



ElasticSearch Performance Tuning Practice

Pei Wang
Mar 2018



Background

eBay Pronto, the platform that hosts Elasticsearch clusters

Currently

- Support Elasticsearch cluster deploy, manage, monitor, remediation, authenticate, decommission.
- Managed 60+ Elasticsearch clusters
- 2000+ nodes
- 18 billion documents ingested per day
- 3.5 billion search requests per day

Considering performance ...

- How to organize the index?
- Shard number?
- Replica number?
- Mapping?
- Routing?
- Cache?
- Any other settings?



- Best practice from eBay Pronto team ...

Agenda

- Optimize Index Design
- Index performance tuning
- Search performance tuning
- Test tools

Optimize Index Design

For example, there is an index which have one billion social media messages, and we have queries like below.

- **Organize your index by date**
 - Cases: Logging / Monitoring / Event
- **Organize your index by field, like region**
 - Cases: Query with **enumerable** filter field
- **Use routing key**
 - Distribute docs with same routing key to same shard
 - Cases: Query with **non-enumerable** filter field

```
GET _search
{
  "query": {
    "bool": {
      "must": [
        { "match": { "title": "Search" } }
      ],
      "filter": [ |
        { "term": { "region": "US" } },
        { "range": { "publish_date": { "gte": "2017-01-01" } } },
        { "match": { "author_id": 10086 } }
      ]
    }
  }
}
```

Optimize Index Design

- **Set mapping explicitly**
 - The default mapping may not fit your case
- **Make shards distributed evenly across nodes**
 - Nodes have more shards than others may become the bottle neck
- **Avoid imbalanced sharding if docs are indexed with routing key or user-defined ID**
 - User-defined ID should be random enough.
 - Imbalanced routing key distribution could cause imbalanced sharding.

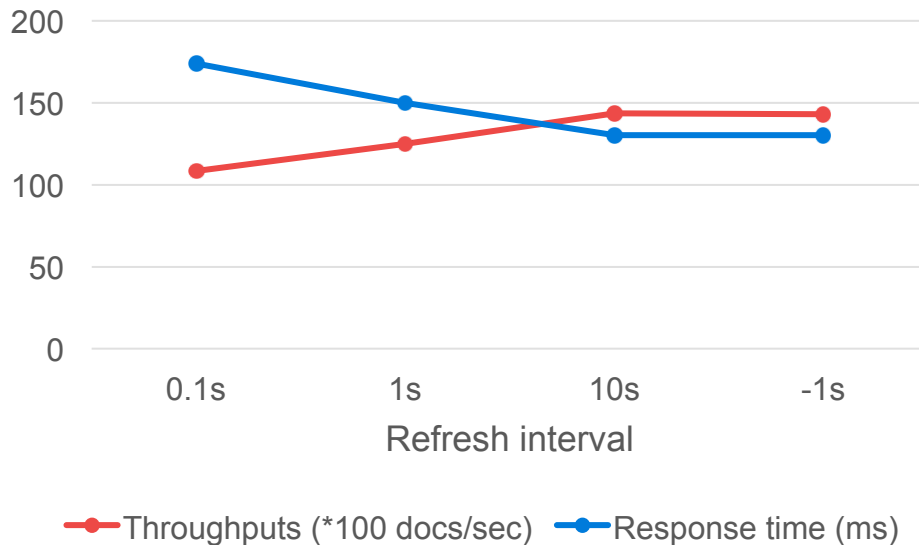


Agenda

- Optimize Index Design
- Index performance tuning
- Search performance tuning
- Test tools

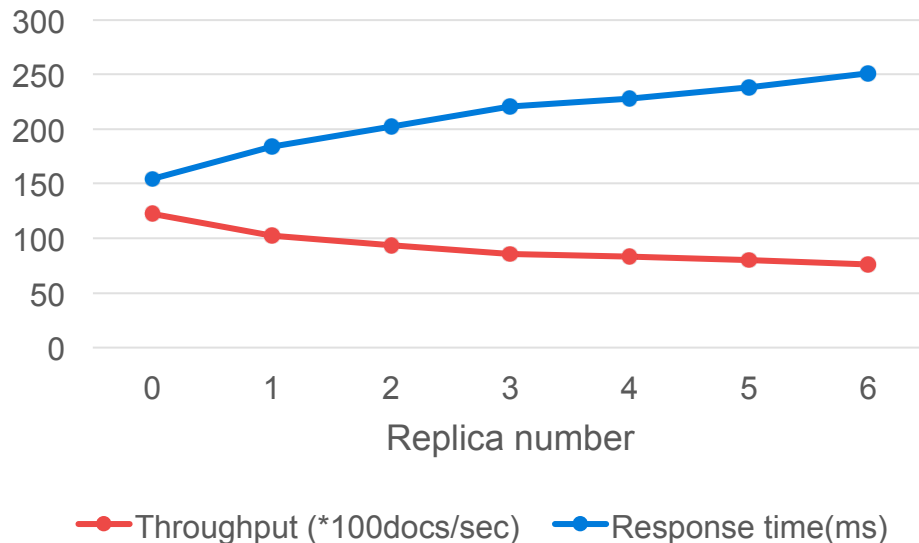
Index performance tuning

- **Increase shard number if node count > (1+ replica) * shard number**
 - Scale out, distribute data into more nodes
- **Increase refresh interval**
 - Elasticsearch create a segment every time refresh event happen. Increase interval would reduce segment count and merge cost.
 - Documents are not available for search until refresh



Index performance tuning

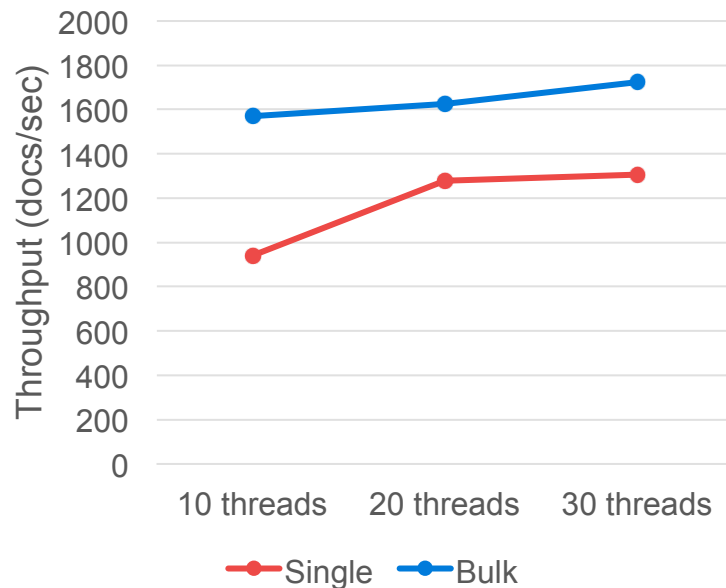
- **Use auto generated IDs if possible**
 - Elasticsearch auto generate ID algorithm can reduce the duplicate ID check and version check cost.
- **Reduce replica number**
 - Elasticsearch need to write primary shard and all replica shards for every index request
 - Replica number should not be 0, otherwise you will have data loss risk



Index performance tuning

Client side

- **Use bulk request**
- **Use multiple threads/workers**



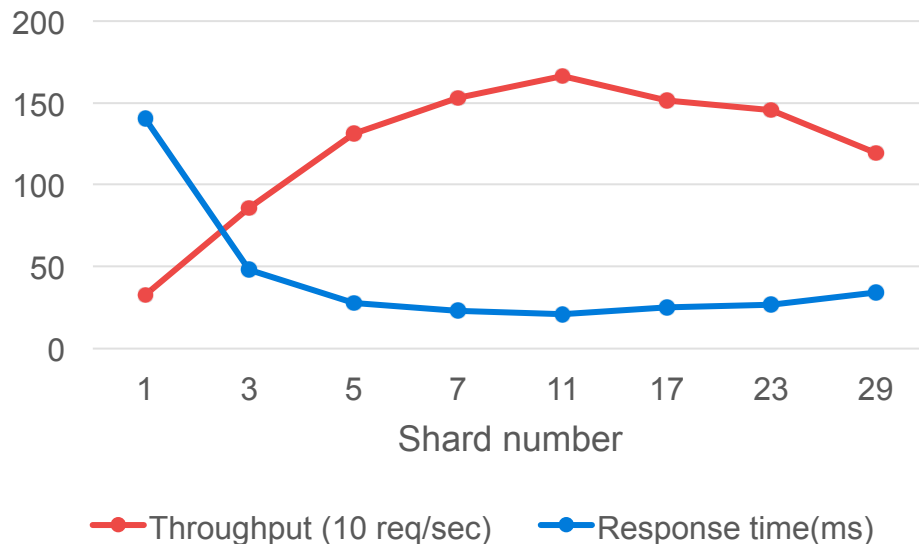
Agenda

- Optimize Index Design
- Index performance tuning
- Search performance tuning
- Test tools

Search performance tuning

- **Choose suitable shard number**

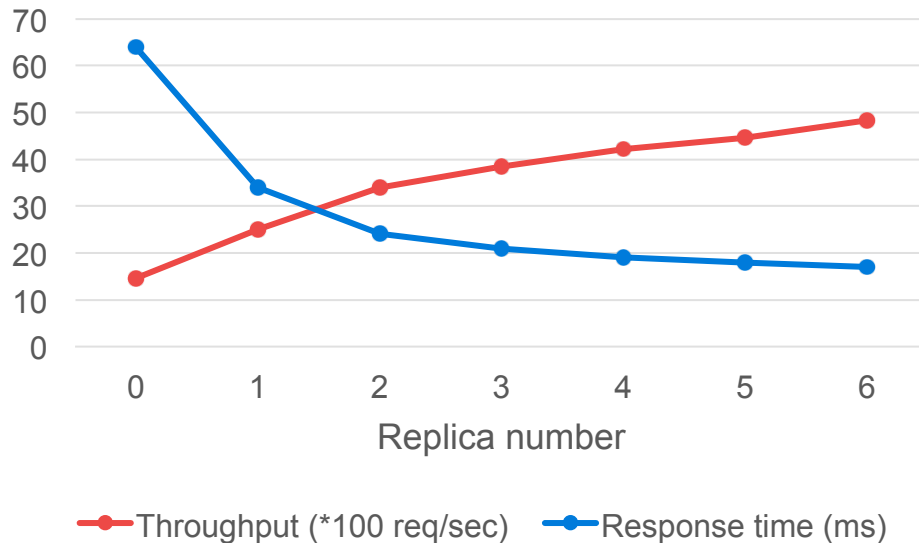
- Too small shard number will make search unable to scale out.
- Too big shard number will hurt performance too.
- Shard size should not exceed 30-50GB
- Notes, in this test, every shard has an exclusive node.



Search performance tuning

- **Increase replica number**

- Search can be performed on either a primary or replica shard.
- Will decrease indexing performance
- Notes, in this test, every shard has an exclusive node.



Search performance tuning

- **Use filter context instead of query context if possible**
 - Elasticsearch do not need to calculate relevancy score for filter context
 - Filtered result can be cached



Search performance tuning

- **Node query cache**
 - Only cache queries used in filter context.
 - Elasticsearch used bit set mechanism to cache filter results.
 - Elasticsearch only enable node query cache for segments have more than 10000 or 3% documents, whichever is larger.

```
GET index_name/_stats?filter_path=indices.**.query_cache
{
  "indices": {
    "index_name": {
      "primaries": {
        "query_cache": {
          "memory_size_in_bytes": 46004616,
          "total_count": 1588886,
          "hit_count": 515001,
          "miss_count": 1073885,
          "cache_size": 630,
          "cache_count": 630,
          "evictions": 0
        }
      },
      "total": {
        "query_cache": {
          "memory_size_in_bytes": 46004616,
          "total_count": 1588886,
          "hit_count": 515001,
          "miss_count": 1073885,
          "cache_size": 630,
          "cache_count": 630,
          "evictions": 0
        }
      }
    }
  }
}
```

Search performance tuning

- **Shard request cache**

- Only cache request size:0, like aggregate, suggestions or hit.totals.
- Use payload body as cache key, so that the payload must be the same, even the JSON body order.
- Do not use now or new Date() to build request body since it will make payload change every time, round your date time.
- Cache are invalidated when refresh happen and the data in the shard has actually changed.

```
GET index_name/_stats?filter_path=indices.**.request_cache
{
  "indices": {
    "index_name": {
      "primaries": {
        "request_cache": {
          "memory_size_in_bytes": 0,
          "evictions": 0,
          "hit_count": 541,
          "miss_count": 514098
        }
      },
      "total": {
        "request_cache": {
          "memory_size_in_bytes": 0,
          "evictions": 0,
          "hit_count": 982,
          "miss_count": 947321
        }
      }
    }
  }
}
```


Search performance tuning

- **Retrieve only necessary fields**
 - Use “_source” or “stored_fields” to let Elasticsearch only return necessary fields
- **Reduce documents count in response if possible**

GET /perf_test_result/_search

```
{
  "size": 5,
  "_source": [
    "job_name",
    "create_at"
  ]
}
```

```
"hits": {
  "total": 12338,
  "max_score": 1,
  "hits": [
    {
      "_index": "perf_test_result",
      "_type": "REQUEST",
      "_id": "AVwv0bQplyIgCso4VHV",
      "_score": 1,
      "_source": {
        "job_name": "single_shard_ingest_test",
        "create_at": "2017-05-22T08:15:29.703Z"
      }
    },
  ],
}
```

Search performance tuning

- Sort by “_doc” explicitly if do not care about the document order in response
 - Elasticsearch use the “_score” field to sort by score as default.
 - Use “sort”: “_doc” to let Elasticsearch return hits by index order
 - Especially helps when scrolling

```
GET /perf_test_result/_search
{
  "sort": ["_doc"]
}
```

Search performance tuning

- **Avoid searching stop words**

- Stop words like “a” and “the” may cause the query hit results count to explode.
- Use stop word filter
- Refine query, use “the AND fox” if you really meant to search “the” word.
- If some words are frequently used in your documents but not in the default list, use `cutoff_frequency` to specify which terms are stop words.

```
GET /_search
{
  "query": {
    "match" : {
      "message" : {
        "query" : "to be or not to be",
        "cutoff_frequency" : 0.001
      }
    }
  }
}
```

Search performance tuning

- **Avoid using script query to calculate hits in flight.**
 - Script query like below is quite time-consuming.
 - Should consider add extra fields in indexing phase if you have a lot of script query

```
{
  "query": {
    "script" : {
      "script" : {
        "inline": "doc['job_name.keyword'].value.startsWith('ebay')"
      }
    }
  }
}
```



```
{
  "query": {
    "match": {
      "job_name_prefix.keyword": "ebay"
    }
  }
}
```



Search performance tuning

- **Avoid wildcard query**
 - wildcard query like below is quite time-consuming, like script query, especially for leading wildcard query like “*ebay”
 - Should consider add extra fields in indexing phase if you have a lot of wildcard query

```
{  
  "query": {  
    "wildcard": {  
      "job_name.keyword": "ebay*"  
    }  
  }  
}
```



```
{  
  "query": {  
    "match": {  
      "job_name_prefix.keyword": "ebay"  
    }  
  }  
}
```



Agenda

- Optimize Index Design
- Index performance tuning
- Search performance tuning
- Test tools

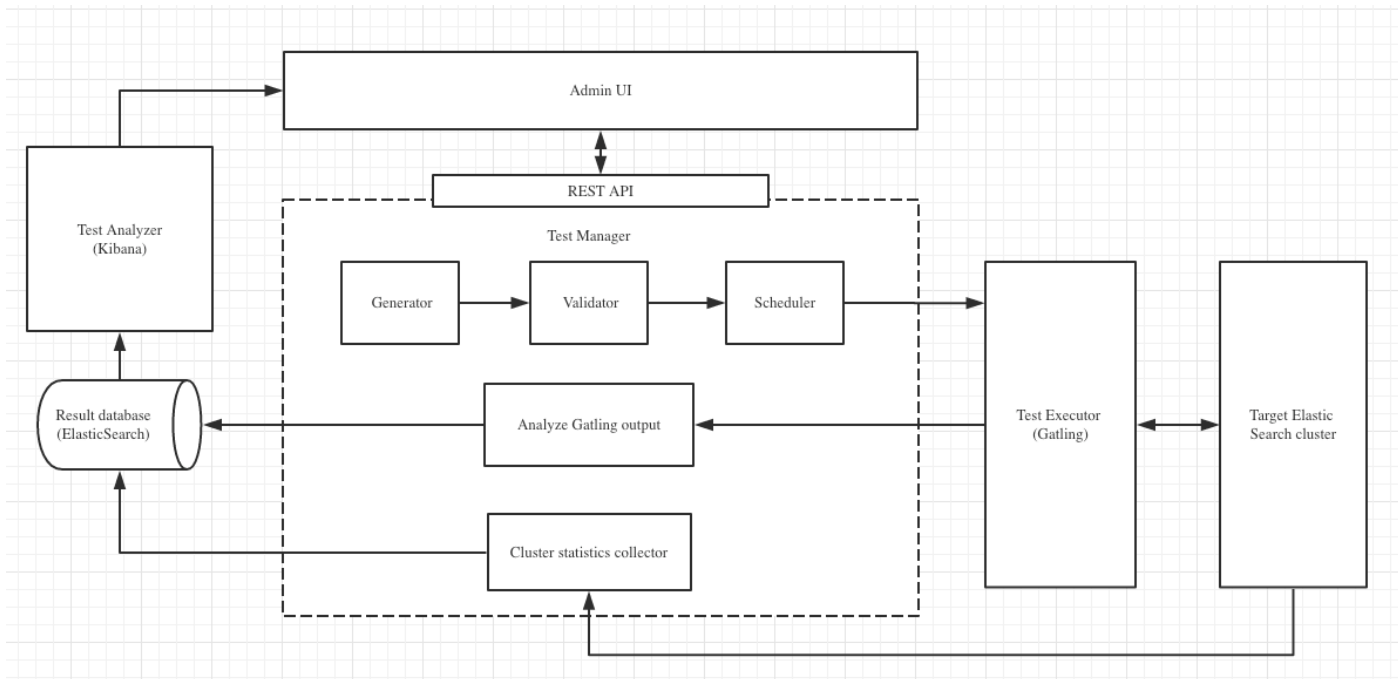
Test tools

Why we develop this tools

- Web UI, easy to access and use, supply performance test service for other teams.
- Run multiple tests with different configurations, change cluster setting and check cluster status when tests running.
- Help compare and analyze test results. Test reports are persisted and can be analyzed by Kibana.
- Rest APIs, easy to integrate with other systems.

Test tools

Architecture



Test tools

Screenshot

Pronto performance test manager

PREVIOUS JOBS

Job name
auth_622_read_vip_load_dis_auth

Id	Name	Simulation
1	search_600	search
2	search_600	search
3	search_600	search

Edit Test

Test Name
search_600

ES endpoint
http://10.137.210.88

Auth header

Index Name
test

Type Name
_doc

☒ Search
☐ Single indexing
☐ Bulk indexing

Search Catalog
ProntoSearch

Search Name
simple_query_city_random

Mode
Fixed request rate

Duration(seconds)
300

Requests per Second
600

SAVE SAVE AS NEW CLOSE

ChangeIndexSetting Actions

No	🗑️ ✎
No	🗑️ ✎
No	🗑️ ✎

ADD TEST START

Test tools

Test output



SeqId	test_name.keyword: Descending Q	Mean(ms)	Min	Max	50%	75%	95%	99%	<800ms(%)	800ms< x <1200ms(%)	> 1200ms(%)	Error percentage	count	throughput
0	50_search_per_s	26	2	1,397	12	21	97	184	98	0	0	2	1	50.088
1	100_search_per_s	37	2	594	15	33	163	235	99	0	0	1	1	99.545
2	200_search_per_s	155	2	1,523	113	217	431	869	97	1	0	2	1	199.286
3	300_search_per_s	1,924	1	3,489	2,125	2,353	2,608	2,803	8	9	81	2	1	299.12
4	400_search_per_s	23,249	0	60,126	14,698	35,755	54,159	59,155	0	0	68	32	1	389.501

Test tools

Monitor



Q&A