

---

# **eldam Documentation**

***Release 1.2.0***

**Srikanth Bemineni**

**Apr 30, 2017**



---

## Contents

---

<b>1</b>	<b>Some requirements kept in mind while designing this library</b>	<b>3</b>
<b>2</b>	<b>Transaction process</b>	<b>5</b>
2.1	Two step commit . . . . .	5
<b>3</b>	<b>API Documentation</b>	<b>7</b>
3.1	<code>connect(settings,default_index = "")</code> . . . . .	7
3.2	<code>connection()</code> . . . . .	7
3.3	<code>add(item)</code> . . . . .	8
3.4	<code>remove(item)</code> . . . . .	8
3.5	<code>update(item)</code> . . . . .	9
3.6	<code>update_by_query(item)</code> . . . . .	10
3.7	<code>delete_by_query(item)</code> . . . . .	11



Elasticsearch data manager (ELDAM) with zope transaction support. Other transaction data manager like sqlalchemy use the database atomic operation feature to rollback if an error occurs during the commit process. This is not achievable in elastic search. To overcome this eldam finalizes the records in the commit call of the two-phase transaction process. At the same time takes backup of the existing data in the index, if it involves removing or updating a record.



---

## Some requirements kept in mind while designing this library

---

- This module is not atomic. If any of the other transactions fail, then there is a brief period where the data still resides in Elasticsearch and will be removed during the `tcp_abort` call.
- **Please use it with your discretion if Elasticsearch is your primary data store.** This library is perfect if Elasticsearch is used only for searching in your application.





---

### Transaction process

---

`abort` - If needed, If any of the previous data managers aborted. This is before even beginning this data manager process. Elasticsearch data manager doesn't do anything in this call. Since we haven't committed anything to Elasticsearch

`tpc_abort` - If any of the other managers voted no after Elasticsearch data manager has committed, then this function will be called for cleanup

### Two step commit

- `tpc_begin` - Prepare for the transaction. Elasticsearch does nothing in this call
- `commit` - Elasticsearch data manager commits the data, at the same time takes backup of the existing data in case of delete or update. Raises an exception if the any of the commits fail.
- `tpc_vote` - After committing, vote and tell the transaction manager, that I am okay to go or not. Elasticsearch data manager votes yes
- `tpc_finish` - Final commit, no turning back after this. Elasticsearch data manager has already committed the data.



### `connect(settings, default_index = "")`

Connect to Elasticsearch database

Parameter

- settings - Dictionary with elasticsearch connections details
  - 'elasticsearch\_hosts' : array of Elasticsearch hosts [localhost, localhost:443, 16.74.45.322]
- default\_index - Default index to store documents. If not specified, then add, remove and update input need to mention the index details.

```
from eldam.elasticdatamanager import ElasticDataManager

edm = new ElasticDataManager()
edm.connect({
    "elasticsearch_hosts": ["16.83.62.232:9200", "16.83.63.114:9200"]
},
"test")
```

### `connection()`

Is class property to get the established connection.

```
from eldam.elasticdatamanager import ElasticDataManager

edm = new ElasticDataManager()
edm.connect({
    "elasticsearch_hosts": ["16.83.62.232:9200", "16.83.63.114:9200"]
},
"test")
```

```
# you can directly run api's defined from python Elasticsearch 5.0.0 and above
print(edm.connection.cluster.health())
```

## **add(item)**

Add a document to be included into Elasticsearch once the transaction is committed. If this instance of data manager was not added to the transaction manager, then this call will automatically add it to the current transaction.

Parameters

- item - Dictionary of the document to be added
  - ‘\_index’ - Index to which this document needs to be included. This is optional if default\_index is set during connect.
  - ‘\_type’ - Type of the document.
  - ‘\_id’ - Id of the document.
  - ‘\_source’ - Source of the document.

```
import transaction
from eldam.elasticdatamanager import ElasticDataManager

edm = ElasticDataManager()
edm.connect({
    "elasticsearch_hosts": ["16.83.62.232:9200", "16.83.63.114:9200"]
},
    "test")
edm.add({'_type': 'group',
        '_id': '2',
        '_source': {
            "gid": 234,
            "owner": "bemeneni",
            "name": "Sammy",
            "grp_hash": "456678",
            "description": "Sammy group"
        }}
transaction.commit()
```

## **remove(item)**

Remove an existing document from Elasticsearch once the transaction is committed. If this instance of data manager was not added to the transaction manager, then this call will automatically add it to the current transaction.

Parameters

- item - Dictionary of the document to be added
  - ‘\_index’ - Index to which this document needs to be included. This is optional if default\_index is set during connect.
  - ‘\_type’ - Type of the document.
  - ‘\_id’ - Id of the document.

```
import transaction
from eldam.elasticdatamanager import ElasticDataManager

edm = ElasticDataManager()
edm.connect({
    "elasticsearch_hosts": ["16.83.62.232:9200", "16.83.63.114:9200"]
},
    "test")
edm.remove({'_type': 'group',
            '_id': '2'})

# commit the transaction
transaction.commit()
```

## update(item)

Update an existing document from Elasticsearch once the transaction is committed. If the document doesn't exist in the index, then this call will automatically convert to an add call, and a new document will be included into ElasticSearch using the `_source` provided in the `item`. If this instance of data manager was not added to the transaction manager, then this call will automatically add it to the current transaction.

### Parameters

- `item` - Dictionary of the document to be added
  - `'_index'` - Index to which this document needs to be included. This is optional if `default_index` is set during `connect`.
  - `'_type'` - Type of the document.
  - `'_id'` - Id of the document.
  - `'_source'` - Source of the document. Can be partial update attributes of the original document.

```
import transaction
from eldam.elasticdatamanager import ElasticDataManager

edm = ElasticDataManager()
edm.connect({
    "elasticsearch_hosts": ["16.83.62.232:9200", "16.83.63.114:9200"]
},
    "test")
edm.update({'_type': 'group',
            '_id': '2',
            '_source': {
                "description": "Sammy study group"
            }}

# commit the transaction
transaction.commit()
```

## update\_by\_query(item)

Update existing documents from Elasticsearch using a query and script once the transaction is committed. If this instance of data manager was not added to the transaction manager, then this call will automatically add it to the current transaction.

### Parameters

- item - Dictionary of the document to be added
  - ‘\_index’ - Index to which this document needs to be included. This is optional if default\_index is set during connect.
  - ‘\_type’ - Type of the document.
  - ‘\_query’ - Elasticsearch query using DSL format. <https://www.elastic.co/guide/en/elasticsearch/reference/current/query-dsl.html>
  - ‘\_script’ - Specify the script to update in any language, as long as the language is enabled in the deployed Elasticsearch. By default “painless” lang is supported in Elasticsearch

Below example searches for fields with id starting with ‘g’, and updates the description and name fields using the script provided. <https://www.elastic.co/guide/en/elasticsearch/reference/current/modules-scripting.html>

```
import transaction
from eldam.elasticdatamanager import ElasticDataManager

edm = ElasticDataManager()
edm.connect({
    "elasticsearch_hosts": ["16.83.62.232:9200", "16.83.63.114:9200"]
},
    "test")
item = {}
item['query'] = {
    "wildcard": {
        "_uid": {
            "value": "g*"
        }
    }
}

item['script'] = {
    "inline": "ctx._source.description = params.description; ctx._source.
↪name = params.name",
    "lang": "painless",
    'params': {
        "description": "eldam group",
        "name": "eldam"
    }
}
edm.update_by_query(item)

# commit the transaction
transaction.commit()

# if you need to access the data immediately, then refresh the indices
edm.connection.indices.refresh(index="_all")
```

## delete\_by\_query(item)

Delete existing documents from Elasticsearch using a query once the transaction is committed. If this instance of data manager was not included in the transaction manager, then this call will automatically add it to the current transaction.

Parameters

- item - Dictionary of the document to be added
  - ‘\_index’ - Index to which this document needs to be included. This is optional if default\_index is set during connect.
  - ‘\_type’ - Type of the document.
  - ‘\_query’ - Elasticsearch query using DSL format. <https://www.elastic.co/guide/en/elasticsearch/reference/current/query-dsl.html>

Below example deletes all the documents with an id starting with ‘g’.

```
import transaction
from eldam.elasticdatamanager import ElasticDataManager

edm = ElasticDataManager()
edm.connect({
    "elasticsearch_hosts": ["16.83.62.232:9200", "16.83.63.114:9200"]
},
"test")
item = {}
item['query'] = {
    "wildcard": {
        "_uid": {
            "value": "g*"
        }
    }
}

edm.delete_by_query(item)

# commit the transaction
transaction.commit()
```