

Elasticsearch Merging Test Results

We did our 'Merge' testing with 2 different sets of settings. This documents lists the steps to do for merging using these and the pros and cons associated with both.

Settings 1:

1. Create an index with settings as shown below.

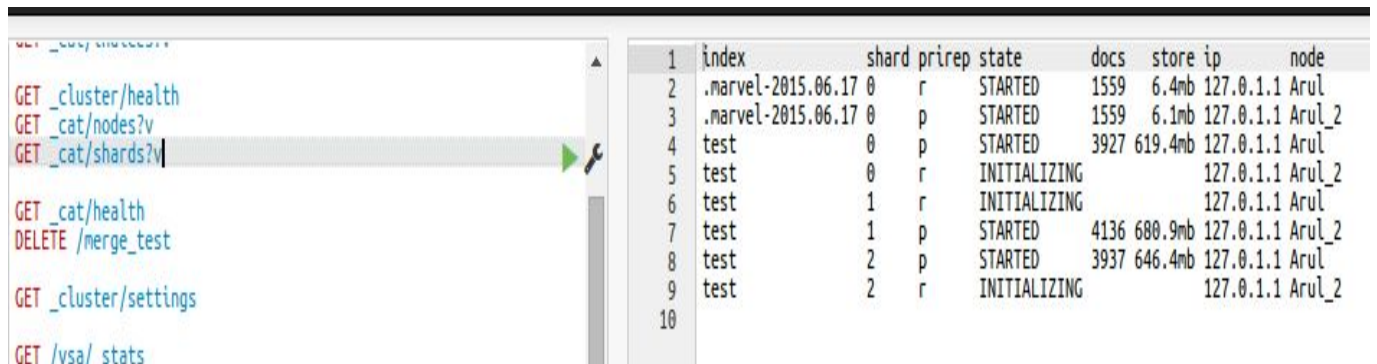
```
PUT /test/
{
  "settings": {
    "index": {
      "store": {
        "type": "default"
      },
      "number_of_shards": 3,[Ideal Value = Number of Team Members]
      "number_of_replicas": 0
    }
  }
}
```

2. Merge the documents.
3. Once all the documents have been merged, update the index settings to have replicas.

```
PUT /test/_settings
{
  "number_of_replicas":2 [Ideal Value = Number of team members - 1]
}
```

4. Monitor your shards' status. The newly created replicas will be in 'INITIALIZING' status.

IMPORTANT: Wait until the status of replica shards change to 'STARTED' state.



	index	shard	pri	rep	state	docs	store	ip	node
1	.marvel-2015.06.17	0	r		STARTED	1559	6.4mb	127.0.1.1	Arul
2	.marvel-2015.06.17	0	p		STARTED	1559	6.1mb	127.0.1.1	Arul_2
3	test	0	p		STARTED	3927	619.4mb	127.0.1.1	Arul
4	test	0	r		INITIALIZING			127.0.1.1	Arul_2
5	test	1	r		INITIALIZING			127.0.1.1	Arul
6	test	1	p		STARTED	4136	680.9mb	127.0.1.1	Arul_2
7	test	2	p		STARTED	3937	646.4mb	127.0.1.1	Arul
8	test	2	r		INITIALIZING			127.0.1.1	Arul_2
9									
10									

Settings 2:

1. Create an index with number of shards equal to the number of team members and number of replicas equal to the team number count - 1.

PUT /test/

```
{
  "settings": {
    "index": {
      "store": {
        "type": "default"
      },
      "number_of_shards": 3, [Ideal Value = Number of Team Members]
      "number_of_replicas": 2 [Ideal Value = Number of Team Members - 1]
    }
  }
}
```

2. Start your merging program and merge the results.

Settings 1	Settings 2
<p>Pro: Less time to merge.</p> <p>Test Results:</p> <ol style="list-style-type: none"> 1. Takes about 20-25 mins to merge 35000 documents. 2. Takes about 20 mins to create replica shards. 	<p>Con: More time to merge.</p> <p>Test Results:</p> <ol style="list-style-type: none"> 1. Takes about 3-4 hours to merge 25000 documents.
<p>Con:</p> <ol style="list-style-type: none"> 1.No way to check the status of replica shards' state transformation from 'INITIALIZING' to 'STARTED'. 2.While the state transformation always succeeds, it takes some time over which we have no control or monitoring ability. <p>Test Results:</p> <p>State Transformation from 'INITIALIZING' to 'STARTED' takes about 6-10 mins for shards of size 700MB.</p>	<p>Pro: No need to create any additional replicas</p>

Index Mapping:

PUT /test/document/_mapping

```
{
  "document": {
    "properties": {
      "docno": {
        "type": "string",
        "store": true,
        "index": "analyzed",
        "term_vector": "with_positions_offsets_payloads"
      },
      "HTTPheader": {
        "type": "string",
        "store": true,
        "index": "not_analyzed"
      },
      "title": {
        "type": "string",
        "store": true,
        "index": "analyzed",
        "term_vector": "with_positions_offsets_payloads"
      },
      "text": {
        "type": "string",
        "store": true,
        "index": "analyzed",
        "term_vector": "with_positions_offsets_payloads"
      },
      "html_Source": {
        "type": "string",
        "store": true,
        "index": "no"
      },
      "in_links": {
        "type": "string",
        "store": true,
        "index": "no"
      },
      "out_links": {
        "type": "string",
        "store": true,
        "index": "no"
      }
    }
  }
}
```

```
"author":{  
  "type": "string",  
  "store": true,  
  "index": "analyzed"  
}  
}  
}  
}
```