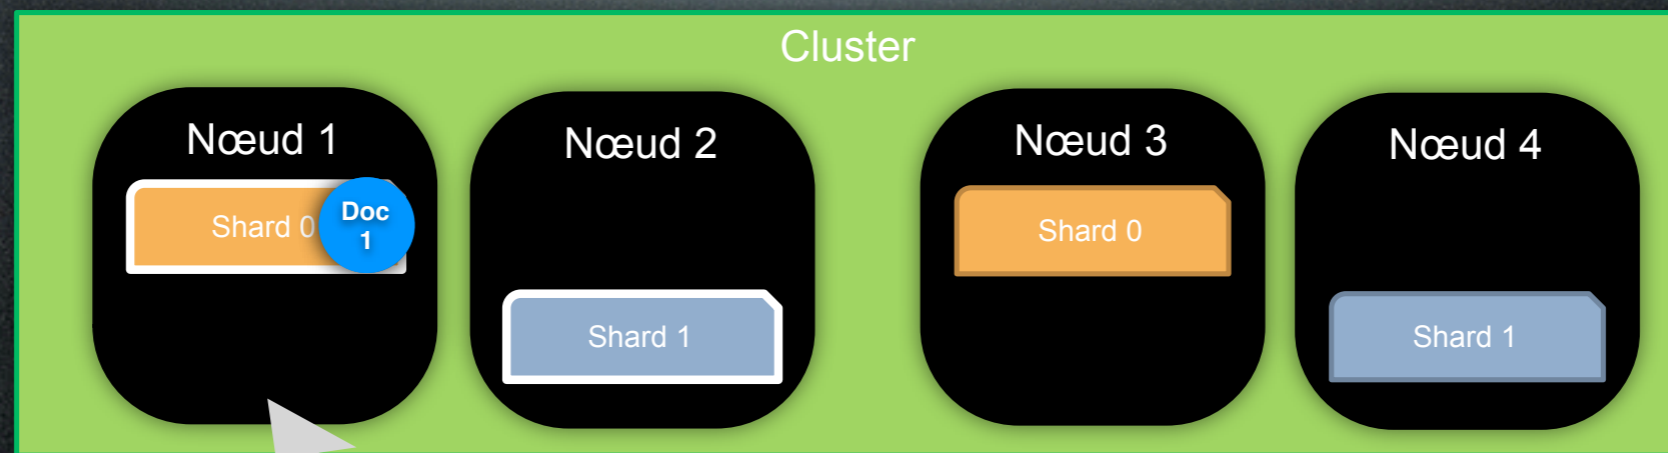


Client
CURL

```
$ curl -XPUT localhost:9200/twitter/tweet/1 -d '{
  "text": "Bienvenue à la conférence #elasticsearch pour #JUG",
  "created_at": "2012-04-06T20:45:36.000Z",
  "source": "Twitter for iPad",
  ...
}'
```



Indexons un document

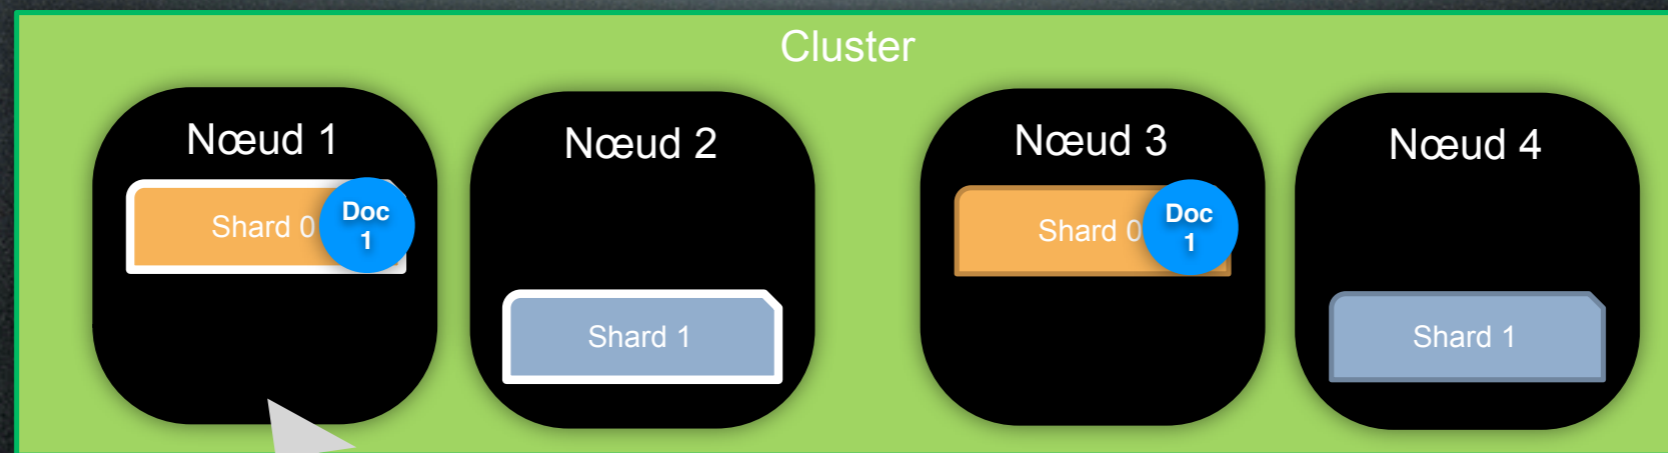


Client
CURL

```
$ curl -XPUT localhost:9200/twitter/tweet/1 -d '{
  "text": "Bienvenue à la conférence #elasticsearch pour #JUG",
  "created_at": "2012-04-06T20:45:36.000Z",
  "source": "Twitter for iPad",
  ...
}'
```



Indexons un document

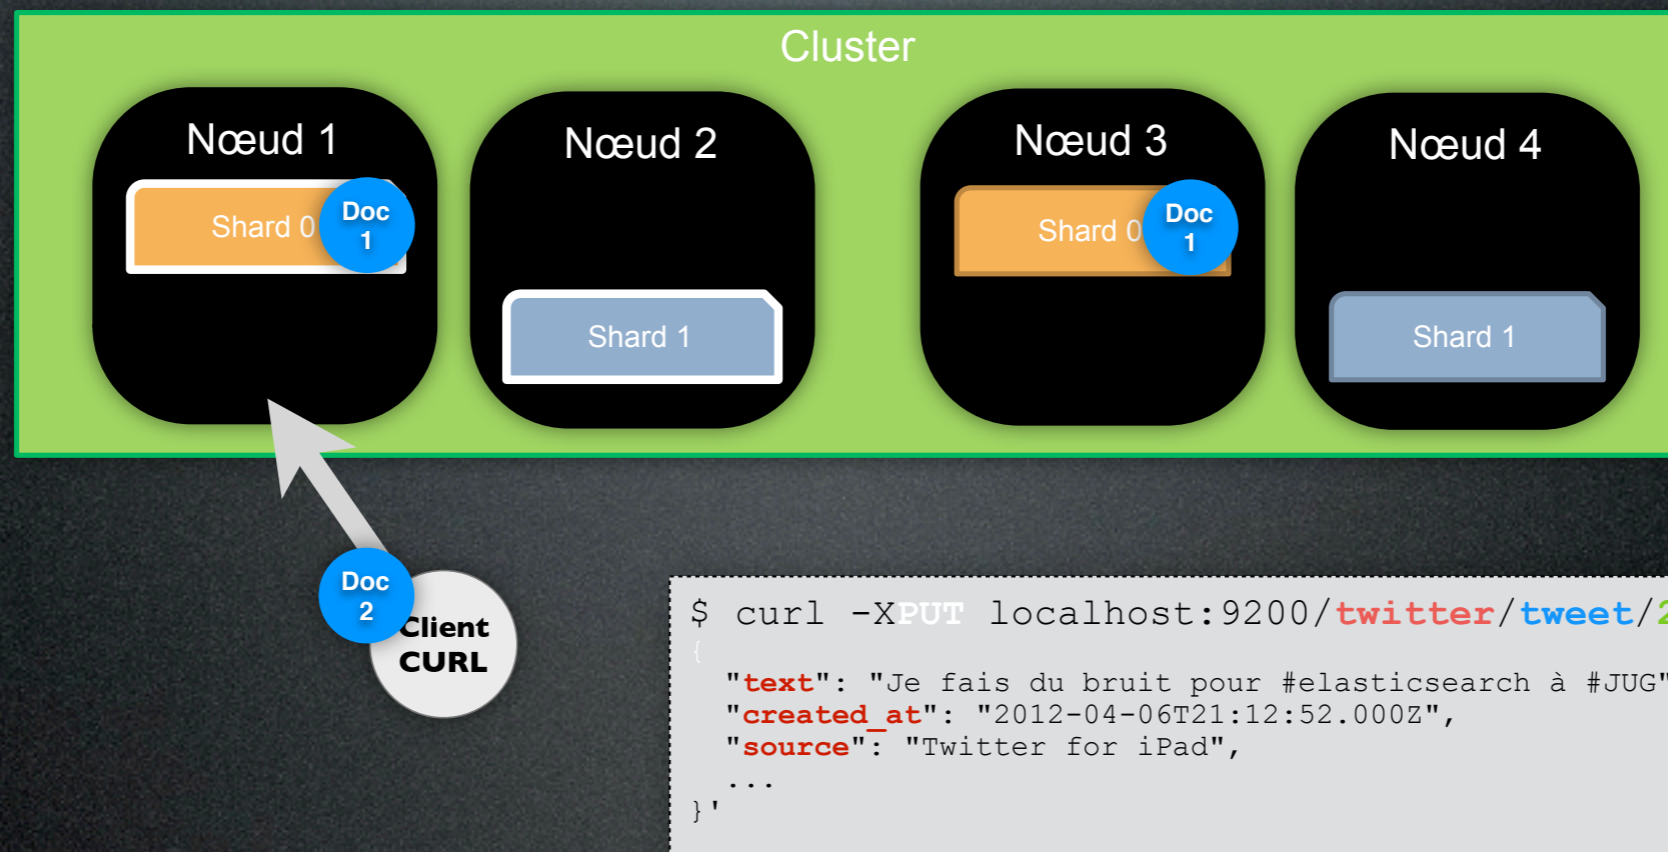


Client
CURL

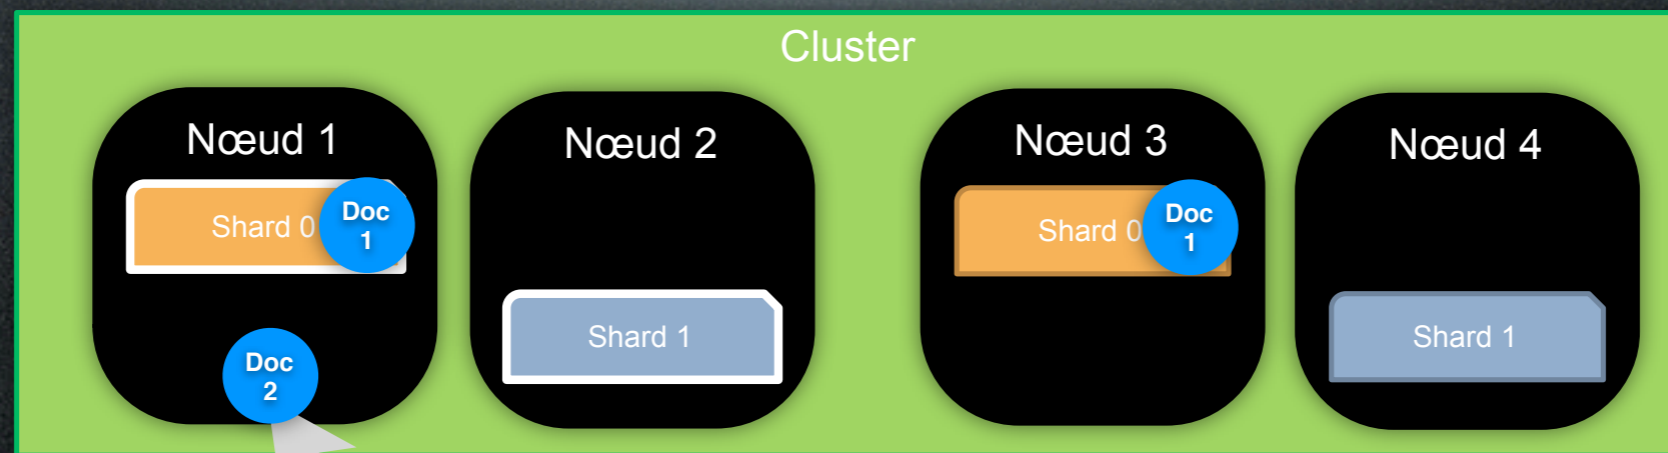
```
$ curl -XPUT localhost:9200/twitter/tweet/1 -d '{
  "text": "Bienvenue à la conférence #elasticsearch pour #JUG",
  "created_at": "2012-04-06T20:45:36.000Z",
  "source": "Twitter for iPad",
  ...
}'
```



Indexons un 2^{ème} document



Indexons un 2^{ème} document

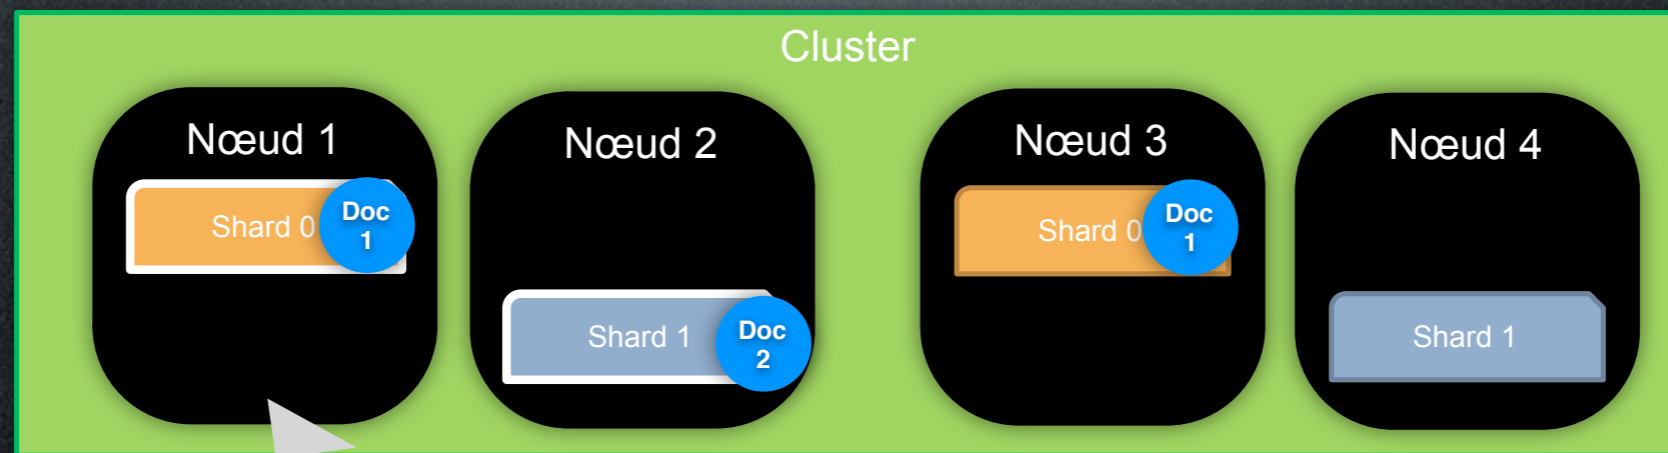


Client
CURL

```
$ curl -XPUT localhost:9200/twitter/tweet/2 -d '{
  "text": "Je fais du bruit pour #elasticsearch à #JUG",
  "created_at": "2012-04-06T21:12:52.000Z",
  "source": "Twitter for iPad",
  ...
}'
```



Indexons un 2^{ème} document

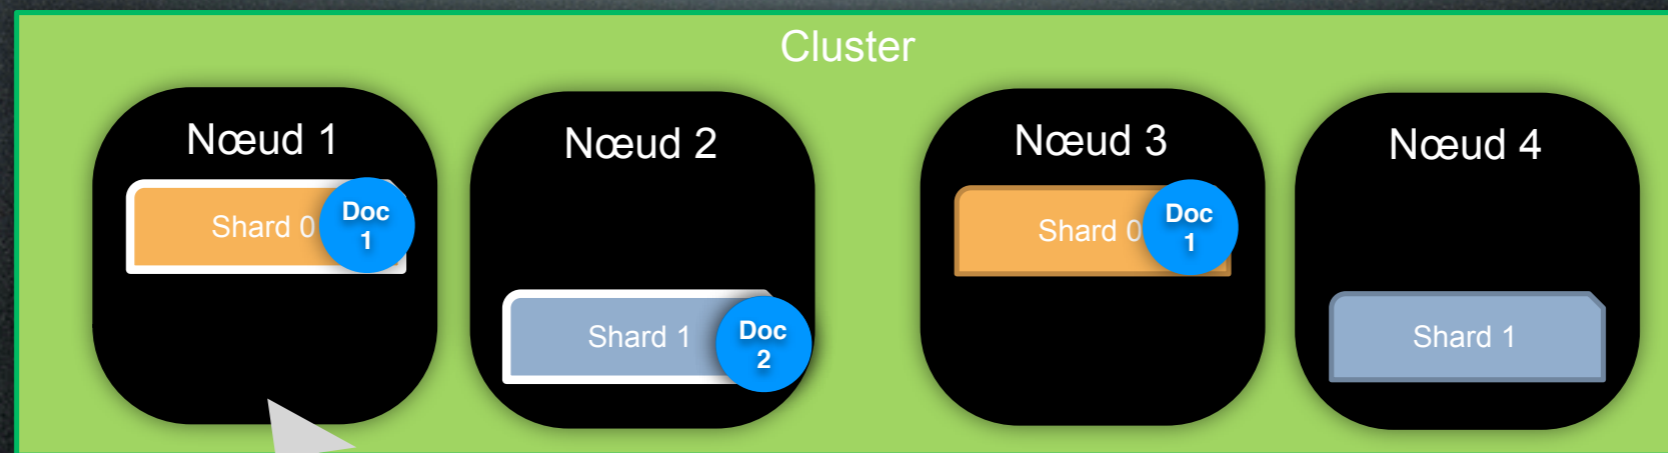


Client
CURL

```
$ curl -XPUT localhost:9200/twitter/tweet/2 -d '{
  "text": "Je fais du bruit pour #elasticsearch à #JUG",
  "created_at": "2012-04-06T21:12:52.000Z",
  "source": "Twitter for iPad",
  ...
}'
```



Indexons un 2^{ème} document

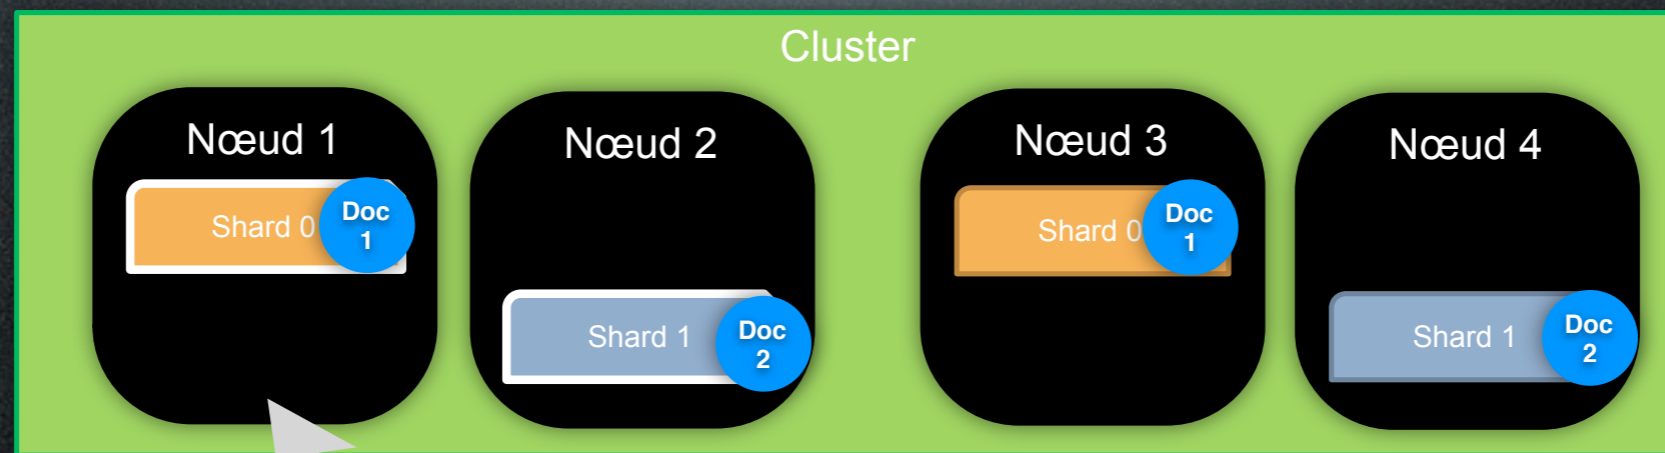


Client
CURL

```
$ curl -XPUT localhost:9200/twitter/tweet/2 -d '{
  "text": "Je fais du bruit pour #elasticsearch à #JUG",
  "created_at": "2012-04-06T21:12:52.000Z",
  "source": "Twitter for iPad",
  ...
}'
```



Indexons un 2^{ème} document

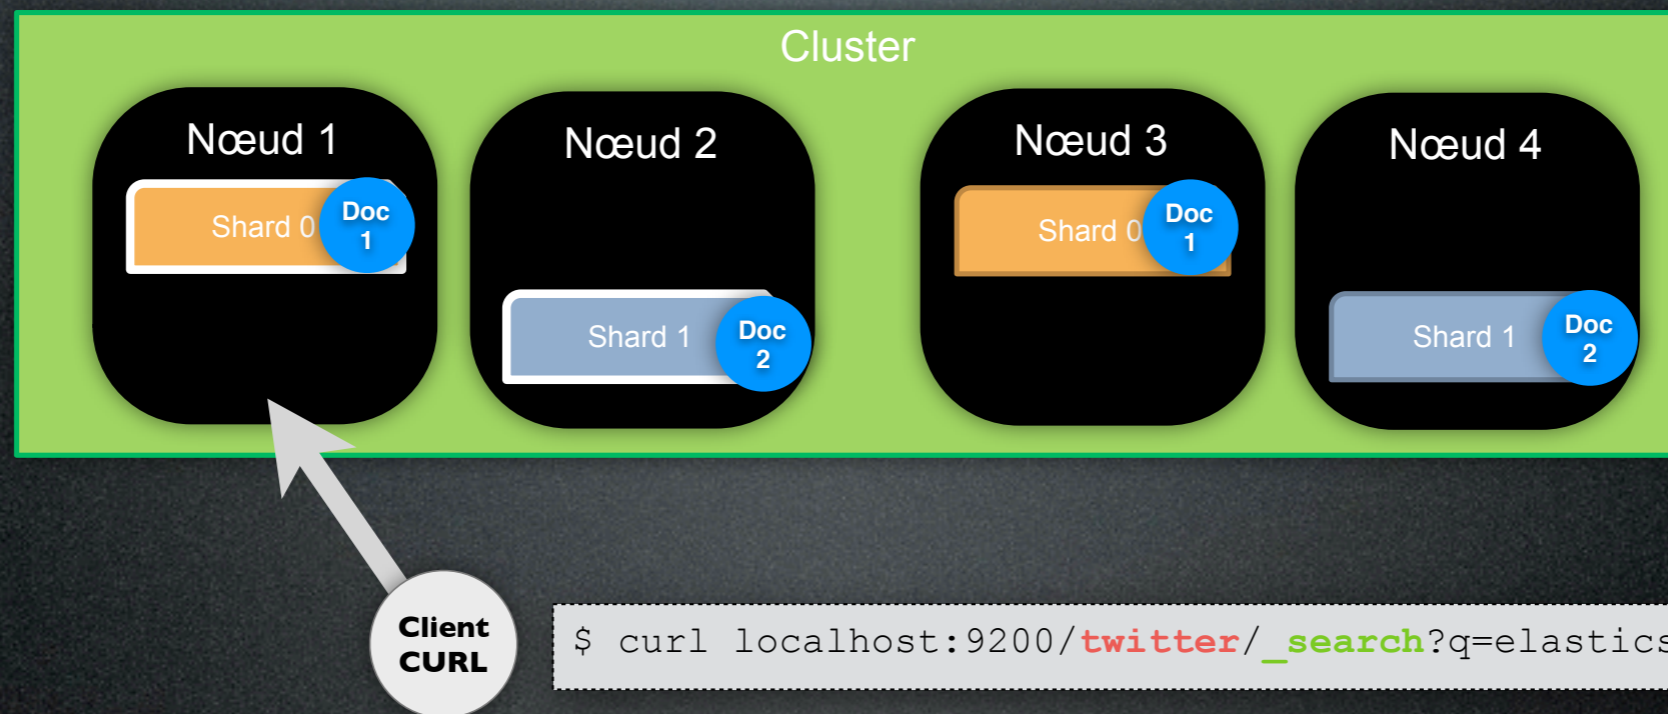


Client
CURL

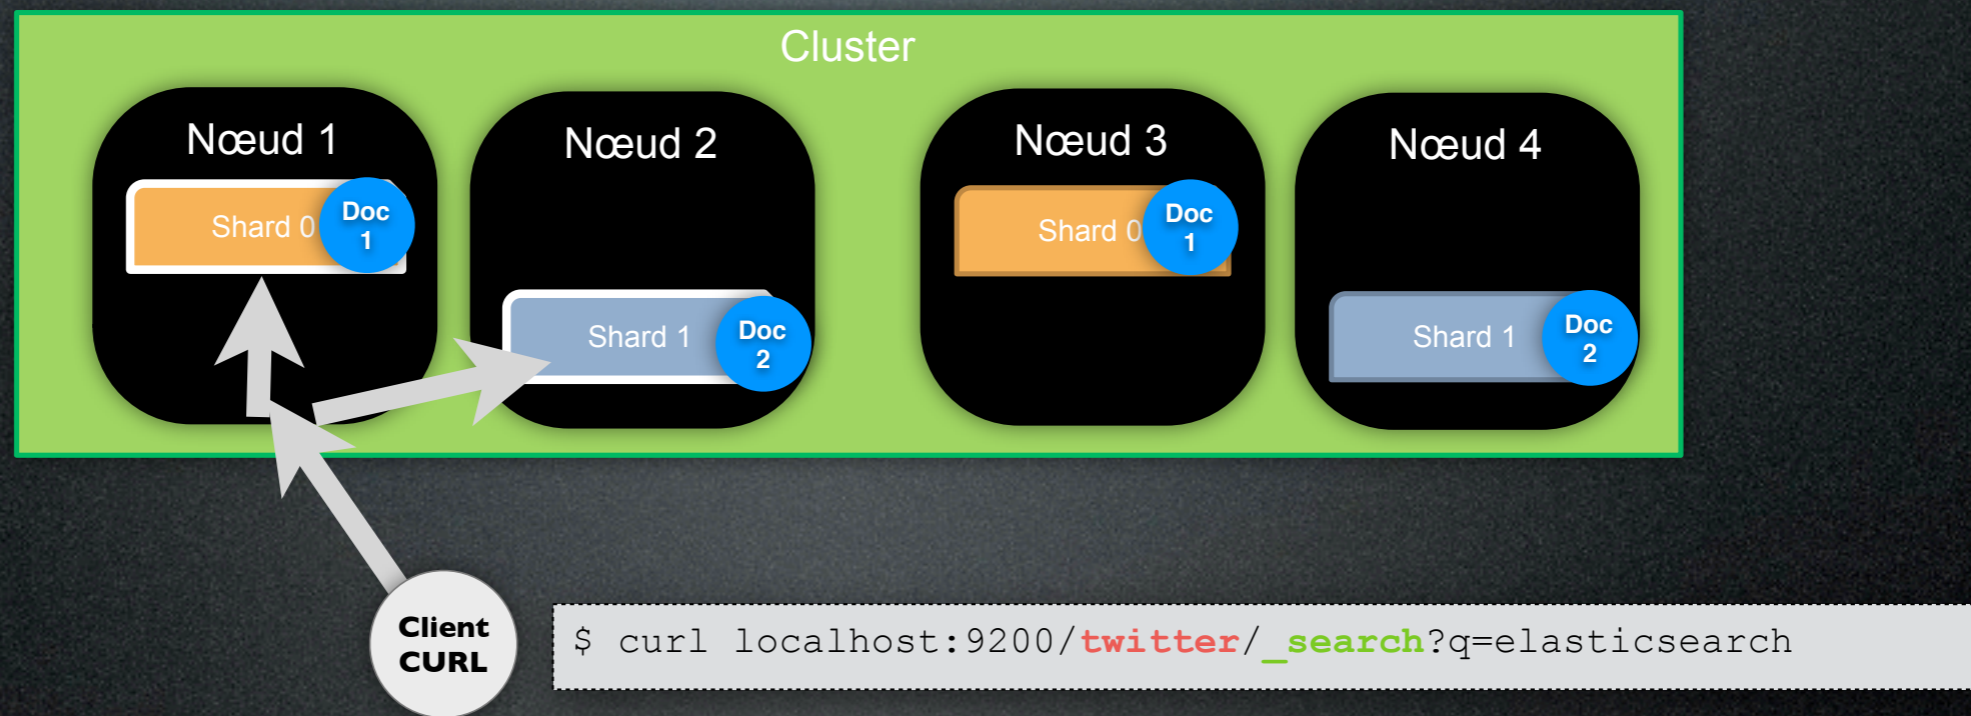
```
$ curl -XPUT localhost:9200/twitter/tweet/2 -d '{
  "text": "Je fais du bruit pour #elasticsearch à #JUG",
  "created_at": "2012-04-06T21:12:52.000Z",
  "source": "Twitter for iPad",
  ...
}'
```



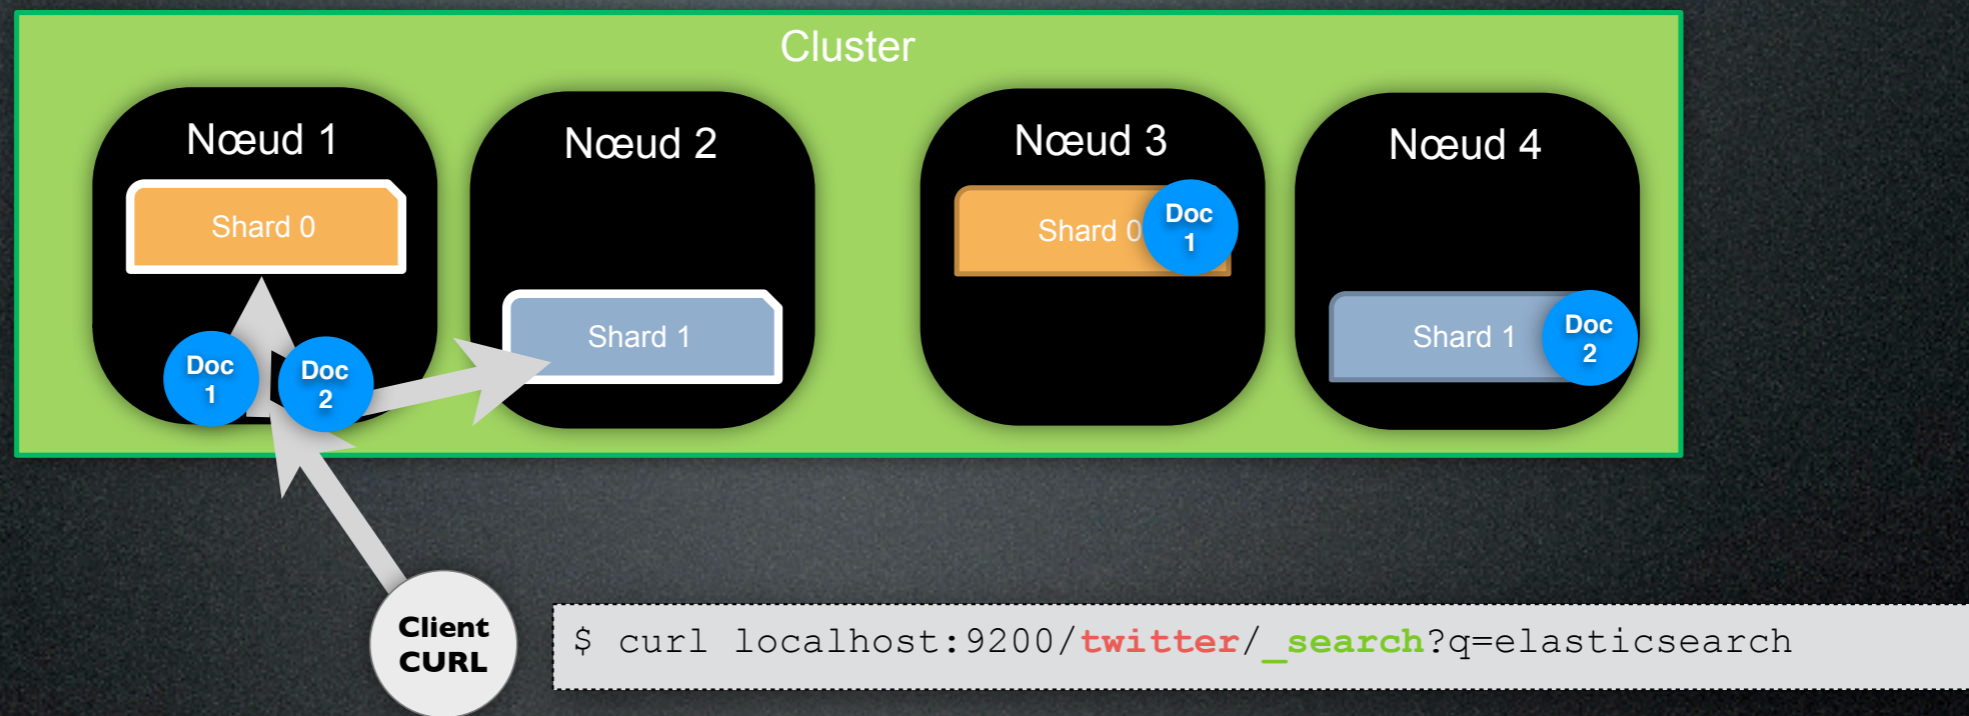
Cherchons !



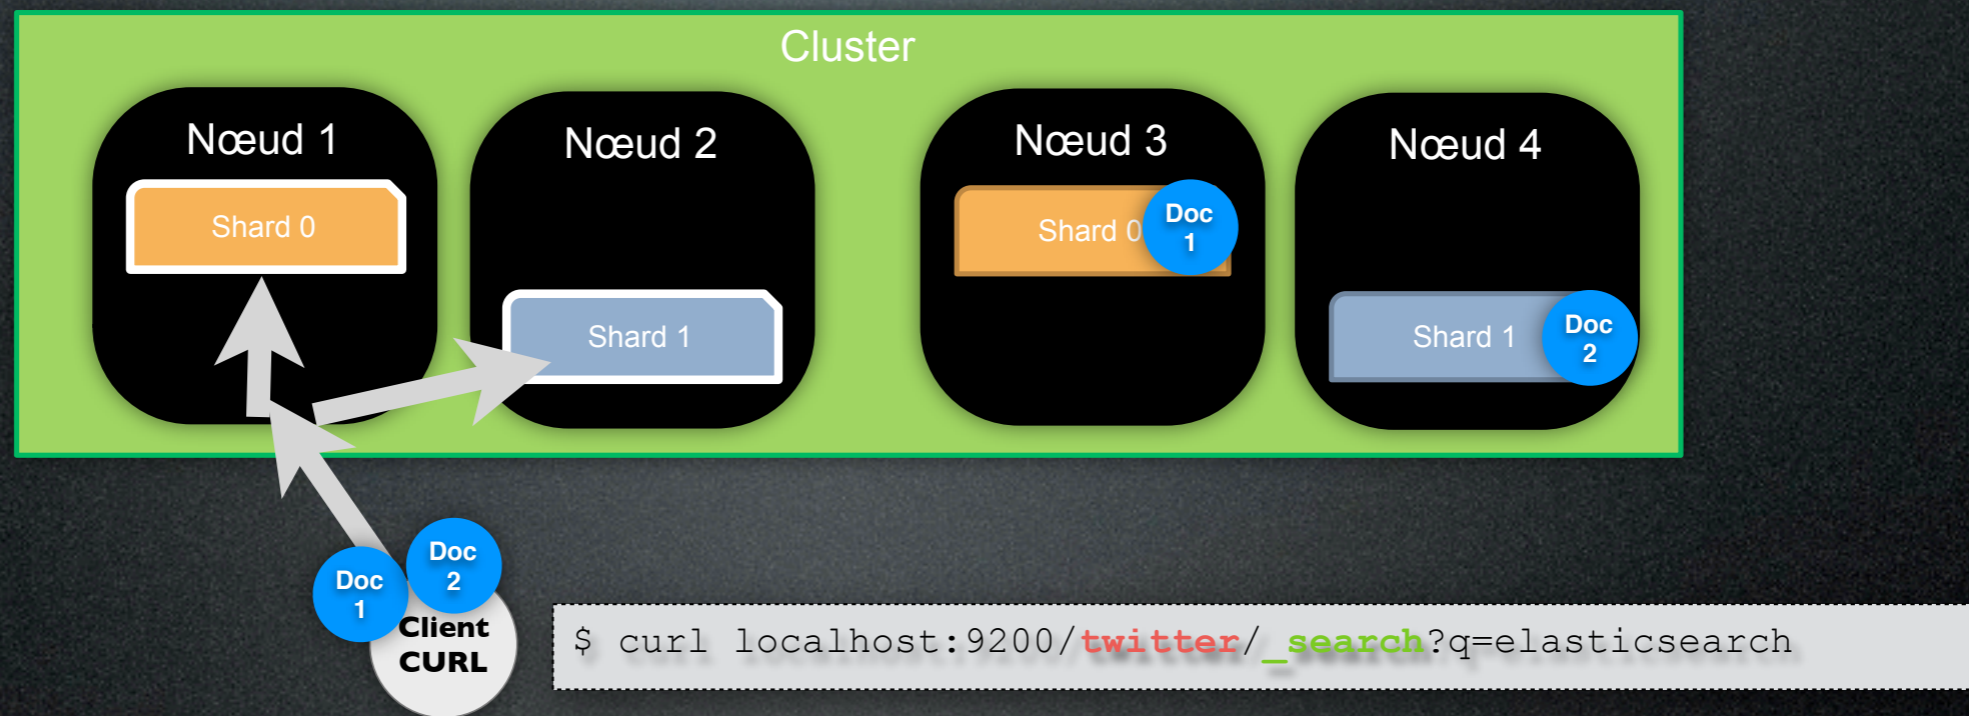
Cherchons !



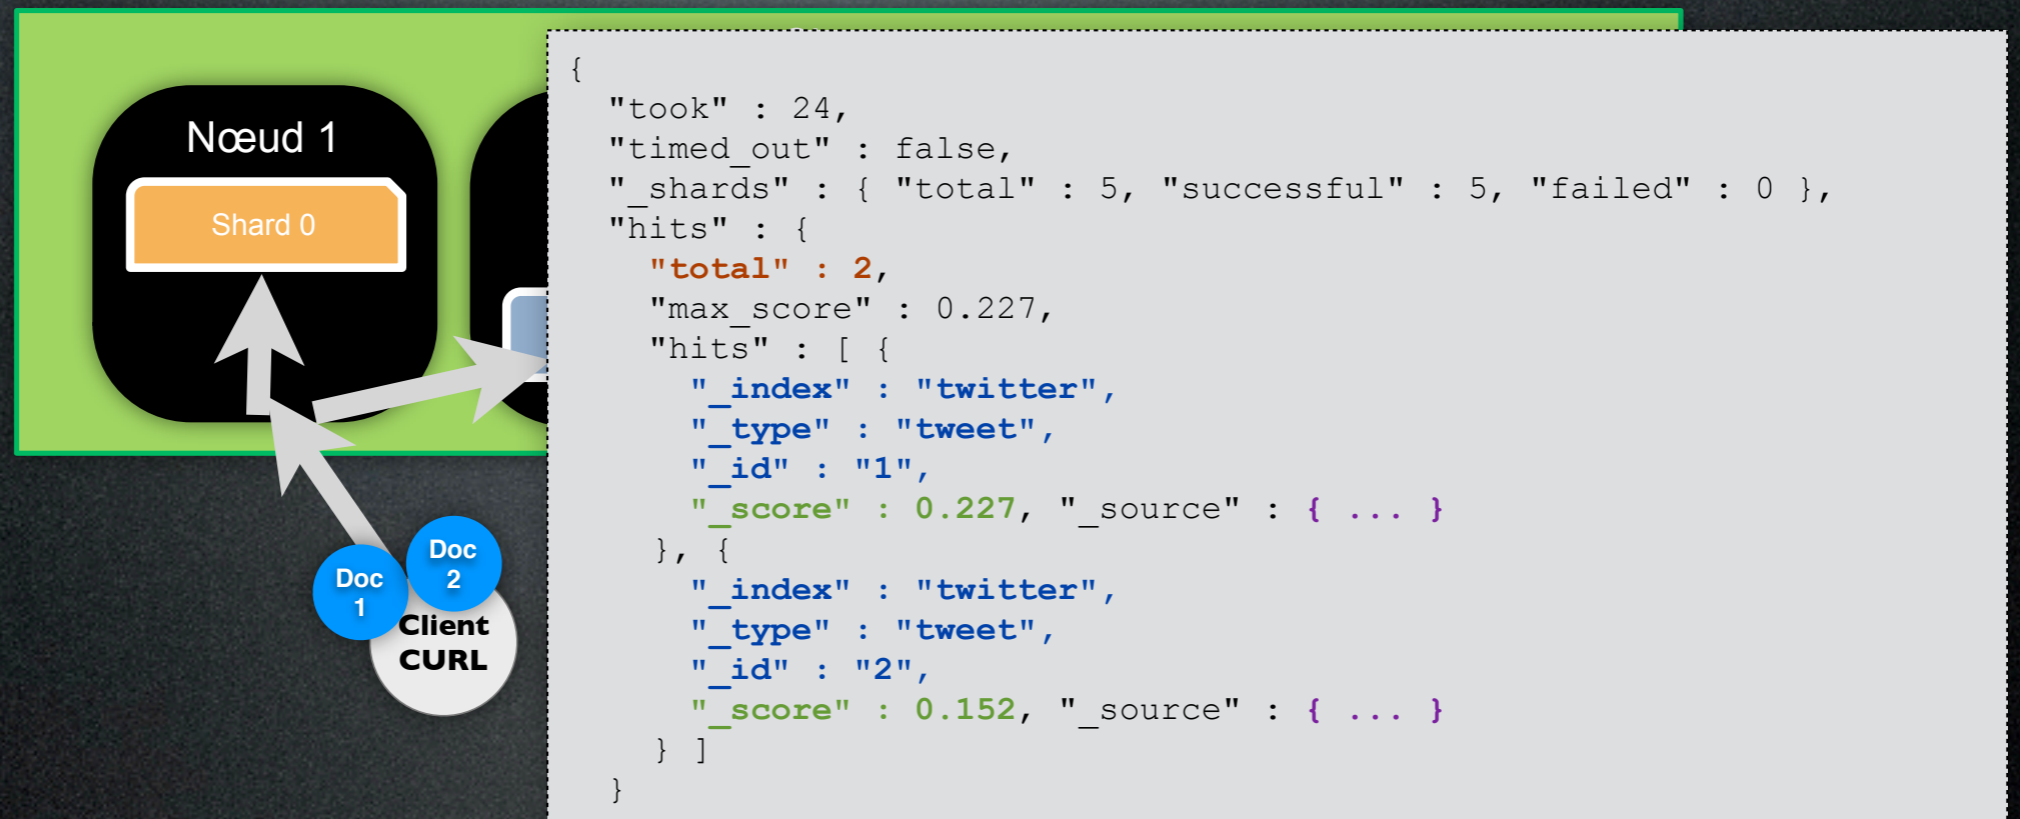
Cherchons !




Cherchons !




Cherchons !



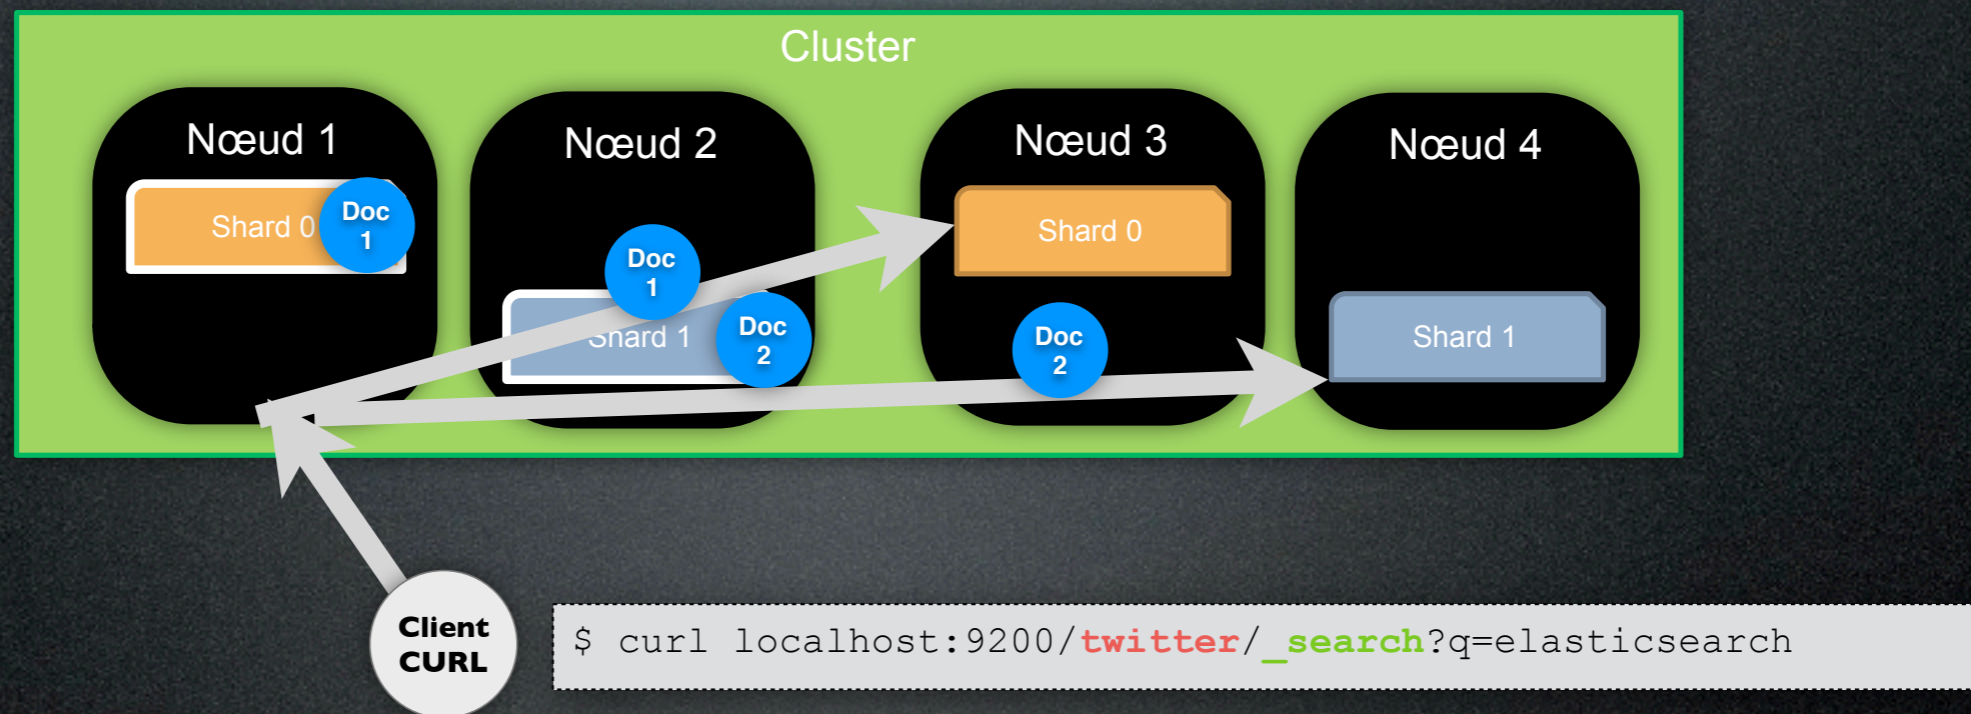
Cherchons encore !



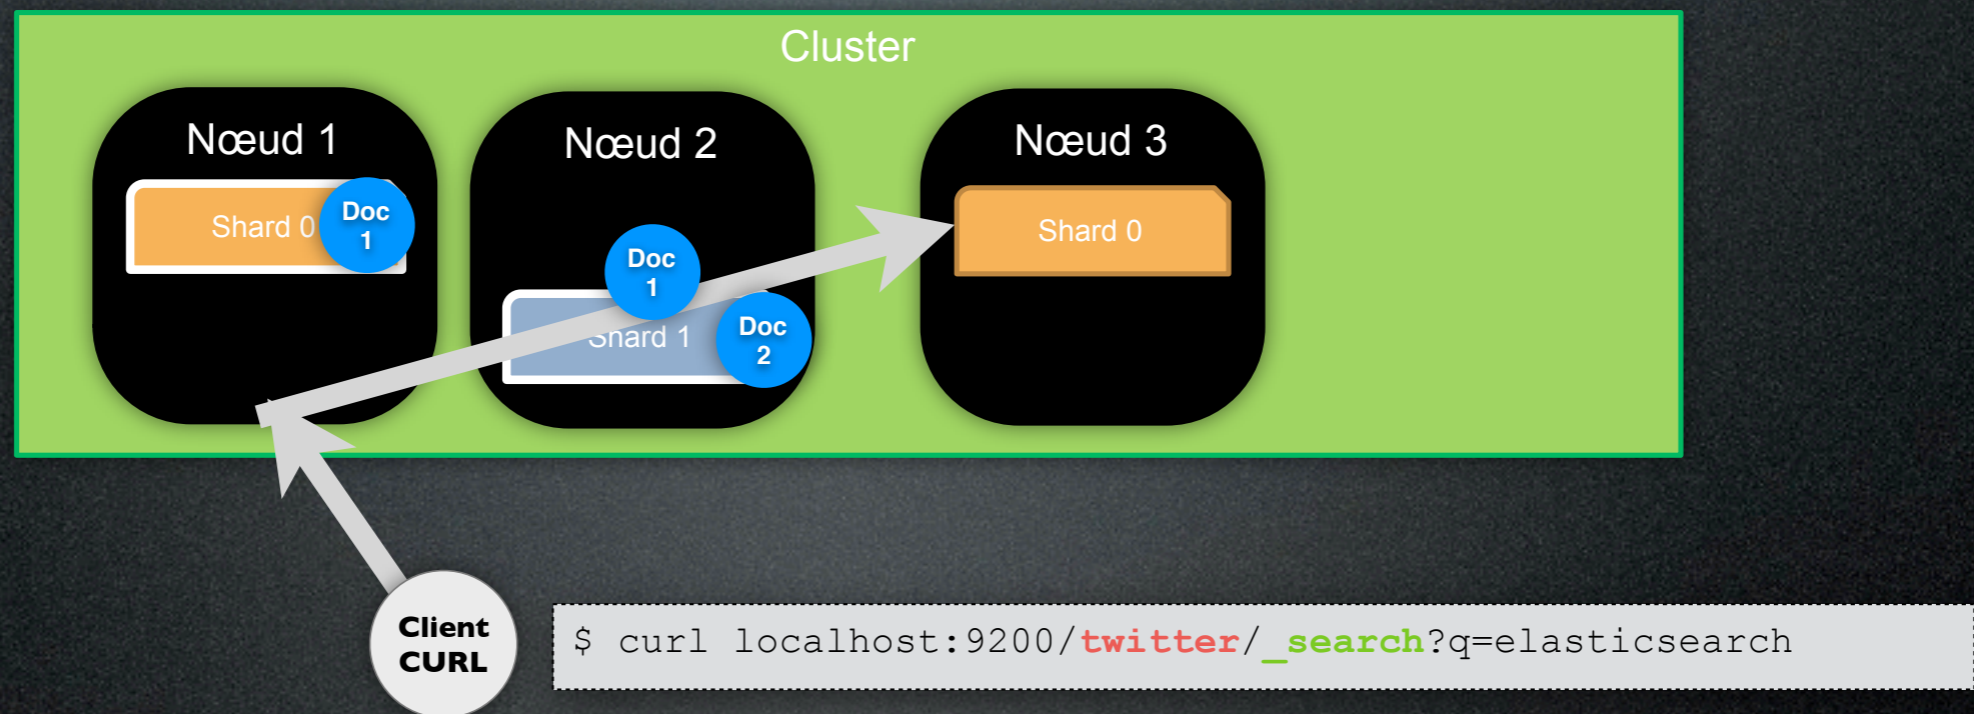
Cherchons encore !



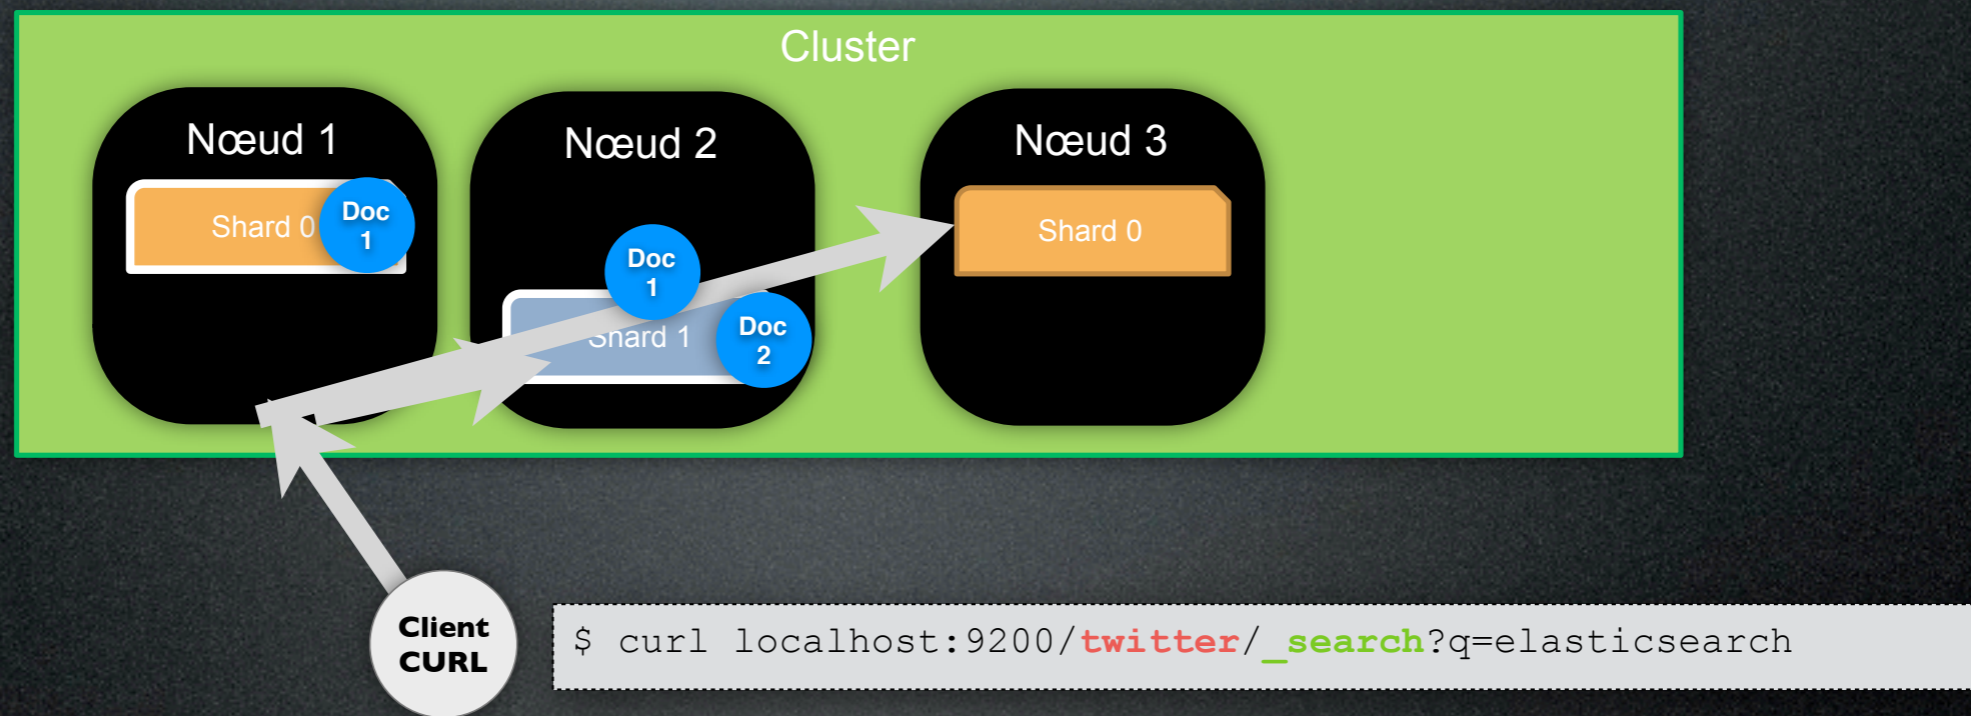
Cherchons encore !




Cherchons encore !



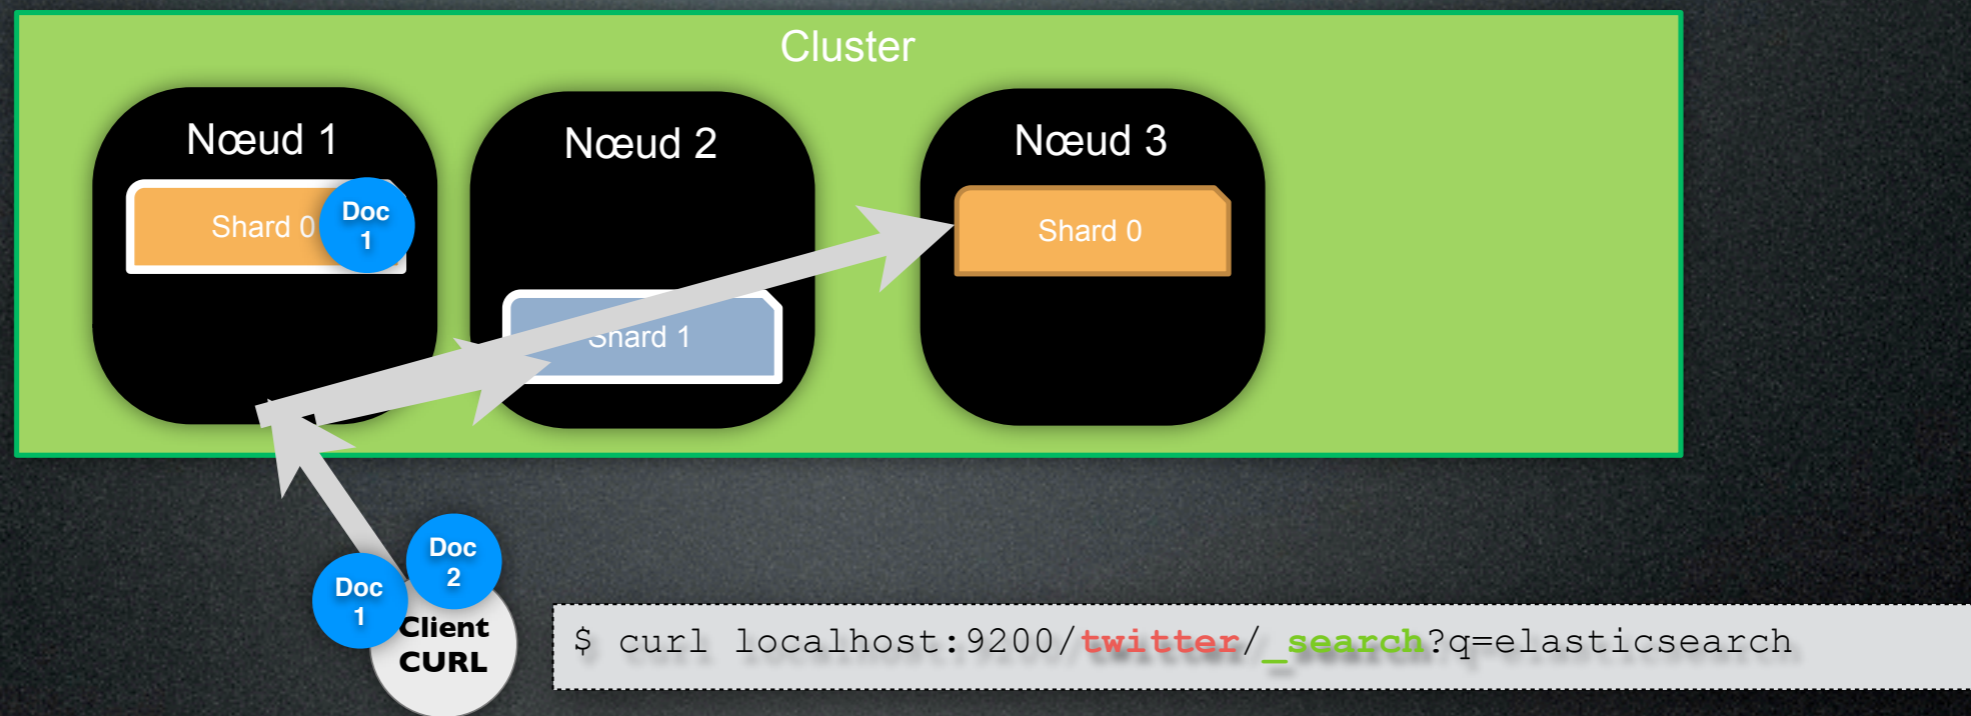
Cherchons encore !



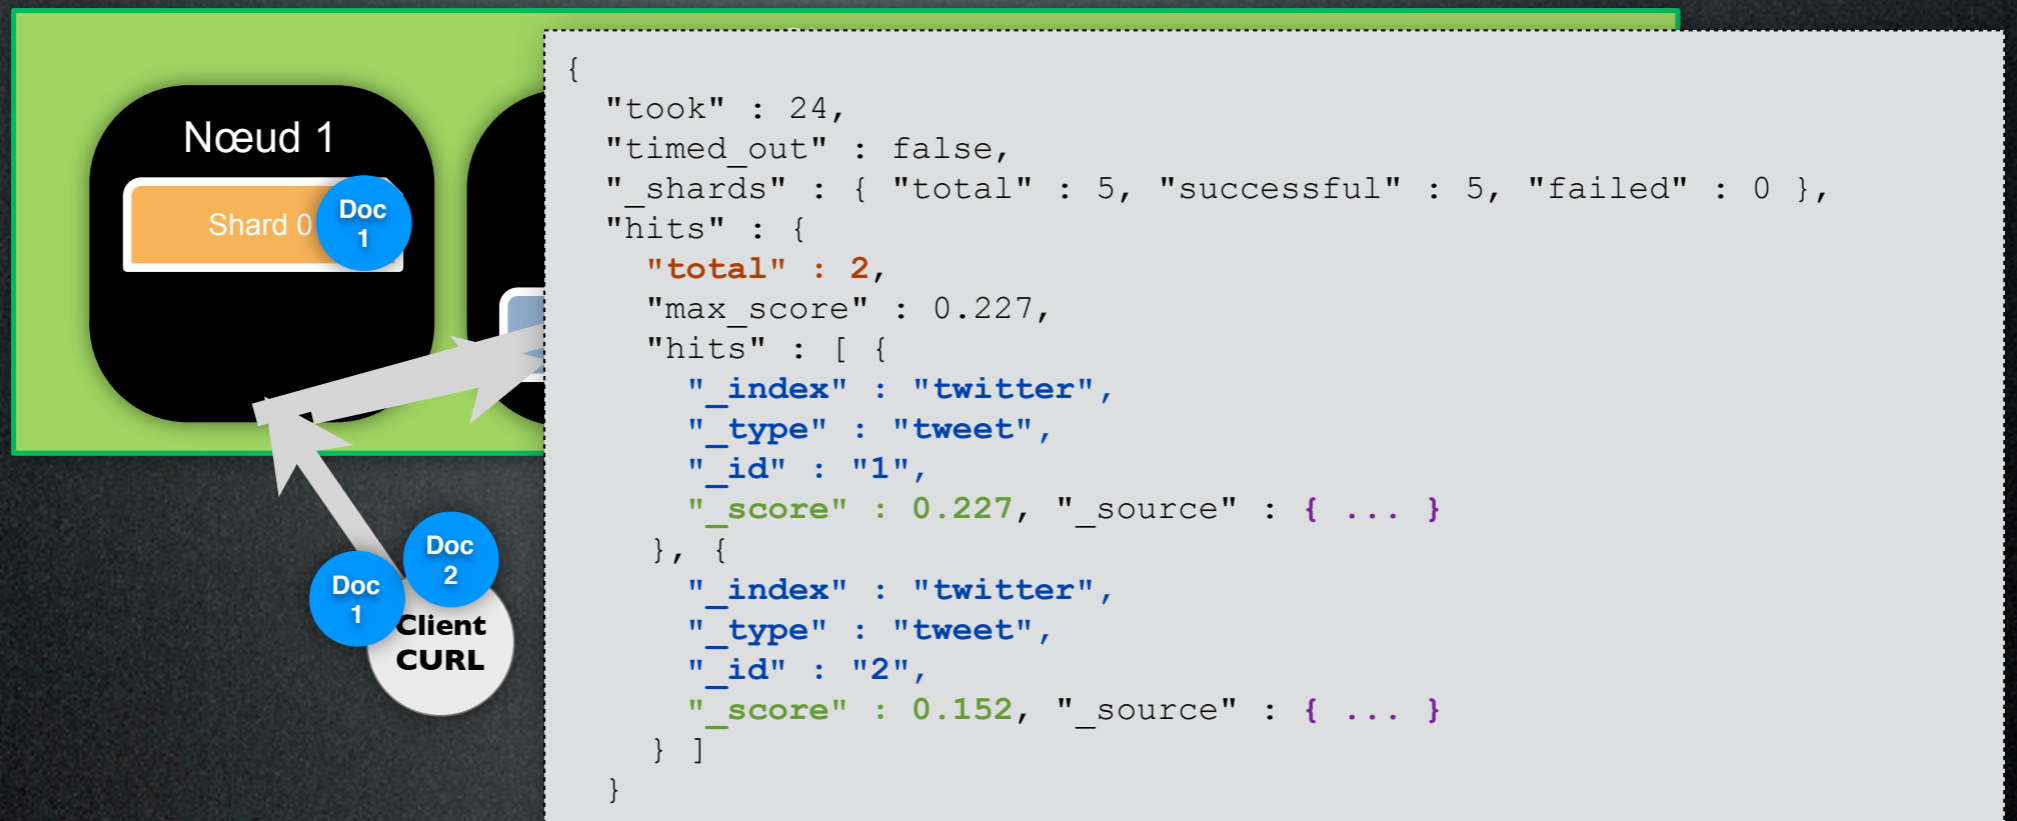
Cherchons encore !





Cherchons encore !



Cherchons encore !



La communauté



~ 100 contributeurs directs au projet (+ de 3200 watchers et + de 620 forks)

La communauté



elasticsearch.



mozilla
FOUNDATION



~ 100 contributeurs directs au projet (+ de 3200 watchers et + de 620 forks)

La communauté



elasticsearch.



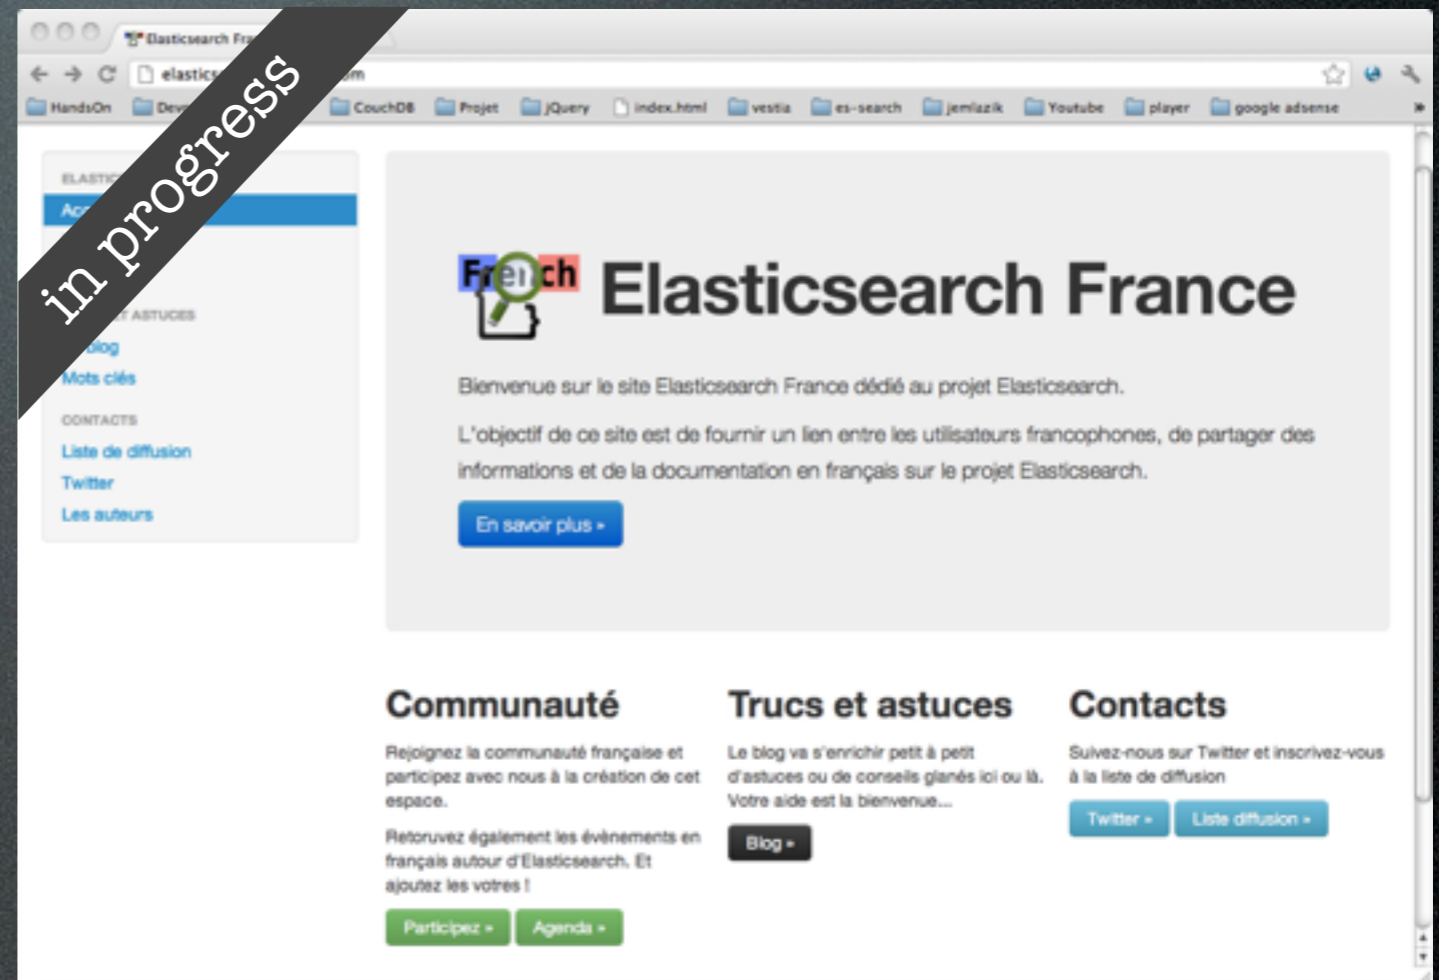
~170 inscrits sur la mailing list, 70 messages / mois, ~270 followers

Rejoignez le mouvement !

@ElasticsearchFR



www.elasticsearch.fr



Questions ?

Posez aussi vos questions sur elasticsearch-fr@googlegroups.com



Prochaines rencontres sur <http://www.meetup.com/elasticsearchfr/>



Slides sur <http://fr.slideshare.net/dadoonet>

Sources sur <https://github.com/elasticsearchfr/talks>

