



# Elasticsearch

Using the completion suggester

Alexander Reelsen

@spinscale

[alexander.reelsen@elasticsearch.com](mailto:alexander.reelsen@elasticsearch.com)

# Suggestion - Introduction



The screenshot shows the GitHub interface for the `elasticsearch/elasticsearch` repository. A search dropdown is open, showing results for the query 'debian'. The results include several issues related to Debian packaging, such as naming conventions, init-scripts, and repository locations. The main content area shows a list of issues, with the top one being a bug report about a startup exception.

**GitHub** This repository ▾ `debian` [Sign up](#) [Sign in](#)

PUBLIC `elasticsearch` [Star](#) 4,683 [Fork](#) 1,097

[Browse Issues](#) [New Issue](#)

Everyone's Issues

**Labels**

- Lucene 4.5 Upgrade
- breaking
- bug
- enhancement
- feature
- non-issue

Search elasticsearch/elasticsearch for 'debian'

Search GitHub for 'debian'

Issue Title	Count	Opened by	Time
elasticsearch/elasticsearch#1726 <code>debian</code> package violates naming convention	1	s1monw	14 hours ago
elasticsearch/elasticsearch#3571 <code>debian</code> package init-script: start-stop-daemon ne	1		
elasticsearch/elasticsearch#1681 <code>Debian</code> pkg	1		
elasticsearch/elasticsearch#3286 There is no official <code>debian/ubuntu</code> repository	1		
elasticsearch/elasticsearch#3500 Elasticsearch should include <code>debian's</code> standard j	1		
elasticsearch/elasticsearch#1526 Moving <code>debian</code> package to maven	1		

Forms #3702

reproducible #3701

**NoShardAvailableActionException in ES 0.90.3 on startup** #3700  
Opened by richardwilly98 a day ago

**Feature Request: Don't reindex the document when updating non-indexed fields** #3696  
Opened by dorian 2 days ago 4 comments

# Suggestion - Introduction

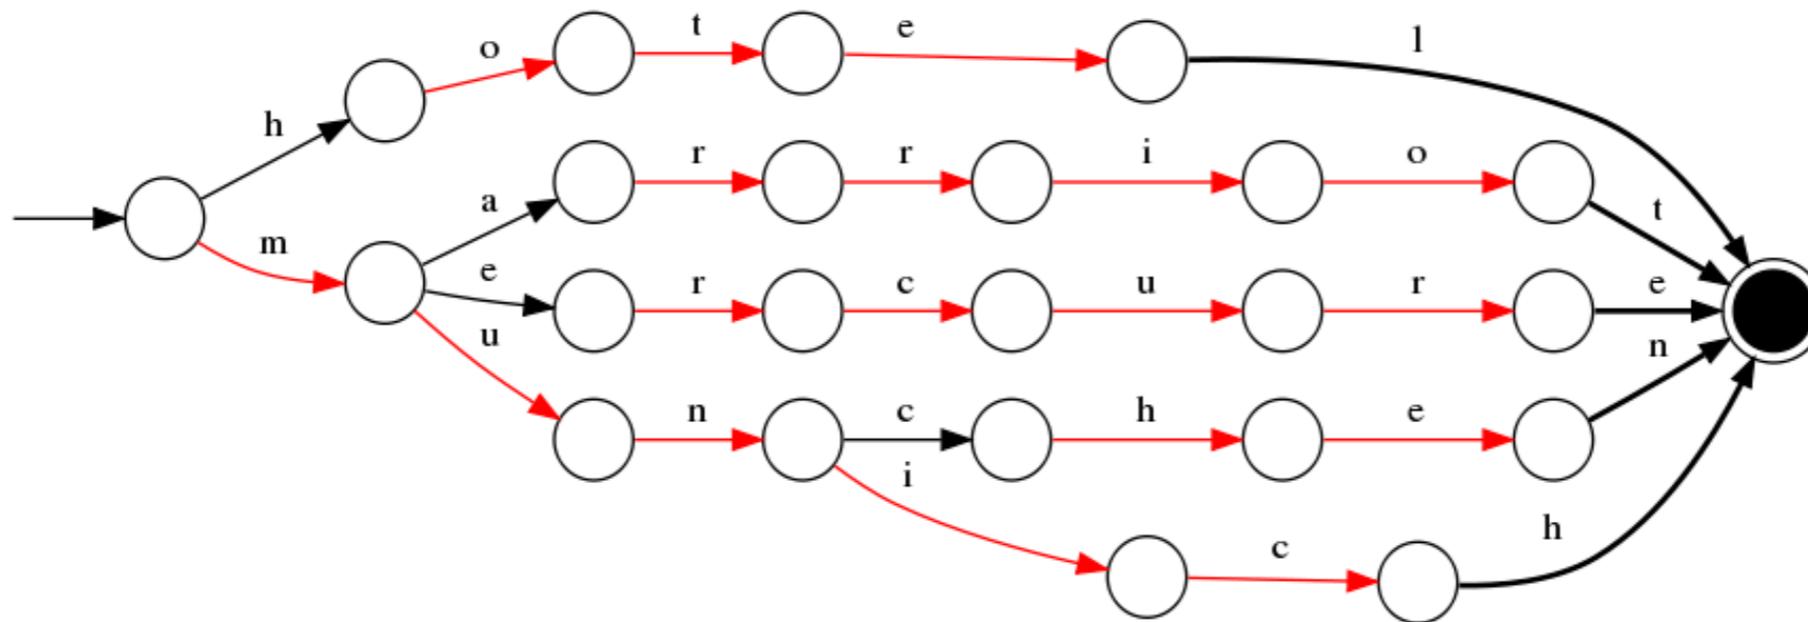


- Different possibilities
  - Using queries (EdgeNGram, NGram, match prefix)
    - Returns the whole document
    - As fast as a query
- Term suggester
- Phrase suggester
  - did-you-mean

# Suggestion - Completion



- Search-as-you-type, **prefix** suggester
- Creates an in-memory structure (speed!)
- FST is written on disk on indexing
- Responses < 1ms



# Indexing



```
curl -X PUT localhost:9200/products/product/1 -d
{
  "name" : "Apple Iphone 5",
  "name_suggest" : {
    "input" : [ "Apple Iphone 5", "Iphone 5" ],
    "output" : "Apple Iphone 5",
    "weight" : 100,
    "payload" : 1
  }
}
```

# Querying - Request



```
curl -X POST localhost:9200/products/_suggest -d '{
  "suggestions" : {
    "text" : "i",
    "completion" : {
      "field" : "name_suggest"
    }
  }
}'
```

# Querying - Response



```
{
  ...
  "options" : [ {
    "text" : "Apple Iphone 5",
    "score" : 100.0
    "payload" : 1
  }, {
    "text" : "iphone 4 GS",
    "score" : 1.0
  }, {
    "text" : "iPhone 6",
    "score" : 1.0
  }, {
    "text" : "iPhone case",
    "score" : 1.0
  } ]
} ]
}
```

# Suggestion - In operation



- Completion suggester needs memory to be fast
- Monitoring per node/index/field
- Keep your payloads small
- Improve relevancy by constant monitoring
  
- Optional: Maintain an own index for suggestions

# Suggestion - Future



- Done: Fuzzy
- Highlighting
- Improved handling of stopwords
- Geo-Suggestions
- More than prefix-only suggestions



# Thanks for listening!

- <http://elasticsearch.org>
- <http://elasticsearch.com>
- <https://groups.google.com/forum/#!forum/elasticsearch>
- <http://www.elasticsearch.org/blog/you-complete-me/>
- <http://www.elasticsearch.org/guide/en/elasticsearch/reference/current/search-suggesters.html>
- <http://www.elasticsearch.org/guide/en/elasticsearch/reference/current/search-suggesters-completion.html>