



Elasticsearch.

Le moteur de recherche
élastique pour tous



David Pilato, Elasticsearch.com, Paris



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

UploadSettings



3 documents found in 78 milliseconds

[NOTICE.txt](#)

by The **Apache** Software Foundation (<http://www.apache.org/>).

 [Apache.pdf](#)

Configurer **Apache** mars 2007 i Ce dossier est public et libre au sens de la licence GNU-FDL : Permission serveur Web **Apache** [10] : ins- tallation et configuration de base, notion de Directive, cre´ation d'un re´pertoire serveur Web **Apache** est un logiciel Open source [6] et multiplate-forme : on peut l'installer aussi bien sur serveur **Apache** doit eˆtre installe´ en tant que ((service)) ; – Si vous faites cohabiter ce serveur avec serveur **Apache** doit alors eˆcouter les requetes 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

Elasticsearch.com



- Créée en 2012 par les auteurs d'Elasticsearch
 - **Formation** (publique et intra)
 - **Consulting** (support dév)
 - Abonnement annuel **support pour la production** avec 3 niveaux de SLA (délai de réponse et disponibilité)

Pour la démo

Faites du bruit sur Twitter

avec le hashtag

#elasticsearch



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	
Web Title	
Url	
Category	Select v
Web Description	
Keywords	
Contact Name	
Contact Email	
Featured Links 	Select v
Cool Links 	Select v
Bold Links	Select v
Icon	☐  ☐  ☐  ☐  ☐  ☐ 
Rating Average 	Select v
Number of Votes	between and
Total Hits	between and
Hits Today	between and
IP Address	
Submission Software Name	





Moteur de recherche ?

- un moteur d'**indexation de documents**





Moteur de recherche ?

- un moteur d'**indexation de documents**



Moteur de recherche ?

- un moteur d'**indexation de documents**



- un moteur de **recherche dans les index**



Elasticsearch





Elasticsearch

- C'est un **moteur** !
- NoSQL orienté document
- Apache Lucene
- HTTP / REST / JSON
- Distribué, Scalable, Cloud ready



Points clés



Points clés

- **Simple:** start in ~~5 minutes~~ 30 seconds



Points clés

- **Simple:** start in ~~5 minutes~~ 30 seconds
- **Efficace:** just start new nodes!



Points clés

- **Simple:** start in ~~5 minutes~~ 30 seconds
- **Efficace:** just start new nodes!
- **Puissant:** 20-300ms!



Points clés

- **Simple:** start in ~~5 minutes~~ 30 seconds
- **Efficace:** just start new nodes!
- **Puissant:** 20-300ms!
- **Complet:** built-in + plugins



Penser « document » !

- **Document** : Un objet représentant les données (au sens NoSQL).
Penser "recherche", c'est oublier le SGBDR et penser "Documents"
- **Type** : Regroupe des documents de même type
- **Index** : Espace logique de stockage des documents dont les types sont fonctionnellement communs



Penser « document » !

- **Document** : Un objet représentant les données (au sens NoSQL).
Penser "recherche", c'est oublier le SGBDR et penser "Documents"

```
{
  "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 !" }
}
```

- **Type** : Regroupe des documents de même type
- **Index** : Espace logique de stockage des documents dont les types sont fonctionnellement communs



Penser « document » !

- **Document** : Un objet représentant les données (au sens NoSQL).
Penser "recherche", c'est oublier le SGBDR et penser "Documents"

```
{
  "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 !" }
}
```

- **Type** : Regroupe des documents de même type



Penser « document » !

- **Document** : Un objet représentant les données (au sens NoSQL).
Penser "recherche", c'est oublier le SGBDR et penser "Documents"

```
{
  "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 !" }
}
```

- **Type** : Regroupe des documents de même type
- **Index** : Espace logique de stockage des documents dont les types sont fonctionnellement communs



Interagir avec Elasticsearch



Interagir avec Elasticsearch

- API REST : `http://host:port/[index]/[type]/[_action/id]`
Méthodes HTTP : **GET, POST, PUT, DELETE**



Interagir avec Elasticsearch

- API REST : `http://host:port/[index]/[type]/[_action/id]`
Méthodes HTTP : **GET**, **POST**, **PUT**, **DELETE**
- Documents
 - `curl -XPUT http://localhost:9200/twitter/tweet/1`



Interagir avec Elasticsearch

- API REST : `http://host:port/[index]/[type]/[_action/id]`
Méthodes HTTP : **GET**, **POST**, **PUT**, **DELETE**
- Documents
 - `curl -XPUT http://localhost:9200/twitter/tweet/1`
 - `curl -XGET http://localhost:9200/twitter/tweet/1`



Interagir avec Elasticsearch

- API REST : `http://host:port/[index]/[type]/[_action/id]`
Méthodes HTTP : **GET**, **POST**, **PUT**, **DELETE**
- Documents
 - `curl -XPUT http://localhost:9200/twitter/tweet/1`
 - `curl -XGET http://localhost:9200/twitter/tweet/1`
 - `curl -XDELETE http://localhost:9200/twitter/tweet/1`



Interagir avec Elasticsearch

- API REST : `http://host:port/[index]/[type]/[_action/id]`
Méthodes HTTP : **GET**, **POST**, **PUT**, **DELETE**
- Documents
 - `curl -XPUT http://localhost:9200/twitter/tweet/1`
 - `curl -XGET http://localhost:9200/twitter/tweet/1`
 - `curl -XDELETE http://localhost:9200/twitter/tweet/1`
- Recherche
 - `curl -XPOST http://localhost:9200/twitter/tweet/_search`



Interagir avec Elasticsearch

- API REST : `http://host:port/[index]/[type]/[_action/id]`
Méthodes HTTP : **GET**, **POST**, **PUT**, **DELETE**
- Documents
 - `curl -XPUT http://localhost:9200/twitter/tweet/1`
 - `curl -XGET http://localhost:9200/twitter/tweet/1`
 - `curl -XDELETE http://localhost:9200/twitter/tweet/1`
- Recherche
 - `curl -XPOST http://localhost:9200/twitter/tweet/_search`
 - `curl -XPOST http://localhost:9200/twitter/_search`



Interagir avec Elasticsearch

- API REST : `http://host:port/[index]/[type]/[_action/id]`
Méthodes HTTP : **GET**, **POST**, **PUT**, **DELETE**
- Documents
 - `curl -XPUT http://localhost:9200/twitter/tweet/1`
 - `curl -XGET http://localhost:9200/twitter/tweet/1`
 - `curl -XDELETE http://localhost:9200/twitter/tweet/1`
- Recherche
 - `curl -XPOST http://localhost:9200/twitter/tweet/_search`
 - `curl -XPOST http://localhost:9200/twitter/_search`
 - `curl -XPOST http://localhost:9200/_search`



Interagir avec Elasticsearch

- API REST : `http://host:port/[index]/[type]/[_action/id]`
Méthodes HTTP : **GET**, **POST**, **PUT**, **DELETE**
- Documents
 - `curl -XPUT http://localhost:9200/twitter/tweet/1`
 - `curl -XGET http://localhost:9200/twitter/tweet/1`
 - `curl -XDELETE http://localhost:9200/twitter/tweet/1`
- Recherche
 - `curl -XPOST http://localhost:9200/twitter/tweet/_search`
 - `curl -XPOST http://localhost:9200/twitter/_search`
 - `curl -XPOST http://localhost:9200/_search`
- Meta-données
 - `curl -XGET http://localhost:9200/twitter/_status`



Interagir avec Elasticsearch

- API REST : `http://host:port/[index]/[type]/[_action/id]`
Méthodes HTTP : **GET**, **POST**, **PUT**, **DELETE**
- Documents
 - `curl -XPUT http://localhost:9200/twitter/tweet/1`
 - `curl -XGET http://localhost:9200/twitter/tweet/1`
 - `curl -XDELETE http://localhost:9200/twitter/tweet/1`
- Recherche
 - `curl -XPOST http://localhost:9200/twitter/tweet/_search`
 - `curl -XPOST http://localhost:9200/twitter/_search`
 - `curl -XPOST http://localhost:9200/_search`
- Meta-données
 - `curl -XGET http://localhost:9200/twitter/_status`
 - `curl -XPOST http://localhost:9200/_shutdown`



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\nAt this time, enjoying NoSQL world : CouchDB, ElasticSearch.\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" }
      }
    ]
  }
}'
```





Injecter les données


Et au milieu coule une rivière




La collecte




La collecte




La collecte




La collecte




La collecte



La collecte



La collecte




La collecte



La collecte



La collecte



La collecte



La collecte



Quelques Rivers...



Quelques Rivers...

- CouchDB River



Quelques Rivers...

- CouchDB River
- CouchBase River




Quelques Rivers...

- CouchDB River
- CouchBase River
- MongoDB River



Quelques Rivers...

- CouchDB River
- CouchBase River
- MongoDB River
- JDBC River



Quelques Rivers...

- CouchDB River
- CouchBase River
- MongoDB River
- JDBC River
- Wikipedia River



Quelques Rivers...

- CouchDB River
- CouchBase River
- MongoDB River
- JDBC River
- Wikipedia River
- Twitter River

twitter



Quelques Rivers...

- CouchDB River
- CouchBase River
- MongoDB River
- JDBC River
- Wikipedia River
- Twitter River
- RabbitMQ River



Quelques Rivers...

- CouchDB River
- CouchBase River
- MongoDB River
- JDBC River
- Wikipedia River
- Twitter River
- RabbitMQ River
- ActiveMQ River



Quelques Rivers...

- CouchDB River
- CouchBase River
- MongoDB River
- JDBC River
- Wikipedia River
- Twitter River
- RabbitMQ River
- ActiveMQ River
- RSS River



Quelques Rivers...

- CouchDB River
- CouchBase River
- MongoDB River
- JDBC River
- Wikipedia River
- Twitter River
- RabbitMQ River
- ActiveMQ River
- RSS River
- LDAP River



Quelques Rivers...

- CouchDB River
- CouchBase River
- MongoDB River
- JDBC River
- Wikipedia River
- Twitter River
- RabbitMQ River
- ActiveMQ River
- RSS River
- LDAP River
- FS River



Quelques Rivers...

- CouchDB River
- CouchBase River
- MongoDB River
- JDBC River
- Wikipedia River
- Twitter River
- RabbitMQ River
- ActiveMQ River
- RSS River
- LDAP River
- FS River
- Dropbox River



Quelques Rivers...

- CouchDB River
- CouchBase River
- MongoDB River
- JDBC River
- Wikipedia River
- Twitter River
- RabbitMQ River
- ActiveMQ River
- RSS River
- LDAP River
- FS River
- Dropbox River
- Dick Rivers



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
0	dadoonet	2012-04-18	1
1	talk	2012-04-18	5
2	elasticsearch	2012-04-18	2
3	dadoonet	2012-04-18	2
4	talk	2012-04-18	6
5	elasticsearch	2012-04-19	3
6	dadoonet	2012-04-19	3
7	talk	2012-04-19	7
8	elasticsearch	2012-04-20	4



Term Facet


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

Username	Count
dadoonet	3
talk	3
elasticsearch	3



Term Facet


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



Term Facet


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

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



Term Facet

	Username	Date	Hashtag
0	dadoonet	2012-0	
1	talk	2012-0	
2	elasticsearch	2012-0	
3	dadoonet	2012-0	
4	talk	2012-0	
5	elasticsearch	2012-0	
6	dadoonet	2012-0	
7	talk	2012-0	
8	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
n	2012-04-18	2
	2012-04-18	2
	2012-04-18	6
n	2012-04-19	3
	2012-04-19	3
	2012-04-19	7
n	2012-04-20	4




Date Histogram Facet

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

Hashtag
1
5
2
2
6
3
3
7
4

```
"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  
    } ] } }
```



Site marchand



Site marchand



Site marchand



Site marchand



Site marchand

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 Infinite
- ☐ 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)
- ☐ XXXL (~ 50 000)

Ranges

Term


Term

Ranges



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	
Web Title	
Url	
Category	Select v
Web Description	
Keywords	
Contact Name	
Contact Email	
Featured Links 	Select v
Cool Links 	Select v
Bold Links	Select v
Icon	☐  ☐  ☐  ☐  ☐  ☐ 
Rating Average 	Select v
Number of Votes	between and
Total Hits	between and
Hits Today	between and
IP Address	
Submission Software Name	



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

Total Hits

Hits Today

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

Total Hits

Hits Today

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



Démonstration

Avez-vous fait du bruit ?




Architecture



Architecture




Architecture



Architecture



Architecture



```
$ curl -XPUT localhost:9200/_river/twitter/_meta -d '{
  "type" : "twitter",
  "twitter" : {
    "user" : "twitter_user",
    "password" : "twitter_password",
    "filter" : { "tracks" : ["elasticsearch"] }
  }
}
```



Architecture



```
$ curl -XPUT localhost:9200/_river/twitter/_meta -d '{
  "type" : "twitter",
  "twitter" : {
    "user" : "twitter_user",
    "password" : "twitter_password",
    "filter" : { "tracks" : ["elasticsearch"] }
  }
}
```



Démonstration

<http://onemilliontweetmap.com/>

<http://www.scrutmydocs.org>

The logo for 'maptimize' features a blue location pin icon above the letter 'a', followed by the word 'maptimize' in a lowercase, sans-serif font.



Démonstration

<http://onemilliontweetmap.com/>

<http://www.scrutmydocs.org>



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  
  }  
'
```




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



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




Cherchons !




Cherchons !




Cherchons !




Cherchons !




Cherchons !




Cherchons encore !




Cherchons encore !




Cherchons encore !




Cherchons encore !




Cherchons encore !




Cherchons encore !



Cherchons encore !



Cherchons encore !



La percolation

Ou la recherche inversée



Usage courant d'un moteur de recherche

- J'indexe un document
- Je cherche de temps en temps si un document m'intéresse
- **Avec de la chance**, il sera bien placé au niveau pertinence dans les résultats. Sinon, il passe **inaperçu** !



La recherche inversée

- Enregistrer ses critères de recherche
- A chaque document indexé, on récupère la liste des recherches qui correspondent
- On a un « **listener** » sur le moteur d'indexation : le **percolator**



Usage du percolator

```
$ curl -XPOST localhost:9200/_percolator/twitter/dadoonet -d '{
  "query" : { "term" : { "user.screen_name" : "dadoonet" } }
}'

$ curl -XPOST localhost:9200/_percolator/twitter/elasticsearch -d '{
  "query" : { "match" : { "hashtag.text" : "elasticsearch" } }
}'

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



Usage du percolator



Usage du percolator

```
$ curl -XPUT localhost:9200/twitter/tweet/1&percolate=* -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
\rDeeJay 4 times a year, just for fun !" }
}'
```



Usage du percolator

```
$ curl -XPUT localhost:9200/twitter/tweet/1&percolate=* -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
\r\nDeeJay 4 times a year, just for fun !" }
}'
```




```
{
  "ok": true,
  "_index": "twitter",
  "_type": "tweet",
  "_id": "1",
  "matches": [
    "dadoonet",
    "elasticsearch"
  ]
}
```



Tout doit être indexé ?

Analyse et mapping



The lazy dog...



The lazy dog...

The quick brown fox
jumped over the lazy Dog



The lazy dog...

The quick brown fox
jumped over the lazy Dog

The quick brown fox
jumped over the lazy dog



Analyseur standard

```
$ curl -XPOST 'localhost:9200/test/_analyze?analyzer=standard&pretty=1' -d  
'The quick brown fox jumped over the lazy Dog'
```



Analyseur standard

```
$ curl -XPOST 'localhost:9200/test/_analyze?analyzer=standard&pretty=1' -d  
'The quick brown fox jumped over the lazy Dog'
```



```
{  
  "tokens" : [ {  
    "token" : "quick",  
    "start_offset": 4, "end_offset": 9, "type": "<ALPHANUM>", "position": 2  
  }, {  
    "token" : "brown",  
    "start_offset": 10, "end_offset": 15, "type": "<ALPHANUM>", "position": 3  
  }, {  
    "token" : "fox",  
    "start_offset": 16, "end_offset": 19, "type": "<ALPHANUM>", "position": 4  
  }, {  
    "token": "jumped",  
    "start_offset": 20, "end_offset": 26, "type": "<ALPHANUM>", "position": 5  
  }, {  
    "token": "over",  
    "start_offset": 27, "end_offset": 31, "type": "<ALPHANUM>", "position": 6  
  }, {  
    "token" : "lazy",  
    "start_offset": 36, "end_offset": 40, "type": "<ALPHANUM>", "position": 8  
  }, {  
    "token" : "dog",  
    "start_offset": 41, "end_offset": 44, "type": "<ALPHANUM>", "position": 9  
  } ] }
```



Analysateur whitespace

```
$ curl -XPOST 'localhost:9200/test/_analyze?analyzer=whitespace&pretty=1' -d  
'The quick brown fox jumped over the lazy Dog'
```



Analysseur whitespace

```
$ curl -XPOST 'localhost:9200/test/_analyze?analyzer=whitespace&pretty=1' -d  
'The quick brown fox jumped over the lazy Dog'
```



```
{  
  "tokens" : [ {  
    "token" : "The", ...  
  }, {  
    "token" : "quick", ...  
  }, {  
    "token" : "brown", ...  
  }, {  
    "token" : "fox", ...  
  }, {  
    "token" : "jumped", ...  
  }, {  
    "token" : "over", ...  
  }, {  
    "token" : "the", ...  
  }, {  
    "token" : "lazy", ...  
  }, {  
    "token" : "Dog", ...  
  } ] }
```



Un analyseur

Un ensemble
de tokenizers et
de filtres



Un tokenizer

- Découpe une chaîne en « mots » et transforme :
- whitespace tokenizer :

"the dog!" -> "the", "dog!"

- standard tokenizer :

"the dog!" -> "the", "dog"



Un filtre

- Supprime ou transforme un token :

- asciifolding filter :

éléphant -> elephant

- stemmer filter (french) :

elephants -> "eleph"

cheval -> "cheval"

chevaux -> "cheval"

- phonetic (plugin) :

quick -> "Q200"

quik -> "Q200"



Analyzer

```
"analysis":{
  "analyzer":{
    "francais":{
      "type":"custom",
      "tokenizer":"standard",
      "filter":["lowercase", "stop_francais", "fr_stemmer", "asciifolding", "elision"]
    }
  },
  "filter":{
    "stop_francais":{
      "type":"stop",
      "stopwords":["_french_", "twitter"]
    },
    "fr_stemmer" : {
      "type" : "stemmer",
      "name" : "french"
    },
    "elision" : {
      "type" : "elision",
      "articles" : ["l", "m", "t", "qu", "n", "s", "j", "d"]
    }
  }
}
```



Mapping

```
"type1" : {  
  "properties" : {  
    "text1" : { "type" : "string", "analyzer" : "simple" },  
    "text2" : { "type" : "string", "index_analyzer" : "simple",  
               "search_analyzer" : "standard"  
    },  
    "text3" : {  
      "type" : "multi_field",  
      "fields" : {  
        "text3" : {  
          "type" : "string",  
          "analyzer" : "standard"  
        },  
        "ngram" : {  
          "type" : "string",  
          "analyzer" : "ngram"  
        },  
        "soundex" : {  
          "type" : "string",  
          "analyzer" : "soundex"  
        }  
      }  
    }  
  }  
}
```



Les types

- string
- integer / long
- float / double
- boolean
- null
- array
- objects
- multi_field
- ip
- geo_point
- geo_shape
- binary
- attachment (plugin)



Champs spéciaux

- `_all` (et `include_in_all`)
- `_source`
- `_ttl`
- `parent / child`
- `nested`



Autres fonctionnalités

- highlighting



The screenshot shows a search interface with a search bar at the top containing the text 'apache'. Below the search bar, it indicates '3 documents found in 78 milliseconds'. The search results are displayed in a list format. The first result is 'NOTICE.txt' by The Apache Software Foundation (http://www.apache.org/). The second result is 'Apache.pdf', which includes a snippet of text: 'Configurer Apache mars 2007 i 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'. The third result is 'LICENSE.txt' Apache License Version 2.0, January 2004 http://www. The word 'Apache' is highlighted in blue in the search bar and in the search results.

Search bar: apache

3 documents found in 78 milliseconds

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

[Apache.pdf](#)
Configurer Apache mars 2007 i 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



Autres fonctionnalités

- highlighting
- scoring
- sort
- explain
- multi get / multi search
- bulk



Démonstrations

CURL est ton ami !

JAVA est aussi ton ami !



La communauté



elasticsearch.



mozilla
FOUNDATION

Revinat
The New Standard for Guest Satisfaction



foursquare



presspeople

KLOUT



**stack
overflow**

CabinetOffice

atlas
The global video & audio index



LiveChat



b:
bazaarvoice

Booking.com

infochimps



**JBoss
Community**

enonic



Udini
by ProQuest

InSTEDD

ataxo

sonian
Cloud Powered Archiving and Search

**need
a property.com**


COGENTA

Quora

D P L A

~80 contributeurs directs au projet (+ de 4000 watchers et + de 840 forks)

La communauté




~250 inscrits sur la mailing list, 70 messages / mois, ~400 followers, ~200 sur meetup

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>

