

# *Intro to Elasticsearch and Kibana*



# *Table of contents*

PARTIE 1 INTRODUCTION

---

PARTIE 2 QUERIES

---

PARTIE 3 AGGREGATIONS

---

PARTIE 4 MAPPING

# Elasticsearch and Kibana



**Elasticsearch**

*Elasticsearch is a distributed, RESTful search and analytics engine capable of addressing a growing number of use cases. As the heart of the Elastic Stack, it centrally stores your data for lightning fast search, fine-tuned relevancy, and powerful analytics that scale with ease.*



**Kibana**

*Kibana is a free and open user interface that lets you visualize your Elasticsearch data and navigate the Elastic Stack. Do anything from tracking query load to understanding the way requests flow through your apps.*

# Downloading Elasticsearch and Kibana (macOS/Linux and Windows)

Step 1: The following links will take you to the download pages for Elasticsearch and Kibana.

- <https://www.elastic.co/downloads/elasticsearch>
- <https://www.elastic.co/downloads/kibana>


Step 2: Relocate downloaded Elasticsearch and unzip Elasticsearch.

Step 3: Relocate downloaded Kibana and unzip Kibana.

Step 4: Start the Elasticsearch/Kibana server and ensure that everything is working properly

- Go to the command prompt (Terminal).
- Go to the installation folder where you have unzipped the file.
- Run “bin\elasticsearch.bat”(“bin/elasticsearch” On Linux/macOs).
- Elasticsearch has now started.
- Run “\bin\kibana.bat”(“bin/kibana” On Linux/macOs).
- Kibana has now started.
- Open your browser and go to “[https://localhost:5601/app/dev\\_tools#console](https://localhost:5601/app/dev_tools#console)”

# Who uses Elasticsearch?



# Create an index

Syntax:

**PUT Name-of-the-Index**

Example:

**PUT favorite\_candy**

```
200 - OK 211 ms
1 {
2   "acknowledged" : true,
3   "shards_acknowledged" : true,
4   "index" : "favorite_candy"
5 }
6
```

# Index a document

When indexing a document, both HTTP verbs POST or PUT can be used.

- Use POST when you want Elasticsearch to autogenerate an id for your document.
- Use PUT when you want to assign a specific id to your document(i.e. if your document has a natural identifier - purchase order number, patient id, & etc). For more detailed explanation, check out this documentation from Elastic

# Index a document using POST

Syntax:

```
POST Name-of-the-Index/_doc
{
  "field": "value"
}
```

Example:

```
POST favorite_candy/_doc
{
  "first_name": "Lisa",
  "candy": "Sour Skittles"
}
```

```
201 - Created 70 ms
1 {
2   "_index" : "favorite_candy",
3   "_type" : "_doc",
4   "_id" : "HIbNUnYBYpCfo_1CFga-",
5   "_version" : 1,
6   "result" : "created",
7   "_shards" : {
8     "total" : 2,
9     "successful" : 1,
10    "failed" : 0
11  },
12  "_seq_no" : 0,
13  "_primary_term" : 1
14 }
15
```



# Index a document using PUT

Syntax:

```
PUT Name-of-the-Index/_doc/id-you-want-to-assign-to-this-document
{
  "field": "value"
}
```

Example:

```
PUT favorite_candy/_doc/1
{
  "first_name": "John",
  "candy": "Starburst"
}
```

# Create Endpoint

When you index a document using an id that already exists, the existing document is overwritten by the new document. If you do not want a existing document to be overwritten, you can use the `_create` endpoint! With the `_create` Endpoint, no indexing will occur and you will get a 409 error message

Syntax:

```
PUT Name-of-the-Index/_create/id
{
  "field": "value"
}
```

Example:

```
PUT favorite_candy/_create/1
{
  "first_name": "Finn",
  "candy": "Jolly Ranchers"
}
```

```
409 - Conflict 22 ms
1 - {
2 -   "error" : {
3 -     "root_cause" : [
4 -       {
5 -         "type" : "version_conflict_engine_exception"
6 -       },
7 -       "reason" : "[1]: version conflict, document
8         already exists (current version [2])",
9         "index_uuid" : "sVls0NrZQS6ZETigG2TM6g",
10        "shard" : "0",
11        "index" : "favorite_candy"
12      },
13      "type" : "version_conflict_engine_exception",
14      "reason" : "[1]: version conflict, document
15        already exists (current version [2])",
16      "index_uuid" : "sVls0NrZQS6ZETigG2TM6g",
17      "shard" : "0",
18      "index" : "favorite_candy"
19    },
20    "status" : 409
}
```

# Read a document

Syntax:

**GET Name-of-the-Index/\_doc/id**

Example:

**GET favorite\_candy/\_doc/1**

```
200 - OK 9 ms
1 - {
2   "_index" : "favorite_candy",
3   "_type" : "_doc",
4   "_id" : "1",
5   "_version" : 1,
6   "_seq_no" : 1,
7   "_primary_term" : 1,
8   "found" : true,
9   "_source" : {
10    | "first_name" : "John",
11    | "candy" : "Starburst"
12  ^ }
13 ^ }
14
```

# Update a document

Syntax:

```
POST Name-of-the-Index/_update/id
{
  "doc": {
    "field1": "value",
    "field2": "value",
  }
}
```

Example:

```
POST favorite_candy/_update/1
{
  "doc": {
    "candy": "M&M's"
  }
}
```

```
200 - OK 26 ms
1 - {
2   "_index" : "favorite_candy",
3   "_type" : "_doc",
4   "_id" : "1",
5   "_version" : 3,
6   "result" : "updated",
7 -  "_shards" : {
8     "total" : 2,
9     "successful" : 1,
10    "failed" : 0
11 - },
12   "_seq_no" : 5,
13   "_primary_term" : 1
14 - }
15
```

# Delete a document

Syntax:

```
DELETE Name-of-the-Index/_doc/id
```

Example:

```
DELETE favorite_candy/_doc/1
```

```
200 - OK 26 ms
1- {
2  "_index" : "favorite_candy",
3  "_type" : "_doc",
4  "_id" : "1",
5  "_version" : 4,
6  "result" : "deleted",
7-  "_shards" : {
8    "total" : 2,
9    "successful" : 1,
10   "failed" : 0
11- },
12  "_seq_no" : 6,
13  "_primary_term" : 1
14- }
15
```

# Searching for search terms

The **match query** is a standard query for performing a full text search. This **query** retrieves documents that contain the search terms. It uses "OR" logic by default, meaning that it will retrieve documents that contain any one of the search terms. The order and the proximity in which the search terms are found(i.e. phrases) are not taken into account

Syntax:

```
GET Enter_name_of_index_here/_search
{
  "query": {
    "match": {
      "Specify the field you want to search": {
        "query": "Enter search terms"
      }
    }
  }
}
```

# Searching for search terms

What happens when you use the **match query** to search for phrases?

Let's search for articles about Ed Sheeran's song "Shape of you" using the **match query**.

Example:

```
GET news_headlines/_search
```

```
{
  "query": {
    "match": {
      "headline": {
        "query": "Shape of you"
      }
    }
  }
}
```

```
    "hits" : {
      "total" : {
        "value" : 10000,
        "relation" : "gte"
      }
    },
    "max_score" : 12.274778,
    "hits" : [
      {
        "_index" : "news_headlines",
        "_type" : "_doc",
        "_id" : "u9g9S3cBwsjPafpA2HGP",
        "_score" : 12.274778,
        "_source" : {
          "date" : "2012-08-30",
          "short_description" : "Get stronger. Practice wall squats (with back against the wall, lower your body until knees are at 90 degrees; hold for 30",
          "@timestamp" : "2012-08-30T00:00:00.000-06:00",
          "link" : "https://www.huffingtonpost.com/entry/fitness-test-women_us_5b9c2c91e4b03a1dcc7cda8e",
          "category" : "WELLNESS",
          "headline" : "Fitness Test: Are You In Shape?",
          "authors" : ""
        }
      }
    ]
  },
}
```

# Searching for a phrase

If the order and the proximity in which the search terms are found(i.e. phrases) are important in determining the relevance of your search, you use the `match_phrase` query.

Syntax:

```
GET Enter_name_of_index_here/_search
{
  "query": {
    "match_phrase": {
      "Specify the field you want to search": {
        "query": "Enter search terms"
      }
    }
  }
}
```



# Searching for a phrase

Example:

```
GET news_headlines/_search
{
  "query": {
    "match_phrase": {
      "headline": {
        "query": "Shape of You"
      }
    }
  }
}
```

When the `match_phrase` parameter is used, all hits must meet the following criteria:

1. the search terms "Shape", "of", and "you" must appear in the field headline .
2. the terms must appear in that order.
3. the terms must appear next to each other.

```
10  "hits" : {
11  |   "total" : {
12  |     "value" : 3,
13  |     "relation" : "eq"
14  |   },
15  |   "max_score" : 12.074881,
16  |   "hits" : [
17  |     {
18  |       "_index" : "news_headlines",
19  |       "_type" : "_doc",
20  |       "_id" : "lMs7S3cBVGnaeHqj6RUZ",
21  |       "_score" : 12.074881,
22  |       "_source" : {
23  |         "date" : "2017-03-20",
24  |         "short_description" : "Puerto Rico's
25  |           Zion & Lennox are behind the new
26  |           version.",
27  |         "@timestamp" : "2017-03-20T00:00:00.000
28  |           -06:00",
29  |         "link" : "https://www.huffingtonpost
30  |           .com/entry/ed-sheerans-zion-lennox-
31  |           -shape-of-you-latin-remix_us_58d03b09e4b0be71dcf72c6f",
32  |         "category" : "LATINO VOICES",
33  |         "headline" : "Ed Sheeran's 'Shape Of
34  |           You' Gets An Unexpected Latin Remix",
35  |         "authors" : "Carolina Moreno"
36  |       }
37  |     },
38  |   ]
39  }
```

# Running a match query against multiple fields

To accommodate these contexts, you can write a `multi_match query`, which searches for terms in multiple fields.

The `multi_match query` runs a `match query` on multiple fields and calculates a score for each field. Then, it assigns the highest score among the fields to the document.

This score will determine the ranking of the document within the search results.

Syntax:

```
GET Enter_the_name_of_the_index_here/_search
{
  "query": {
    "multi_match": {
      "query": "Enter search terms here",
      "fields": [
        "List the field you want to search over",
        "List the field you want to search over",
        "List the field you want to search over"
      ]
    }
  }
}
```

# Running a match query against multiple fields

Example:

```
GET news_headlines/_search
{
  "query": {
    "multi_match": {
      "query": "Michelle Obama",
      "fields": [
        "headline",
        "short_description",
        "authors"
      ]
    }
  }
}
```

```
10- "hits" : {
11-   "total" : {
12-     "value" : 3044,
13-     "relation" : "eq"
14-   },
15-   "max_score" : 16.937054,
16-   "hits" : [
17-     {
18-       "_index" : "news_headlines",
19-       "_type" : "_doc",
20-       "_id" : "idc8S3cBwsjPafpAJViR",
21-       "_score" : 16.937054,
22-       "_source" : {
23-         "date" : "2016-07-26",
24-         "short_description" : "But Michelle
25-         Obama stole the show.",
26-         "@timestamp" : "2016-07-26T00:00:00
27-         .000-06:00",
28-         "link" : "https://www.huffingtonpost
29-         .com/entry/tuesdays-morning-email-
30-         sanders-supporters-make-dnc-feel-
31-         the-bern_us_579749ace4b02d5d5ed2bd1d",
32-         "category" : "POLITICS",
33-         "headline" : "Tuesday's Morning Email:
34-         Sanders Supporters Make DNC Feel The
35-         Bern",
36-         "authors" : "Lauren Weber"
37-       }
38-     }
39-   ]
40- }
41- }
```

But this  
headline  
is about  
Bernie  
Sanders!

# Per-field boosting

To improve the precision of your search, you can designate one field to carry more weight than the others.

This can be done by boosting the score of the field headline(**per-field boosting**). This is notated by adding a carat(^) symbol and number 2 to the desired field as shown below.

Syntax:

```
GET Enter_the_name_of_the_index_here/_search
{
  "query": {
    "multi_match": {
      "query": "Enter search terms",
      "fields": [
        "List field you want to boost^2",
        "List field you want to search over",
        "List field you want to search over"
      ]
    }
  }
}
```

# Per-field boosting

Example:

```
GET news_headlines/_search
{
  "query": {
    "multi_match": {
      "query": "Michelle Obama",
      "fields": [
        "headline^2",
        "short_description",
        "authors"
      ]
    }
  }
}
```

```
10 ~ "hits" : {
11 ~   "total" : {
12 ~     "value" : 5128,
13 ~     "relation" : "eq"
14 ~   },
15 ~   "max_score" : 26.837097,
16 ~   "hits" : [
17 ~     {
18 ~       "_index" : "news_headlines",
19 ~       "_type" : "_doc",
20 ~       "_id" :
21 ~         "5dc8S3cBwsjPafpAtLG7",
22 ~       "_score" : 26.837097,
23 ~       "_source" : {
24 ~         "date" : "2015-03-25",
25 ~         "short_description" : "",
26 ~         "@timestamp" : "2015-03-
27 ~           -25T00:00:00.000-06:00",
28 ~         "link" : "https://www
29 ~           .huffingtonpost.com/entry
30 ~           /michelle-obama
31 ~           -jeopardy_n_6939122.html"
32 ~       },
33 ~       "category" : "POLITICS",
34 ~       "headline" : "Michelle
35 ~         Obama Appears On
36 ~         'Jeopardy!'",
37 ~       "authors" : "Amber
38 ~         Ferguson"
39 ~     }
40 ~   ]
41 ~ }
```

# What happens when you use the `multi_match` query to search for a phrase?

Example:

```
GET news_headlines/_search
{
  "query": {
    "multi_match": {
      "query": "party planning",
      "fields": [
        "headline^2",
        "short_description"
      ]
    }
  }
}
```

```
1- {
2   "took" : 4,
3   "timed_out" : false,
4   "_shards" : {
5     "total" : 1,
6     "successful" : 1,
7     "skipped" : 0,
8     "failed" : 0
9   }
10  "hits" : {
11    "total" : {
12      "value" : 2846,
13      "relation" : "eq"
14    },
15    "max_score" : 28.425034,
16    "hits" : [
79     {
80       "_index" : "news_headlines",
81       "_type" : "_doc",
82       "_id" : "s8s7S3cBVGnaeHqj4g_g",
83       "_score" : 20.281507,
84       "_source" : {
85         "date" : "2017-04-07",
86         "short_description" : "Democratic
87         officials are expected to join
88         Sanders on the road.",
89         "@timestamp" : "2017-04-07T00:00:00
90         .000-06:00",
91         "link" : "https://www.huffingtonpost
92         .com/entry/bernie-sanders-national
93         -tour_us_58e6f28be4b0cdad578e7a52",
94         "category" : "POLITICS",
95         "headline" : "Bernie Sanders And Tom
96         Perez Planning National Tour To
97         Boost Candidates, Grassroots Party
98         Activism",
99         "authors" : "Ryan Grim"
100      }
101    }
102  ]
103 }
```

Not quite the party planning we are looking for...

# Improving precision with phrase type match

You can improve the precision of a `multi_match` query by adding the `"type":"phrase"` to the query. The phrase type performs a `match_phrase` query on each field and calculates a score for each field. Then, it assigns the highest score among the fields to the document

Syntax:

```
GET Enter_the_name_of_the_index_here/_search
{
  "query": {
    "multi_match": {
      "query": "Enter search phrase",
      "fields": [
        "List field you want to boost^2",
        "List field you want to search over",
        "List field you want to search over"
      ],
      "type": "phrase"
    }
  }
}
```

# Improving precision with phrase type match

Example:

```
GET news_headlines/_search
{
  "query": {
    "multi_match": {
      "query": "party planning",
      "fields": [
        "headline^2",
        "short_description"
      ],
      "type": "phrase"
    }
  }
}
```

```
10- "hits" : {
11-   "total" : {
12-     "value" : 6,
13-     "relation" : "eq"
14-   },
15-   "max_score" : 28.425034,
16-   "hits" : [
17-     {
18-       "_index" : "news_headlines",
19-       "_type" : "_doc",
20-       "_id" : "cdg9S3cBwsjPafpAMROy",
21-       "_score" : 28.425034,
22-       "_source" : {
23-         "date" : "2013-12-31",
24-         "short_description" : "If you are
    opting out of the over-the-top party
    this year and entertaining friends
    and family at home, I've rounded up
    a few New Years Eve party-planning
    musts.",
25-         "@timestamp" : "2013-12-31T00:00:00
    .000-07:00",
26-         "link" : "https://www.huffingtonpost
    .com/entry/new-years-eve-party
    -planning_us_5b9dbe4fe4b03a1dcc8c53a
    2",
27-         "category" : "HOME & LIVING",
28-         "headline" : "New Years Eve Party
    Planning Ideas",
29-         "authors" : ""Lauren Nelson,
    Contributor
30- Founder, Lauren-Nelson.com""
31-       }
    }
```



# Bool Query

The bool query retrieves documents matching boolean combinations of other queries.

With the **bool query**, you can combine multiple queries into one request and further specify boolean clauses to narrow down your search results. There are four clauses to choose from:

1. *must*
2. *must\_not*
3. *should*
4. *filter*

You can build combinations of one or more of these clauses.

Each clause can contain one or multiple queries that specify the criteria of each clause.

These clauses are optional and can be mixed and matched to cater to your use case.

The order in which they appear does not matter either!

# Bool Query

*Syntax:*

```
GET name_of_index/_search
```

```
{
```

```
  "query": {
```

```
    "bool": {
```

```
      "must": [{"One or more queries can be specified here. A document MUST match all of these queries to be considered as a hit.}],
```

```
      "must_not": [{"A document must NOT match any of the queries specified here. If it does, it is excluded from the search results.}],
```

```
      "should": [{"A document does not have to match any queries specified here. However, if it does match, this document is given a higher score.}],
```

```
      "filter": [{"These filters(queries) place documents in either yes or no category. Ones that fall into the yes category are included in the hits. }]
```

```
    }
```

```
  }
```

```
}
```

# Aggregations Request

*Syntax:*

```
GET Enter_name_of_the_index_here/_search
{
  "aggs": {
    "Name your aggregations here": {
      "Specify the aggregation type here": {
        "field": "Name the field you want to aggregate on here"
      }
    }
  }
}
```

# Metric Aggregations

**Metric** aggregations are used to compute numeric values based on your dataset. It can be used to calculate the values of **sum**, **min**, **max**, **avg**, unique count(**cardinality**) and etc.

Compute the **sum** of all unit prices in the index

Syntax:

```
GET Enter_name_of_the_index_here/_search
{
  "aggs": {
    "Name your aggregations here": {
      "sum": {
        "field": "Name the field you want to aggregate on here"
      }
    }
  }
}
```

# Metric Aggregations

Example:

```
GET ecommerce_data/_search
{
  "aggs": {
    "sum_unit_price": {
      "sum": {
        "field": "UnitPrice"
      }
    }
  }
}
```

```
1 {
2   "took" : 33,
3   "timed_out" : false,
4   "_shards" : {
5     "total" : 1,
6     "successful" : 1,
7     "skipped" : 0,
8     "failed" : 0
9   },
10  "hits" : { },
179  "aggregations" : {
180    "sum_unit_price" : {
181      "value" : 1876200.97
182    }
183  }
184 }
185
```

# Metric Aggregations

**Stats** Aggregation: Compute the count, min, max, avg, sum in one go

Syntax:

```
GET Enter_name_of_the_index_here/_search
{
  "size": 0,
  "aggs": {
    "Name your aggregations here": {
      "stats": {
        "field": "Name the field you want to aggregate on here"
      }
    }
  }
}
```

# Metric Aggregations

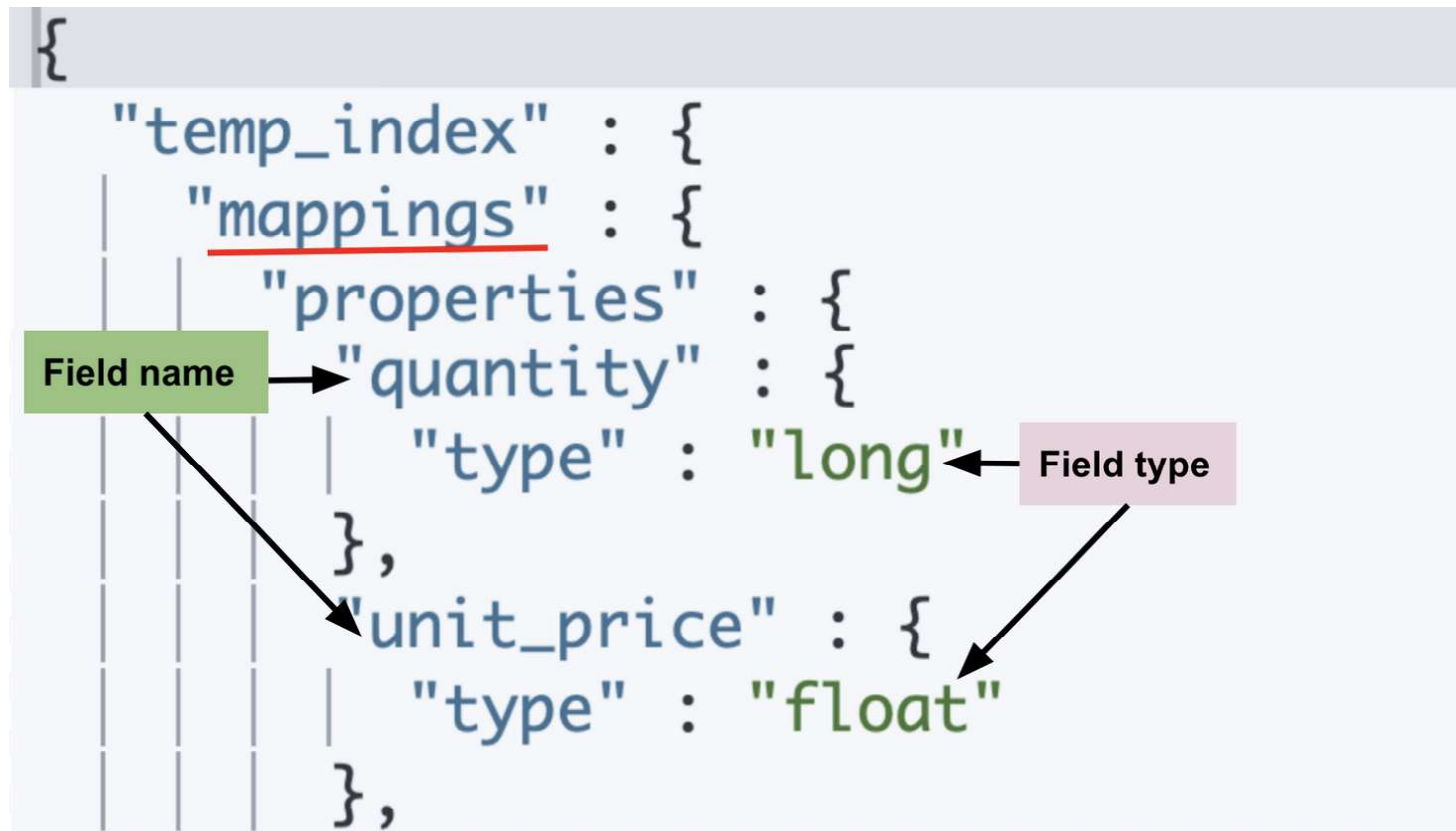
Example:

```
GET ecommerce_data/_search
{
  "size": 0,
  "aggs": {
    "all_stats_unit_price": {
      "stats": {
        "field": "UnitPrice"
      }
    }
  }
}
```

```
1 {
2   "took" : 48,
3   "timed_out" : false,
4   "_shards" : {
5     "total" : 1,
6     "successful" : 1,
7     "skipped" : 0,
8     "failed" : 0
9   },
10  "hits" : {
11    "total" : {
12      "value" : 10000,
13      "relation" : "gte"
14    },
15    "max_score" : null,
16    "hits" : [ ]
17  },
18  "aggregations" : {
19    "all_stats_unit_price" : {
20      "count" : 426841,
21      "min" : 1.01,
22      "max" : 498.79,
23      "avg" : 4.39555002916777,
24      "sum" : 1876200.97
25    }
26  }
27 }
```

# Mapping Explained

Mapping determines how a document and its fields are indexed and stored by defining the type of each field.





# View the Mapping

Syntax:

**GET Enter\_name\_of\_the\_index\_here/\_mapping**

Example:

**GET temp\_index/\_mapping**

```
{
  "temp_index" : {
    "mappings" : {
      "properties" : {
        "botanical_name" : {
          "type" : "text",
          "fields" : {
            "keyword" : {
              "type" : "keyword",
              "ignore_above" : 256
            }
          }
        },
        "country_of_origin" : {
          "type" : "text",
          "fields" : {
            "keyword" : {
              "type" : "keyword",
              "ignore_above" : 256
            }
          }
        },
        "date_purchased" : {
          "type" : "date"
        },
        "description" : {
          "type" : "text",
          "fields" : {
            "keyword" : {
              "type" : "keyword",
              "ignore_above" : 256
            }
          }
        }
      }
    }
  },
}
```

# References

- <https://www.elastic.co/guide/index.html>
- Blog: <https://dev.to/lisahjung>

For more courses feel free to check our website:  
<https://sdadclub.tech/courses/courses.html>

