Elasticsearch Documentation

Release 1.1.1

Honza Král

July 13, 2014

Contents

Official low-level client for Elasticsearch. It's goal is to provide common ground for all Elasticsearch-related code in Python; because of this it tries to be opinion-free and very extendable.

Compatibility

The library is compatible with both Elasticsearch 1.x and 0.90.x but you **have to use a matching version**.

For **Elasticsearch 1.0** and later, use the major version 1 (`1.x.y`) of the library.

For **Elasticsearch 0.90.x**, use a version from `0.4.x` releases of the library.

The recommended way to set your requirements in your *setup.py* or *requirements.txt* is:

```
# Elasticsearch 1.0
elasticsearch>=1.0.0,<2.0.0

# Elasticsearch 0.90
elasticsearch<1.0.0
```

The development is happening on `master` and `0.4` branches, respectively.

Example Usage

```python
from datetime import datetime
from elasticsearch import Elasticsearch
es = Elasticsearch()

doc = {
    'author': 'kimchy',
    'text': 'Elasticsearch: cool. bonsai cool.',
    'timestamp': datetime(2010, 10, 10, 10, 10, 10)
}
res = es.index(index="test-index", doc_type='tweet', id=1, body=doc)
print(res['created'])

res = es.get(index="test-index", doc_type='tweet', id=1)
print(res['_source'])

es.indices.refresh(index="test-index")

res = es.search(index="test-index", body={"query": {"match_all": {}}})
print("Got %d Hits:" % res['hits']['total'])
for hit in res['hits']['hits']:
    print("%(timestamp)s %(author)s: %(text)s" % hit["_source"])
```

Features

This client was designed as very thin wrapper around Elasticseach's REST API to allow for maximum flexibility. This means that there are no opinions in this client; it also means that some of the APIs are a little cumbersome to use from Python. We have created some *Helpers* to help with this issue.

3.1 Persistent Connections

`elasticsearch-py` uses persistent connections inside of individual connection pools (one per each configured or sniffed node). Out of the box you can choose to use `http`, `thrift` or an experimental `memcached` protocol to communicate with the elasticsearch nodes. See *Transport classes* for more information.

The transport layer will create an instance of the selected connection class per node and keep track of the health of individual nodes - if a node becomes unresponsive (throwing exceptions while connecting to it) it's put on a timeout by the `ConnectionPool` class and only returned to the circulation after the timeout is over (or when no live nodes are left). By default node are randomized before passed into the pool and round-robin strategy is used for load balancing.

You can customize this behavior by passing parameters to the *Connection Layer API* (all keyword arguments to the `Elasticsearch` class will be passed through). If what you want to accomplish is not supported you should be able to create a subclass of the relevant component and pass it in as a parameter to be used instead of the default implementation.

3.2 Sniffing

The client can be configured to inspect the cluster state to get a list of nodes upon startup, periodically and/or on failure. See `Transport` parameters for details.

Some example configurations:

```python
from elasticsearch import Elasticsearch

# by default we don't sniff, ever
es = Elasticsearch()

# you can specify to sniff on startup to inspect the cluster and load
# balance across all nodes
es = Elasticsearch(["seed1", "seed2"], sniff_on_start=True)

# you can also sniff periodically and/or after failure:
es = Elasticsearch(["seed1", "seed2"], sniff_on_start=True, sniff_on_connection_fail=True, sniffer_t:
```

3.3 Logging

`elasticsearch-py` uses the standard logging library from python to define two loggers: `elasticsearch` and `elasticsearch.trace`. `elasticsearch` is used by the client to log standard activity, depending on the log level. `elasticsearch.trace` can be used to log requests to the server in the form of `curl` commands using pretty-printed json that can then be executed from command line. The trace logger doesn't inherit from the base one - it needs to be activated separately.

Contents

4.1 API Documentation

All the API calls map the raw REST api as closely as possible, including the distinction between required and optional arguments to the calls. This means that the code makes distinction between positional and keyword arguments; we, however, recommend that people **use keyword arguments for all calls for consistency and safety**.

An API call is considered successful (and will return a response) if elasticsearch returns a 2XX response. Otherwise an instance of `TransportError` (or a more specific subclass) will be raised. You can see other exception and error states in *Exceptions*. If you do not wish an exception to be raised you can always pass in an `ignore` parameter with either a single status code that should be ignored or a list of them:

```python
from elasticsearch import Elasticsearch
es = Elasticsearch()

# ignore 400 cause by IndexAlreadyExistsException when creating an index
es.indices.create(index='test-index', ignore=400)

# ignore 404 and 400
es.indices.delete(index='test-index', ignore=[400, 404])
```

You can also specify `request_timeout` (float) as part of any API call, this value will get passed to the `perform_request` method of the connection class:

```python
# only wait for 1 second, regardless of the client's default
es.cluster.health(wait_for_status='yellow', request_timeout=1)
```

Note: for compatibility with the Python ecosystem we use `from_` instead of `from` and `doc_type` instead of `type` as parameter names.

4.1.1 Elasticsearch

class elasticsearch.**Elasticsearch**(*hosts=None*, *transport_class=<class 'elasticsearch.transport.Transport'>*, ***kwargs*)
Elasticsearch low-level client. Provides a straightforward mapping from Python to ES REST endpoints.

The instance has attributes *cat*, *cluster*, *indices*, *nodes* and *snapshot* that provide access to instances of `CatClient`, `ClusterClient`, `IndicesClient`, `NodesClient` and `SnapshotClient` respectively. This is the preferred (and only supported) way to get access to those classes and their methods.

Some examples:

```
# create connection to localhost using the ThriftConnection and it's
# default port (9500)
es = Elasticsearch(connection_class=ThriftConnection)

# create connection that will automatically inspect the cluster to get
# the list of active nodes. Start with nodes 'esnode1' and 'esnode2'
es = Elasticsearch(
    ['esnode1', 'esnode2'],
    # sniff before doing anything
    sniff_on_start=True,
    # refresh nodes after a node fails to respond
    sniff_on_connection_fail=True,
    # and also every 60 seconds
    sniffer_timeout=60
)

# connect to localhost directly and another node using SSL on port 443
# and an url_prefix
es = Elasticsearch([
    {'host': 'localhost'},
    {'host': 'othernode', 'port': 443, 'url_prefix': 'es', 'use_ssl': True},
])
```

> **Parameters**
>
> - **hosts** – list of nodes we should connect to. Node should be a dictionary ({"host": "local-host", "port": 9200}), the entire dictionary will be passed to the `Connection` class as kwargs, or a string in the format of `host[:port]` which will be translated to a dictionary automatically. If no value is given the `Connection` class defaults will be used.
>
> - **transport_class** – `Transport` subclass to use.
>
> - **kwargs** – any additional arguments will be passed on to the `Transport` class and, subsequently, to the `Connection` instances.

abort_benchmark(*args*, ***kwargs*)

Aborts a running benchmark. http://www.elasticsearch.org/guide/en/elasticsearch/reference/master/search-benchmark.html

> **Parameters name** – A benchmark name

benchmark(*args*, ***kwargs*)

The benchmark API provides a standard mechanism for submitting queries and measuring their performance relative to one another. http://www.elasticsearch.org/guide/en/elasticsearch/reference/master/search-benchmark.html

> **Parameters**
>
> - **index** – A comma-separated list of index names; use *_all* or empty string to perform the operation on all indices
>
> - **doc_type** – The name of the document type
>
> - **body** – The search definition using the Query DSL
>
> - **verbose** – Specify whether to return verbose statistics about each iteration (default: false)

bulk(*args*, ***kwargs*)

Perform many index/delete operations in a single API call. http://www.elasticsearch.org/guide/en/elasticsearch/reference/curl bulk.html

See the `bulk()` helper function for a more friendly API.

> **Parameters**
>
> - **body** – The operation definition and data (action-data pairs), as either a newline separated string, or a sequence of dicts to serialize (one per row).
> - **index** – Default index for items which don't provide one
> - **doc_type** – Default document type for items which don't provide one
> - **consistency** – Explicit write consistency setting for the operation
> - **refresh** – Refresh the index after performing the operation
> - **routing** – Specific routing value
> - **replication** – Explicitly set the replication type (default: sync)
> - **timeout** – Explicit operation timeout

clear_scroll(**args*, ***kwargs*)

Clear the scroll request created by specifying the scroll parameter to search. http://www.elasticsearch.org/guide/en/elasticsearch/reference/current/search-request-scroll.html

> **Parameters**
>
> - **scroll_id** – The scroll ID or a list of scroll IDs
> - **body** – A comma-separated list of scroll IDs to clear if none was specified via the scroll_id parameter

count(**args*, ***kwargs*)

Execute a query and get the number of matches for that query. http://www.elasticsearch.org/guide/en/elasticsearch/reference/current/search-count.html

> **Parameters**
>
> - **index** – A comma-separated list of indices to restrict the results
> - **doc_type** – A comma-separated list of types to restrict the results
> - **body** – A query to restrict the results (optional)
> - **allow_no_indices** – Whether to ignore if a wildcard indices expression resolves into no concrete indices. (This includes _all string or when no indices have been specified)
> - **expand_wildcards** – Whether to expand wildcard expression to concrete indices that are open, closed or both., default 'open'
> - **ignore_unavailable** – Whether specified concrete indices should be ignored when unavailable (missing or closed)
> - **min_score** – Include only documents with a specific _score value in the result
> - **preference** – Specify the node or shard the operation should be performed on (default: random)
> - **q** – Query in the Lucene query string syntax
> - **routing** – Specific routing value
> - **source** – The URL-encoded query definition (instead of using the request body)

count_percolate(**args*, ***kwargs*)

> The percolator allows to register queries against an index, and then send percolate requests which include a doc, and getting back the queries that match on that doc out of the set of registered queries. http://www.elasticsearch.org/guide/en/elasticsearch/reference/current/search-percolate.html

> **Parameters**

> * **index** – The index of the document being count percolated.
> * **doc_type** – The type of the document being count percolated.
> * **id** – Substitute the document in the request body with a document that is known by the specified id. On top of the id, the index and type parameter will be used to retrieve the document from within the cluster.
> * **body** – The count percolator request definition using the percolate DSL
> * **allow_no_indices** – Whether to ignore if a wildcard indices expression resolves into no concrete indices. (This includes _all string or when no indices have been specified)
> * **expand_wildcards** – Whether to expand wildcard expression to concrete indices that are open, closed or both., default 'open'
> * **ignore_unavailable** – Whether specified concrete indices should be ignored when unavailable (missing or closed)
> * **percolate_index** – The index to count percolate the document into. Defaults to index.
> * **percolate_type** – The type to count percolate document into. Defaults to type.
> * **preference** – Specify the node or shard the operation should be performed on (default: random)
> * **routing** – A comma-separated list of specific routing values
> * **version** – Explicit version number for concurrency control
> * **version_type** – Specific version type

create(**args*, ***kwargs*)

> Adds a typed JSON document in a specific index, making it searchable. Behind the scenes this method calls index(..., op_type='create') http://www.elasticsearch.org/guide/en/elasticsearch/reference/current/docs-index_.html

> **Parameters**

> * **index** – The name of the index
> * **doc_type** – The type of the document
> * **id** – Document ID
> * **body** – The document
> * **consistency** – Explicit write consistency setting for the operation
> * **id** – Specific document ID (when the POST method is used)
> * **parent** – ID of the parent document
> * **percolate** – Percolator queries to execute while indexing the document
> * **refresh** – Refresh the index after performing the operation
> * **replication** – Specific replication type (default: sync)
> * **routing** – Specific routing value

- **timeout** – Explicit operation timeout
- **timestamp** – Explicit timestamp for the document
- **ttl** – Expiration time for the document
- **version** – Explicit version number for concurrency control
- **version_type** – Specific version type

delete(**args*, ***kwargs*)

Delete a typed JSON document from a specific index based on its id. http://www.elasticsearch.org/guide/en/elasticsearch/reference/current/docs-delete.html

> **Parameters**
>
> - **index** – The name of the index
> - **doc_type** – The type of the document
> - **id** – The document ID
> - **consistency** – Specific write consistency setting for the operation
> - **parent** – ID of parent document
> - **refresh** – Refresh the index after performing the operation
> - **replication** – Specific replication type (default: sync)
> - **routing** – Specific routing value
> - **timeout** – Explicit operation timeout
> - **version** – Explicit version number for concurrency control
> - **version_type** – Specific version type

delete_by_query(**args*, ***kwargs*)

Delete documents from one or more indices and one or more types based on a query. http://www.elasticsearch.org/guide/en/elasticsearch/reference/current/docs-delete-by-query.html

> **Parameters**
>
> - **index** – A comma-separated list of indices to restrict the operation; use *_all* to perform the operation on all indices
> - **doc_type** – A comma-separated list of types to restrict the operation
> - **body** – A query to restrict the operation specified with the Query DSL
> - **allow_no_indices** – Whether to ignore if a wildcard indices expression resolves into no concrete indices. (This includes *_all* string or when no indices have been specified)
> - **analyzer** – The analyzer to use for the query string
> - **consistency** – Specific write consistency setting for the operation
> - **default_operator** – The default operator for query string query (AND or OR), default u'OR'
> - **df** – The field to use as default where no field prefix is given in the query string
> - **expand_wildcards** – Whether to expand wildcard expression to concrete indices that are open, closed or both., default u'open'
> - **ignore_unavailable** – Whether specified concrete indices should be ignored when unavailable (missing or closed)

- **q** – Query in the Lucene query string syntax

- **replication** – Specific replication type, default u'sync'

- **routing** – Specific routing value

- **source** – The URL-encoded query definition (instead of using the request body)

- **timeout** – Explicit operation timeout

exists(**args*, ***kwargs*)

Returns a boolean indicating whether or not given document exists in Elasticsearch. http://www.elasticsearch.org/guide/en/elasticsearch/reference/current/docs-get.html

Parameters

- **index** – The name of the index

- **id** – The document ID

- **doc_type** – The type of the document (uses *_all* by default to fetch the first document matching the ID across all types)

- **parent** – The ID of the parent document

- **preference** – Specify the node or shard the operation should be performed on (default: random)

- **realtime** – Specify whether to perform the operation in realtime or search mode

- **refresh** – Refresh the shard containing the document before performing the operation

- **routing** – Specific routing value

explain(**args*, ***kwargs*)

The explain api computes a score explanation for a query and a specific document. This can give useful feedback whether a document matches or didn't match a specific query. http://www.elasticsearch.org/guide/en/elasticsearch/reference/current/search-explain.html

Parameters

- **index** – The name of the index

- **doc_type** – The type of the document

- **id** – The document ID

- **body** – The query definition using the Query DSL

- **_source** – True or false to return the _source field or not, or a list of fields to return

- **_source_exclude** – A list of fields to exclude from the returned _source field

- **_source_include** – A list of fields to extract and return from the _source field

- **analyze_wildcard** – Specify whether wildcards and prefix queries in the query string query should be analyzed (default: false)

- **analyzer** – The analyzer for the query string query

- **default_operator** – The default operator for query string query (AND or OR), (default: OR)

- **df** – The default field for query string query (default: _all)

- **fields** – A comma-separated list of fields to return in the response

- **lenient** – Specify whether format-based query failures (such as providing text to a numeric field) should be ignored

- **lowercase_expanded_terms** – Specify whether query terms should be lowercased

- **parent** – The ID of the parent document

- **preference** – Specify the node or shard the operation should be performed on (default: random)

- **q** – Query in the Lucene query string syntax

- **routing** – Specific routing value

- **source** – The URL-encoded query definition (instead of using the request body)

get (*args*, **kwargs*)

Get a typed JSON document from the index based on its id. http://www.elasticsearch.org/guide/en/elasticsearch/reference/current/docs-get.html

Parameters

- **index** – The name of the index

- **id** – The document ID

- **doc_type** – The type of the document (uses _*all* by default to fetch the first document matching the ID across all types)

- **_source** – True or false to return the _source field or not, or a list of fields to return

- **_source_exclude** – A list of fields to exclude from the returned _source field

- **_source_include** – A list of fields to extract and return from the _source field

- **fields** – A comma-separated list of fields to return in the response

- **parent** – The ID of the parent document

- **preference** – Specify the node or shard the operation should be performed on (default: random)

- **realtime** – Specify whether to perform the operation in realtime or search mode

- **refresh** – Refresh the shard containing the document before performing the operation

- **routing** – Specific routing value

- **version** – Explicit version number for concurrency control

- **version_type** – Explicit version number for concurrency control

get_source (*args*, **kwargs*)

Get the source of a document by it's index, type and id. http://www.elasticsearch.org/guide/en/elasticsearch/reference/current/docs-get.html

Parameters

- **index** – The name of the index

- **doc_type** – The type of the document (uses _*all* by default to fetch the first document matching the ID across all types)

- **id** – The document ID

- **_source** – True or false to return the _source field or not, or a list of fields to return

- **_source_exclude** – A list of fields to exclude from the returned _source field

- **_source_include** – A list of fields to extract and return from the _source field
- **parent** – The ID of the parent document
- **preference** – Specify the node or shard the operation should be performed on (default: random)
- **realtime** – Specify whether to perform the operation in realtime or search mode
- **refresh** – Refresh the shard containing the document before performing the operation
- **routing** – Specific routing value
- **version** – Explicit version number for concurrency control
- **version_type** – Explicit version number for concurrency control

index(**args*, ***kwargs*)

Adds or updates a typed JSON document in a specific index, making it searchable. http://www.elasticsearch.org/guide/en/elasticsearch/reference/current/docs-index_.html

Parameters

- **index** – The name of the index
- **doc_type** – The type of the document
- **body** – The document
- **id** – Document ID
- **consistency** – Explicit write consistency setting for the operation
- **op_type** – Explicit operation type (default: index)
- **parent** – ID of the parent document
- **refresh** – Refresh the index after performing the operation
- **replication** – Specific replication type (default: sync)
- **routing** – Specific routing value
- **timeout** – Explicit operation timeout
- **timestamp** – Explicit timestamp for the document
- **ttl** – Expiration time for the document
- **version** – Explicit version number for concurrency control
- **version_type** – Specific version type

info(**args*, ***kwargs*)

Get the basic info from the current cluster.

list_benchmarks(**args*, ***kwargs*)

View the progress of long-running benchmarks. http://www.elasticsearch.org/guide/en/elasticsearch/reference/master/search-benchmark.html

Parameters

- **index** – A comma-separated list of index names; use *_all* or empty string to perform the operation on all indices
- **doc_type** – The name of the document type

mget (**args*, ***kwargs*)

> Get multiple documents based on an index, type (optional) and ids. http://www.elasticsearch.org/guide/en/elasticsearch/reference/current/docs-multi-get.html

> > **Parameters**
> >
> > - **body** – Document identifiers; can be either *docs* (containing full document information) or *ids* (when index and type is provided in the URL.
> >
> > - **index** – The name of the index
> >
> > - **doc_type** – The type of the document
> >
> > - **_source** – True or false to return the _source field or not, or a list of fields to return
> >
> > - **_source_exclude** – A list of fields to exclude from the returned _source field
> >
> > - **_source_include** – A list of fields to extract and return from the _source field
> >
> > - **fields** – A comma-separated list of fields to return in the response
> >
> > - **parent** – The ID of the parent document
> >
> > - **preference** – Specify the node or shard the operation should be performed on (default: random)
> >
> > - **realtime** – Specify whether to perform the operation in realtime or search mode
> >
> > - **refresh** – Refresh the shard containing the document before performing the operation
> >
> > - **routing** – Specific routing value

mlt (**args*, ***kwargs*)

> Get documents that are "like" a specified document. http://www.elasticsearch.org/guide/en/elasticsearch/reference/current/search-more-like-this.html

> > **Parameters**
> >
> > - **index** – The name of the index
> >
> > - **doc_type** – The type of the document (use *_all* to fetch the first document matching the ID across all types)
> >
> > - **id** – The document ID
> >
> > - **body** – A specific search request definition
> >
> > - **boost_terms** – The boost factor
> >
> > - **include** – Whether to include the queried document from the response
> >
> > - **max_doc_freq** – The word occurrence frequency as count: words with higher occurrence in the corpus will be ignored
> >
> > - **max_query_terms** – The maximum query terms to be included in the generated query
> >
> > - **max_word_length** – The minimum length of the word: longer words will be ignored
> >
> > - **min_doc_freq** – The word occurrence frequency as count: words with lower occurrence in the corpus will be ignored
> >
> > - **min_term_freq** – The term frequency as percent: terms with lower occurence in the source document will be ignored
> >
> > - **min_word_length** – The minimum length of the word: shorter words will be ignored
> >
> > - **mlt_fields** – Specific fields to perform the query against

- **percent_terms_to_match** – How many terms have to match in order to consider the document a match (default: 0.3)

- **routing** – Specific routing value

- **search_from** – The offset from which to return results

- **search_indices** – A comma-separated list of indices to perform the query against (default: the index containing the document)

- **search_query_hint** – The search query hint

- **search_scroll** – A scroll search request definition

- **search_size** – The number of documents to return (default: 10)

- **search_source** – A specific search request definition (instead of using the request body)

- **search_type** – Specific search type (eg. *dfs_then_fetch*, *count*, etc)

- **search_types** – A comma-separated list of types to perform the query against (default: the same type as the document)

- **stop_words** – A list of stop words to be ignored

mpercolate(**args*, ***kwargs*)

The percolator allows to register queries against an index, and then send percolate requests which include a doc, and getting back the queries that match on that doc out of the set of registered queries. http://www.elasticsearch.org/guide/en/elasticsearch/reference/current/search-percolate.html

Parameters

- **index** – The index of the document being count percolated to use as default

- **doc_type** – The type of the document being percolated to use as default.

- **body** – The percolate request definitions (header & body pair), separated by newlines

- **allow_no_indices** – Whether to ignore if a wildcard indices expression resolves into no concrete indices. (This includes *_all* string or when no indices have been specified)

- **expand_wildcards** – Whether to expand wildcard expression to concrete indices that are open, closed or both., default 'open'

- **ignore_unavailable** – Whether specified concrete indices should be ignored when unavailable (missing or closed)

msearch(**args*, ***kwargs*)

Execute several search requests within the same API. http://www.elasticsearch.org/guide/en/elasticsearch/reference/current/search-multi-search.html

Parameters

- **body** – The request definitions (metadata-search request definition pairs), as either a newline separated string, or a sequence of dicts to serialize (one per row).

- **index** – A comma-separated list of index names to use as default

- **doc_type** – A comma-separated list of document types to use as default

- **search_type** – Search operation type

mtermvectors(**args*, ***kwargs*)

Multi termvectors API allows to get multiple termvectors based on an index, type and id. http://www.elasticsearch.org/guide/en/elasticsearch/reference/master/docs-multi-termvectors.html

Parameters

- **index** – The index in which the document resides.

- **doc_type** – The type of the document.

- **body** – Define ids, parameters or a list of parameters per document here. You must at least provide a list of document ids. See documentation.

- **field_statistics** – Specifies if document count, sum of document frequencies and sum of total term frequencies should be returned. Applies to all returned documents unless otherwise specified in body "params" or "docs"., default True

- **fields** – A comma-separated list of fields to return. Applies to all returned documents unless otherwise specified in body "params" or "docs".

- **ids** – A comma-separated list of documents ids. You must define ids as parameter or set "ids" or "docs" in the request body

- **offsets** – Specifies if term offsets should be returned. Applies to all returned documents unless otherwise specified in body "params" or "docs"., default True

- **parent** – Parent id of documents. Applies to all returned documents unless otherwise specified in body "params" or "docs".

- **payloads** – Specifies if term payloads should be returned. Applies to all returned documents unless otherwise specified in body "params" or "docs"., default True

- **positions** – Specifies if term positions should be returned. Applies to all returned documents unless otherwise specified in body "params" or "docs"., default True

- **preference** – Specify the node or shard the operation should be performed on (default: random) .Applies to all returned documents unless otherwise specified in body "params" or "docs".

- **routing** – Specific routing value. Applies to all returned documents unless otherwise specified in body "params" or "docs".

- **term_statistics** – Specifies if total term frequency and document frequency should be returned. Applies to all returned documents unless otherwise specified in body "params" or "docs"., default False

percolate(**args*, ***kwargs*)

The percolator allows to register queries against an index, and then send percolate requests which include a doc, and getting back the queries that match on that doc out of the set of registered queries. http://www.elasticsearch.org/guide/en/elasticsearch/reference/current/search-percolate.html

Parameters

- **index** – The index of the document being percolated.

- **doc_type** – The type of the document being percolated.

- **id** – Substitute the document in the request body with a document that is known by the specified id. On top of the id, the index and type parameter will be used to retrieve the document from within the cluster.

- **body** – The percolator request definition using the percolate DSL

- **allow_no_indices** – Whether to ignore if a wildcard indices expression resolves into no concrete indices. (This includes _all string or when no indices have been specified)

- **expand_wildcards** – Whether to expand wildcard expression to concrete indices that are open, closed or both., default 'open'

- **ignore_unavailable** – Whether specified concrete indices should be ignored when un-available (missing or closed)

- **percolate_format** – Return an array of matching query IDs instead of objects

- **percolate_index** – The index to percolate the document into. Defaults to index.

- **percolate_type** – The type to percolate document into. Defaults to type.

- **preference** – Specify the node or shard the operation should be performed on (default: random)

- **routing** – A comma-separated list of specific routing values

- **version** – Explicit version number for concurrency control

- **version_type** – Specific version type

ping(**args*, ***kwargs*)
   Returns True if the cluster is up, False otherwise.

scroll(**args*, ***kwargs*)
   Scroll a search request created by specifying the scroll parameter. http://www.elasticsearch.org/guide/en/elasticsearch/reference/current/search-request-scroll.html

   Parameters

   - **scroll_id** – The scroll ID

   - **scroll** – Specify how long a consistent view of the index should be maintained for scrolled search

search(**args*, ***kwargs*)
   Execute a search query and get back search hits that match the query. http://www.elasticsearch.org/guide/en/elasticsearch/reference/current/search-search.html

   Parameters

   - **index** – A comma-separated list of index names to search; use *_all* or empty string to perform the operation on all indices

   - **doc_type** – A comma-separated list of document types to search; leave empty to perform the operation on all types

   - **body** – The search definition using the Query DSL

   - **_source** – True or false to return the _source field or not, or a list of fields to return

   - **_source_exclude** – A list of fields to exclude from the returned _source field

   - **_source_include** – A list of fields to extract and return from the _source field

   - **analyze_wildcard** – Specify whether wildcard and prefix queries should be analyzed (default: false)

   - **analyzer** – The analyzer to use for the query string

   - **default_operator** – The default operator for query string query (AND or OR) (default: OR)

   - **df** – The field to use as default where no field prefix is given in the query string

   - **explain** – Specify whether to return detailed information about score computation as part of a hit

   - **fields** – A comma-separated list of fields to return as part of a hit

- **indices_boost** – Comma-separated list of index boosts

- **lenient** – Specify whether format-based query failures (such as providing text to a numeric field) should be ignored

- **allow_no_indices** – Whether to ignore if a wildcard indices expression resolves into no concrete indices. (This includes _all string or when no indices have been specified)

- **expand_wildcards** – Whether to expand wildcard expression to concrete indices that are open, closed or both., default 'open'

- **ignore_unavailable** – Whether specified concrete indices should be ignored when unavailable (missing or closed)

- **lowercase_expanded_terms** – Specify whether query terms should be lowercased

- **from** – Starting offset (default: 0)

- **preference** – Specify the node or shard the operation should be performed on (default: random)

- **q** – Query in the Lucene query string syntax

- **routing** – A comma-separated list of specific routing values

- **scroll** – Specify how long a consistent view of the index should be maintained for scrolled search

- **search_type** – Search operation type

- **size** – Number of hits to return (default: 10)

- **sort** – A comma-separated list of <field>:<direction> pairs

- **source** – The URL-encoded request definition using the Query DSL (instead of using request body)

- **stats** – Specific 'tag' of the request for logging and statistical purposes

- **suggest_field** – Specify which field to use for suggestions

- **suggest_mode** – Specify suggest mode (default: missing)

- **suggest_size** – How many suggestions to return in response

- **suggest_text** – The source text for which the suggestions should be returned

- **timeout** – Explicit operation timeout

- **version** – Specify whether to return document version as part of a hit

search_shards(*args*, ***kwargs*)

The search shards api returns the indices and shards that a search request would be executed against. This can give useful feedback for working out issues or planning optimizations with routing and shard preferences. http://www.elasticsearch.org/guide/en/elasticsearch/reference/master/search-shards.html

Parameters

- **index** – The name of the index

- **doc_type** – The type of the document

- **allow_no_indices** – Whether to ignore if a wildcard indices expression resolves into no concrete indices. (This includes _all string or when no indices have been specified)

- **expand_wildcards** – Whether to expand wildcard expression to concrete indices that are open, closed or both. (default: "'open'")

- **ignore_unavailable** – Whether specified concrete indices should be ignored when unavailable (missing or closed)

- **local** – Return local information, do not retrieve the state from master node (default: false)

- **preference** – Specify the node or shard the operation should be performed on (default: random)

- **routing** – Specific routing value

search_template(**args*, ***kwargs*)

A query that accepts a query template and a map of key/value pairs to fill in template parameters. http://www.elasticsearch.org/guide/en/elasticsearch/reference/master/query-dsl-template-query.html

Parameters

- **index** – A comma-separated list of index names to search; use *_all* or empty string to perform the operation on all indices

- **doc_type** – A comma-separated list of document types to search; leave empty to perform the operation on all types

- **body** – The search definition template and its params

- **allow_no_indices** – Whether to ignore if a wildcard indices expression resolves into no concrete indices. (This includes *_all* string or when no indices have been specified)

- **expand_wildcards** – Whether to expand wildcard expression to concrete indices that are open, closed or both., default 'open'

- **ignore_unavailable** – Whether specified concrete indices should be ignored when unavailable (missing or closed)

- **preference** – Specify the node or shard the operation should be performed on (default: random)

- **routing** – A comma-separated list of specific routing values

- **scroll** – Specify how long a consistent view of the index should be maintained for scrolled search

- **search_type** – Search operation type

suggest(**args*, ***kwargs*)

The suggest feature suggests similar looking terms based on a provided text by using a suggester. http://www.elasticsearch.org/guide/en/elasticsearch/reference/current/search-search.html

Parameters

- **index** – A comma-separated list of index names to restrict the operation; use *_all* or empty string to perform the operation on all indices

- **body** – The request definition

- **allow_no_indices** – Whether to ignore if a wildcard indices expression resolves into no concrete indices. (This includes *_all* string or when no indices have been specified)

- **expand_wildcards** – Whether to expand wildcard expression to concrete indices that are open, closed or both., default 'open'

- **ignore_unavailable** – Whether specified concrete indices should be ignored when unavailable (missing or closed)

- **preference** – Specify the node or shard the operation should be performed on (default: random)

- **routing** – Specific routing value
- **source** – The URL-encoded request definition (instead of using request body)

termvector(**args*, ***kwargs*)

Added in 1. http://www.elasticsearch.org/guide/en/elasticsearch/reference/master/search-termvectors.html

> **Parameters**
>
> - **index** – The index in which the document resides.
> - **doc_type** – The type of the document.
> - **id** – The id of the document.
> - **body** – Define parameters. See documentation.
> - **field_statistics** – Specifies if document count, sum of document frequencies and sum of total term frequencies should be returned., default True
> - **fields** – A comma-separated list of fields to return.
> - **offsets** – Specifies if term offsets should be returned., default True
> - **parent** – Parent id of documents.
> - **payloads** – Specifies if term payloads should be returned., default True
> - **positions** – Specifies if term positions should be returned., default True
> - **preference** – Specify the node or shard the operation should be performed on (default: random).
> - **routing** – Specific routing value.
> - **term_statistics** – Specifies if total term frequency and document frequency should be returned., default False

update(**args*, ***kwargs*)

Update a document based on a script or partial data provided. http://www.elasticsearch.org/guide/en/elasticsearch/reference/current/docs-update.html

> **Parameters**
>
> - **index** – The name of the index
> - **doc_type** – The type of the document
> - **id** – Document ID
> - **body** – The request definition using either *script* or partial *doc*
> - **consistency** – Explicit write consistency setting for the operation
> - **fields** – A comma-separated list of fields to return in the response
> - **lang** – The script language (default: mvel)
> - **parent** – ID of the parent document
> - **refresh** – Refresh the index after performing the operation
> - **replication** – Specific replication type (default: sync)
> - **retry_on_conflict** – Specify how many times should the operation be retried when a conflict occurs (default: 0)
> - **routing** – Specific routing value

- **script** – The URL-encoded script definition (instead of using request body)

- **timeout** – Explicit operation timeout

- **timestamp** – Explicit timestamp for the document

- **ttl** – Expiration time for the document

- **version** – Explicit version number for concurrency control

- **version_type** – Explicit version number for concurrency control

4.1.2 Indices

class elasticsearch.client.**IndicesClient**(*client*)

> **analyze**(**args*, ***kwargs*)
> Perform the analysis process on a text and return the tokens breakdown of the text.
> http://www.elasticsearch.org/guide/en/elasticsearch/reference/current/indices-analyze.html
>
> > **Parameters**
> >
> > - **index** – The name of the index to scope the operation
> >
> > - **body** – The text on which the analysis should be performed
> >
> > - **analyzer** – The name of the analyzer to use
> >
> > - **char_filters** – A comma-separated list of character filters to use for the analysis
> >
> > - **field** – Use the analyzer configured for this field (instead of passing the analyzer name)
> >
> > - **filters** – A comma-separated list of filters to use for the analysis
> >
> > - **format** – Format of the output, default u'detailed'
> >
> > - **index** – The name of the index to scope the operation
> >
> > - **prefer_local** – With *true*, specify that a local shard should be used if available, with *false*, use a random shard (default: true)
> >
> > - **text** – The text on which the analysis should be performed (when request body is not used)
> >
> > - **tokenizer** – The name of the tokenizer to use for the analysis

> **clear_cache**(**args*, ***kwargs*)
> Clear either all caches or specific cached associated with one ore more indices.
> http://www.elasticsearch.org/guide/en/elasticsearch/reference/current/indices-clearcache.html
>
> > **Parameters**
> >
> > - **index** – A comma-separated list of index name to limit the operation
> >
> > - **field_data** – Clear field data
> >
> > - **fielddata** – Clear field data
> >
> > - **fields** – A comma-separated list of fields to clear when using the *field_data* parameter (default: all)
> >
> > - **filter** – Clear filter caches
> >
> > - **filter_cache** – Clear filter caches
> >
> > - **filter_keys** – A comma-separated list of keys to clear when using the *filter_cache* parameter (default: all)

- **id** – Clear ID caches for parent/child
- **id_cache** – Clear ID caches for parent/child
- **allow_no_indices** – Whether to ignore if a wildcard indices expression resolves into no concrete indices. (This includes _all string or when no indices have been specified)
- **expand_wildcards** – Whether to expand wildcard expression to concrete indices that are open, closed or both.
- **ignore_indices** – When performed on multiple indices, allows to ignore *missing* ones (default: none)
- **ignore_unavailable** – Whether specified concrete indices should be ignored when unavailable (missing or closed)
- **index** – A comma-separated list of index name to limit the operation
- **recycler** – Clear the recycler cache

close(**args*, ***kwargs*)
  Close an index to remove it's overhead from the cluster. Closed index is blocked for read/write operations. http://www.elasticsearch.org/guide/en/elasticsearch/reference/current/indices-open-close.html

  Parameters

- **index** – A comma-separated list of indices to close; use _all or '*' to close all indices
- **allow_no_indices** – Whether to ignore if a wildcard indices expression resolves into no concrete indices. (This includes _all string or when no indices have been specified)
- **expand_wildcards** – Whether to expand wildcard expression to concrete indices that are open, closed or both., default u'open'
- **ignore_unavailable** – Whether specified concrete indices should be ignored when unavailable (missing or closed)
- **master_timeout** – Specify timeout for connection to master
- **timeout** – Explicit operation timeout

create(**args*, ***kwargs*)
  Create an index in Elasticsearch. http://www.elasticsearch.org/guide/en/elasticsearch/reference/current/indices-create-index.html

  Parameters

- **index** – The name of the index
- **body** – The configuration for the index (*settings* and *mappings*)
- **master_timeout** – Specify timeout for connection to master
- **timeout** – Explicit operation timeout

delete(**args*, ***kwargs*)
  Delete an index in Elasticsearch http://www.elasticsearch.org/guide/en/elasticsearch/reference/current/indices-delete-index.html

  Parameters

- **index** – A comma-separated list of indices to delete; use _all or '*' to delete all indices
- **master_timeout** – Specify timeout for connection to master
- **timeout** – Explicit operation timeout

delete_alias(**args*, ***kwargs*)
 Delete specific alias.   http://www.elasticsearch.org/guide/en/elasticsearch/reference/current/indices-aliases.html

  Parameters

- **index** – A comma-separated list of index names (supports wildcards); use *_all* for all indices

- **name** – A comma-separated list of aliases to delete (supports wildcards); use *_all* to delete all aliases for the specified indices.

- **master_timeout** – Specify timeout for connection to master

- **timeout** – Explicit timestamp for the document

delete_mapping(**args*, ***kwargs*)
 Delete a mapping (type) along with its data. http://www.elasticsearch.org/guide/en/elasticsearch/reference/current/indices-delete-mapping.html

  Parameters

- **index** – A comma-separated list of index names (supports wildcard); use *_all* for all indices

- **doc_type** – A comma-separated list of document types to delete (supports wildcards); use *_all* to delete all document types in the specified indices.

- **master_timeout** – Specify timeout for connection to master

delete_template(**args*, ***kwargs*)
 Delete an index template by its name. http://www.elasticsearch.org/guide/en/elasticsearch/reference/current/indices-templates.html

  Parameters

- **name** – The name of the template

- **master_timeout** – Specify timeout for connection to master

- **timeout** – Explicit operation timeout

delete_warmer(**args*, ***kwargs*)
 Delete an index warmer.   http://www.elasticsearch.org/guide/en/elasticsearch/reference/current/indices-warmers.html

  Parameters

- **index** – A comma-separated list of index names to delete warmers from (supports wildcards); use *_all* to perform the operation on all indices.

- **name** – A comma-separated list of warmer names to delete (supports wildcards); use *_all* to delete all warmers in the specified indices.

- **master_timeout** – Specify timeout for connection to master

exists(**args*, ***kwargs*)
 Return a boolean indicating whether given index exists. http://www.elasticsearch.org/guide/en/elasticsearch/reference/current/indices-exists.html

  Parameters

- **index** – A list of indices to check

- **allow_no_indices** – Whether to ignore if a wildcard indices expression resolves into no concrete indices. (This includes *_all* string or when no indices have been specified)

- **expand_wildcards** – Whether to expand wildcard expression to concrete indices that are open, closed or both., default u'open'

- **ignore_unavailable** – Whether specified concrete indices should be ignored when unavailable (missing or closed)

- **local** – Return local information, do not retrieve the state from master node (default: false)

exists_alias(**args*, ***kwargs*)

Return a boolean indicating whether given alias exists. http://www.elasticsearch.org/guide/en/elasticsearch/reference/current/ aliases.html

Parameters

- **name** – A comma-separated list of alias names to return

- **index** – A comma-separated list of index names to filter aliases

- **allow_no_indices** – Whether to ignore if a wildcard indices expression resolves into no concrete indices. (This includes _all string or when no indices have been specified)

- **expand_wildcards** – Whether to expand wildcard expression to concrete indices that are open, closed or both.

- **ignore_indices** – When performed on multiple indices, allows to ignore *missing* ones (default: none)

- **ignore_unavailable** – Whether specified concrete indices should be ignored when unavailable (missing or closed)

- **local** – Return local information, do not retrieve the state from master node (default: false)

exists_template(**args*, ***kwargs*)

Return a boolean indicating whether given template exists. http://www.elasticsearch.org/guide/en/elasticsearch/reference/cur templates.html

Parameters

- **name** – The name of the template

- **local** – Return local information, do not retrieve the state from master node (default: false)

exists_type(**args*, ***kwargs*)

Check if a type/types exists in an index/indices. http://www.elasticsearch.org/guide/en/elasticsearch/reference/current/indices types-exists.html

Parameters

- **index** – A comma-separated list of index names; use _all to check the types across all indices

- **doc_type** – A comma-separated list of document types to check

- **allow_no_indices** – Whether to ignore if a wildcard indices expression resolves into no concrete indices. (This includes _all string or when no indices have been specified)

- **expand_wildcards** – Whether to expand wildcard expression to concrete indices that are open, closed or both.

- **ignore_indices** – When performed on multiple indices, allows to ignore *missing* ones (default: none)

- **ignore_unavailable** – Whether specified concrete indices should be ignored when unavailable (missing or closed)

- **local** – Return local information, do not retrieve the state from master node (default: false)

flush(**args*, ***kwargs*)
    Explicitly flush one or more indices. http://www.elasticsearch.org/guide/en/elasticsearch/reference/current/indices-flush.html

        Parameters

- **index** – A comma-separated list of index names; use *_all* or empty string for all indices
- **force** – Whether a flush should be forced even if it is not necessarily needed ie. if no changes will be committed to the index.
- **full** – If set to true a new index writer is created and settings that have been changed related to the index writer will be refreshed.
- **allow_no_indices** – Whether to ignore if a wildcard indices expression resolves into no concrete indices. (This includes *_all* string or when no indices have been specified)
- **expand_wildcards** – Whether to expand wildcard expression to concrete indices that are open, closed or both.
- **ignore_indices** – When performed on multiple indices, allows to ignore *missing* ones (default: none)
- **ignore_unavailable** – Whether specified concrete indices should be ignored when unavailable (missing or closed)

get_alias(**args*, ***kwargs*)
    Retrieve a specified alias. http://www.elasticsearch.org/guide/en/elasticsearch/reference/current/indices-aliases.html

        Parameters

- **name** – A comma-separated list of alias names to return
- **index** – A comma-separated list of index names to filter aliases
- **allow_no_indices** – Whether to ignore if a wildcard indices expression resolves into no concrete indices. (This includes *_all* string or when no indices have been specified)
- **expand_wildcards** – Whether to expand wildcard expression to concrete indices that are open, closed or both.
- **ignore_indices** – When performed on multiple indices, allows to ignore *missing* ones, default u'none'
- **ignore_unavailable** – Whether specified concrete indices should be ignored when unavailable (missing or closed)
- **local** – Return local information, do not retrieve the state from master node (default: false)

get_aliases(**args*, ***kwargs*)
    Retrieve specified aliases http://www.elasticsearch.org/guide/en/elasticsearch/reference/current/indices-aliases.html

        Parameters

- **index** – A comma-separated list of index names to filter aliases
- **name** – A comma-separated list of alias names to filter
- **local** – Return local information, do not retrieve the state from master node (default: false)
- **timeout** – Explicit operation timeout

get_field_mapping(**args*, ***kwargs*)

Retrieve mapping definition of a specific field. http://www.elasticsearch.org/guide/en/elasticsearch/reference/current/indices-get-field-mapping.html

> **Parameters**
>
> - **index** – A comma-separated list of index names; use *_all* or empty string for all indices
> - **doc_type** – A comma-separated list of document types
> - **field** – A comma-separated list of fields to retrieve the mapping for
> - **include_defaults** – A boolean indicating whether to return default values
> - **allow_no_indices** – Whether to ignore if a wildcard indices expression resolves into no concrete indices. (This includes *_all* string or when no indices have been specified)
> - **expand_wildcards** – Whether to expand wildcard expression to concrete indices that are open, closed or both.
> - **ignore_unavailable** – Whether specified concrete indices should be ignored when unavailable (missing or closed)
> - **local** – Return local information, do not retrieve the state from master node (default: false)

get_mapping(**args*, ***kwargs*)

Retrieve mapping definition of index or index/type. http://www.elasticsearch.org/guide/en/elasticsearch/reference/current/indices-get-mapping.html

> **Parameters**
>
> - **index** – A comma-separated list of index names; use *_all* or empty string for all indices
> - **doc_type** – A comma-separated list of document types
> - **allow_no_indices** – Whether to ignore if a wildcard indices expression resolves into no concrete indices. (This includes *_all* string or when no indices have been specified)
> - **expand_wildcards** – Whether to expand wildcard expression to concrete indices that are open, closed or both.
> - **ignore_unavailable** – Whether specified concrete indices should be ignored when unavailable (missing or closed)
> - **local** – Return local information, do not retrieve the state from master node (default: false)

get_settings(**args*, ***kwargs*)

Retrieve settings for one or more (or all) indices. http://www.elasticsearch.org/guide/en/elasticsearch/reference/current/indices-get-settings.html

> **Parameters**
>
> - **index** – A comma-separated list of index names; use *_all* or empty string to perform the operation on all indices
> - **name** – The name of the settings that should be included
> - **expand_wildcards** – Whether to expand wildcard expression to concrete indices that are open, closed or both.
> - **ignore_indices** – When performed on multiple indices, allows to ignore *missing* ones, default u'none'
> - **ignore_unavailable** – Whether specified concrete indices should be ignored when unavailable (missing or closed)

- **flat_settings** – Return settings in flat format (default: false)

- **local** – Return local information, do not retrieve the state from master node (default: false)

get_template(**args*, ***kwargs*)

Retrieve an index template by its name. http://www.elasticsearch.org/guide/en/elasticsearch/reference/current/indices-templates.html

Parameters

- **name** – The name of the template

- **flat_settings** – Return settings in flat format (default: false)

- **local** – Return local information, do not retrieve the state from master node (default: false)

get_warmer(**args*, ***kwargs*)

Retreieve an index warmer. http://www.elasticsearch.org/guide/en/elasticsearch/reference/current/indices-warmers.html

Parameters

- **index** – A comma-separated list of index names to restrict the operation; use *_all* to perform the operation on all indices

- **doc_type** – A comma-separated list of document types to restrict the operation; leave empty to perform the operation on all types

- **name** – The name of the warmer (supports wildcards); leave empty to get all warmers

- **allow_no_indices** – Whether to ignore if a wildcard indices expression resolves into no concrete indices. (This includes *_all* string or when no indices have been specified)

- **expand_wildcards** – Whether to expand wildcard expression to concrete indices that are open, closed or both., default u'open'

- **ignore_unavailable** – Whether specified concrete indices should be ignored when unavailable (missing or closed)

- **local** – Return local information, do not retrieve the state from master node (default: false)

open(**args*, ***kwargs*)

Open a closed index to make it available for search. http://www.elasticsearch.org/guide/en/elasticsearch/reference/current/ind open-close.html

Parameters

- **index** – The name of the index

- **master_timeout** – Specify timeout for connection to master

- **timeout** – Explicit operation timeout

- **allow_no_indices** – Whether to ignore if a wildcard indices expression resolves into no concrete indices. (This includes *_all* string or when no indices have been specified)

- **expand_wildcards** – Whether to expand wildcard expression to concrete indices that are open, closed or both.

- **ignore_unavailable** – Whether specified concrete indices should be ignored when unavailable (missing or closed)

optimize(**args*, ***kwargs*)

Explicitly optimize one or more indices through an API. http://www.elasticsearch.org/guide/en/elasticsearch/reference/curren optimize.html

Parameters

- **index** – A comma-separated list of index names; use *_all* or empty string to perform the operation on all indices

- **flush** – Specify whether the index should be flushed after performing the operation (default: true)

- **allow_no_indices** – Whether to ignore if a wildcard indices expression resolves into no concrete indices. (This includes *_all* string or when no indices have been specified)

- **expand_wildcards** – Whether to expand wildcard expression to concrete indices that are open, closed or both.

- **ignore_indices** – When performed on multiple indices, allows to ignore *missing* ones, default u'none'

- **ignore_unavailable** – Whether specified concrete indices should be ignored when unavailable (missing or closed)

- **max_num_segments** – The number of segments the index should be merged into (default: dynamic)

- **only_expunge_deletes** – Specify whether the operation should only expunge deleted documents

- **operation_threading** – TODO: ?

- **wait_for_merge** – Specify whether the request should block until the merge process is finished (default: true)

put_alias(**args*, ***kwargs*)

Create an alias for a specific index/indices. http://www.elasticsearch.org/guide/en/elasticsearch/reference/current/indices-aliases.html

Parameters

- **index** – A comma-separated list of index names the alias should point to (supports wildcards); use *_all* or omit to perform the operation on all indices.

- **name** – The name of the alias to be created or updated

- **body** – The settings for the alias, such as *routing* or *filter*

- **master_timeout** – Specify timeout for connection to master

- **timeout** – Explicit timestamp for the document

put_mapping(**args*, ***kwargs*)

Register specific mapping definition for a specific type. http://www.elasticsearch.org/guide/en/elasticsearch/reference/current/put-mapping.html

Parameters

- **index** – A comma-separated list of index names the alias should point to (supports wildcards); use *_all* or omit to perform the operation on all indices.

- **doc_type** – The name of the document type

- **body** – The mapping definition

- **allow_no_indices** – Whether to ignore if a wildcard indices expression resolves into no concrete indices. (This includes *_all* string or when no indices have been specified)

- **expand_wildcards** – Whether to expand wildcard expression to concrete indices that are open, closed or both., default u'open'

- **ignore_conflicts** – Specify whether to ignore conflicts while updating the mapping (default: false)

- **ignore_unavailable** – Whether specified concrete indices should be ignored when unavailable (missing or closed)

- **master_timeout** – Specify timeout for connection to master

- **timeout** – Explicit operation timeout

put_settings(*args*, **kwargs*)

Change specific index level settings in real time. http://www.elasticsearch.org/guide/en/elasticsearch/reference/current/indice update-settings.html

Parameters

- **body** – The index settings to be updated

- **index** – A comma-separated list of index names; use *_all* or empty string to perform the operation on all indices

- **allow_no_indices** – Whether to ignore if a wildcard indices expression resolves into no concrete indices. (This includes *_all* string or when no indices have been specified)

- **expand_wildcards** – Whether to expand wildcard expression to concrete indices that are open, closed or both., default u'open'

- **flat_settings** – Return settings in flat format (default: false)

- **ignore_unavailable** – Whether specified concrete indices should be ignored when unavailable (missing or closed)

- **master_timeout** – Specify timeout for connection to master

put_template(*args*, **kwargs*)

Create an index template that will automatically be applied to new indices created. http://www.elasticsearch.org/guide/en/elasticsearch/reference/current/indices-templates.html

Parameters

- **name** – The name of the template

- **body** – The template definition

- **create** – Whether the index template should only be added if new or can also replace an existing one

- **order** – The order for this template when merging multiple matching ones (higher numbers are merged later, overriding the lower numbers)

- **master_timeout** – Specify timeout for connection to master

- **timeout** – Explicit operation timeout

- **flat_settings** – Return settings in flat format (default: false)

put_warmer(*args*, **kwargs*)

Create an index warmer to run registered search requests to warm up the index before it is available for search. http://www.elasticsearch.org/guide/en/elasticsearch/reference/current/indices-warmers.html

Parameters

- **name** – The name of the warmer

- **body** – The search request definition for the warmer (query, filters, facets, sorting, etc)

- **index** – A comma-separated list of index names to register the warmer for; use _all or omit to perform the operation on all indices

- **doc_type** – A comma-separated list of document types to register the warmer for; leave empty to perform the operation on all types

- **allow_no_indices** – Whether to ignore if a wildcard indices expression resolves into no concrete indices in the search request to warm. (This includes _all string or when no indices have been specified)

- **expand_wildcards** – Whether to expand wildcard expression to concrete indices that are open, closed or both, in the search request to warm., default u'open'

- **ignore_unavailable** – Whether specified concrete indices should be ignored when unavailable (missing or closed) in the search request to warm

- **master_timeout** – Specify timeout for connection to master

recovery(*args, **kwargs*)
The indices recovery API provides insight into on-going shard recoveries. Recovery status may be reported for specific indices, or cluster-wide. http://www.elasticsearch.org/guide/en/elasticsearch/reference/master/indices-recovery.html

   Parameters

- **index** – A comma-separated list of index names; use _all or empty string to perform the operation on all indices

- **active_only** – Display only those recoveries that are currently on- going (default: 'false')

- **detailed** – Whether to display detailed information about shard recovery (default: 'false')

- **human** – Whether to return time and byte values in human-readable format. (default: 'false')

refresh(*args, **kwargs*)
Explicitly refresh one or more index, making all operations performed since the last refresh available for search. http://www.elasticsearch.org/guide/en/elasticsearch/reference/current/indices-refresh.html

   Parameters

- **index** – A comma-separated list of index names; use _all or empty string to perform the operation on all indices

- **allow_no_indices** – Whether to ignore if a wildcard indices expression resolves into no concrete indices. (This includes _all string or when no indices have been specified)

- **expand_wildcards** – Whether to expand wildcard expression to concrete indices that are open, closed or both.

- **ignore_indices** – When performed on multiple indices, allows to ignore *missing* ones, default u'none'

- **ignore_unavailable** – Whether specified concrete indices should be ignored when unavailable (missing or closed)

- **force** – Force a refresh even if not required

segments(*args, **kwargs*)
Provide low level segments information that a Lucene index (shard level) is built with. http://www.elasticsearch.org/guide/en/elasticsearch/reference/current/indices-segments.html

   Parameters

- **index** – A comma-separated list of index names; use *_all* or empty string to perform the operation on all indices

- **allow_no_indices** – Whether to ignore if a wildcard indices expression resolves into no concrete indices. (This includes *_all* string or when no indices have been specified)

- **expand_wildcards** – Whether to expand wildcard expression to concrete indices that are open, closed or both.

- **ignore_indices** – When performed on multiple indices, allows to ignore *missing* ones, default u'none'

- **ignore_unavailable** – Whether specified concrete indices should be ignored when unavailable (missing or closed)

- **human** – Whether to return time and byte values in human-readable format (default: false)

snapshot_index(**args*, ***kwargs*)
Explicitly perform a snapshot through the gateway of one or more indices (backup them). http://www.elasticsearch.org/guide/en/elasticsearch/reference/current/indices-gateway-snapshot.html

Parameters

- **index** – A comma-separated list of index names; use *_all* or empty string for all indices

- **allow_no_indices** – Whether to ignore if a wildcard indices expression resolves into no concrete indices. (This includes *_all* string or when no indices have been specified)

- **expand_wildcards** – Whether to expand wildcard expression to concrete indices that are open, closed or both.

- **ignore_indices** – When performed on multiple indices, allows to ignore *missing* ones (default: none)

- **ignore_unavailable** – Whether specified concrete indices should be ignored when unavailable (missing or closed)

stats(**args*, ***kwargs*)
Retrieve statistics on different operations happening on an index. http://www.elasticsearch.org/guide/en/elasticsearch/reference/current/indices-stats.html

Parameters

- **index** – A comma-separated list of index names; use *_all* or empty string to perform the operation on all indices

- **metric** – A comma-separated list of metrics to display. Possible values: "_all", "completion", "docs", "fielddata", "filter_cache", "flush", "get", "id_cache", "indexing", "merge", "percolate", "refresh", "search", "segments", "store", "warmer"

- **completion_fields** – A comma-separated list of fields for *completion* metric (supports wildcards)

- **fielddata_fields** – A comma-separated list of fields for *fielddata* metric (supports wildcards)

- **fields** – A comma-separated list of fields for *fielddata* and *completion* metric (supports wildcards)

- **groups** – A comma-separated list of search groups for *search* statistics

- **allow_no_indices** – Whether to ignore if a wildcard indices expression resolves into no concrete indices. (This includes *_all* string or when no indices have been specified)

- **expand_wildcards** – Whether to expand wildcard expression to concrete indices that are open, closed or both.

- **ignore_indices** – When performed on multiple indices, allows to ignore *missing* ones (default: none)

- **ignore_unavailable** – Whether specified concrete indices should be ignored when unavailable (missing or closed)

- **human** – Whether to return time and byte values in human-readable format.

- **level** – Return stats aggregated at cluster, index or shard level. ("cluster", "indices" or "shards", default: "indices")

- **types** – A comma-separated list of document types for the *indexing* index metric

status(*args*, **kwargs*)

Get a comprehensive status information of one or more indices. http://elasticsearch.org/guide/reference/api/admin-indices-_/

Parameters

- **index** – A comma-separated list of index names; use *_all* or empty string to perform the operation on all indices

- **allow_no_indices** – Whether to ignore if a wildcard indices expression resolves into no concrete indices. (This includes *_all* string or when no indices have been specified)

- **expand_wildcards** – Whether to expand wildcard expression to concrete indices that are open, closed or both.

- **ignore_indices** – When performed on multiple indices, allows to ignore *missing* ones, default u'none'

- **ignore_unavailable** – Whether specified concrete indices should be ignored when unavailable (missing or closed)

- **operation_threading** – TODO: ?

- **recovery** – Return information about shard recovery

- **snapshot** – TODO: ?

- **human** – Whether to return time and byte values in human-readable format.

update_aliases(*args*, **kwargs*)

Update specified aliases. http://www.elasticsearch.org/guide/en/elasticsearch/reference/current/indices-aliases.html

Parameters

- **body** – The definition of *actions* to perform

- **master_timeout** – Specify timeout for connection to master

- **timeout** – Request timeout

validate_query(*args*, **kwargs*)

Validate a potentially expensive query without executing it. http://www.elasticsearch.org/guide/en/elasticsearch/reference/cur validate.html

Parameters

- **index** – A comma-separated list of index names to restrict the operation; use *_all* or empty string to perform the operation on all indices

- **doc_type** – A comma-separated list of document types to restrict the operation; leave empty to perform the operation on all types

- **body** – The query definition

- **explain** – Return detailed information about the error

- **allow_no_indices** – Whether to ignore if a wildcard indices expression resolves into no concrete indices. (This includes _all string or when no indices have been specified)

- **expand_wildcards** – Whether to expand wildcard expression to concrete indices that are open, closed or both.

- **ignore_indices** – When performed on multiple indices, allows to ignore *missing* ones (default: none)

- **ignore_unavailable** – Whether specified concrete indices should be ignored when unavailable (missing or closed)

- **operation_threading** – TODO: ?

- **q** – Query in the Lucene query string syntax

- **source** – The URL-encoded query definition (instead of using the request body)

4.1.3 Cluster

class elasticsearch.client.**ClusterClient**(*client*)

> **get_settings**(**args, **kwargs*)
>> Get cluster settings. http://www.elasticsearch.org/guide/en/elasticsearch/reference/current/cluster-update-settings.html
>>
>>> Parameters
>>>
>>> - **flat_settings** – Return settings in flat format (default: false)
>>>
>>> - **master_timeout** – Explicit operation timeout for connection to master node
>>>
>>> - **timeout** – Explicit operation timeout

> **health**(**args, **kwargs*)
>> Get a very simple status on the health of the cluster. http://www.elasticsearch.org/guide/en/elasticsearch/reference/current/clu health.html
>>
>>> Parameters
>>>
>>> - **index** – Limit the information returned to a specific index
>>>
>>> - **level** – Specify the level of detail for returned information, default u'cluster'
>>>
>>> - **local** – Return local information, do not retrieve the state from master node (default: false)
>>>
>>> - **master_timeout** – Explicit operation timeout for connection to master node
>>>
>>> - **timeout** – Explicit operation timeout
>>>
>>> - **wait_for_active_shards** – Wait until the specified number of shards is active
>>>
>>> - **wait_for_nodes** – Wait until the specified number of nodes is available
>>>
>>> - **wait_for_relocating_shards** – Wait until the specified number of relocating shards is finished
>>>
>>> - **wait_for_status** – Wait until cluster is in a specific state, default None

pending_tasks(**args*, ***kwargs*)
The pending cluster tasks API returns a list of any cluster-level changes (e.g. create index, update mapping, allocate or fail shard) which have not yet been executed. http://www.elasticsearch.org/guide/en/elasticsearch/reference/current/cluster-pending.html

> **Parameters**
>
> - **local** – Return local information, do not retrieve the state from master node (default: false)
> - **master_timeout** – Specify timeout for connection to master

put_settings(**args*, ***kwargs*)
Update cluster wide specific settings. http://www.elasticsearch.org/guide/en/elasticsearch/reference/current/cluster-update-settings.html

> **Parameters**
>
> - **body** – The settings to be updated. Can be either *transient* or *persistent* (survives cluster restart).
> - **flat_settings** – Return settings in flat format (default: false)

reroute(**args*, ***kwargs*)
Explicitly execute a cluster reroute allocation command including specific commands. http://www.elasticsearch.org/guide/en/elasticsearch/reference/current/cluster-reroute.html

> **Parameters**
>
> - **body** – The definition of *commands* to perform (*move*, *cancel*, *allocate*)
> - **dry_run** – Simulate the operation only and return the resulting state
> - **explain** – Return an explanation of why the commands can or cannot be executed
> - **filter_metadata** – Don't return cluster state metadata (default: false)
> - **master_timeout** – Explicit operation timeout for connection to master node
> - **timeout** – Explicit operation timeout

state(**args*, ***kwargs*)
Get a comprehensive state information of the whole cluster. http://www.elasticsearch.org/guide/en/elasticsearch/reference/current/cluster-state.html

> **Parameters**
>
> - **metric** – Limit the information returned to the specified metrics. Possible values: "_all", "blocks", "index_templates", "metadata", "nodes", "routing_table", "master_node", "version"
> - **index** – A comma-separated list of index names; use *_all* or empty string to perform the operation on all indices
> - **index_templates** – A comma separated list to return specific index templates when returning metadata.
> - **local** – Return local information, do not retrieve the state from master node (default: false)
> - **master_timeout** – Specify timeout for connection to master
> - **flat_settings** – Return settings in flat format (default: false)

stats(**args*, ***kwargs*)
The Cluster Stats API allows to retrieve statistics from a cluster wide perspective. The API returns basic index metrics and information about the current nodes that form the cluster. http://www.elasticsearch.org/guide/en/elasticsearch/reference/current/cluster-stats.html

Parameters

- **node_id** – A comma-separated list of node IDs or names to limit the returned information; use *_local* to return information from the node you're connecting to, leave empty to get information from all nodes

- **flat_settings** – Return settings in flat format (default: false)

- **human** – Whether to return time and byte values in human-readable format.

4.1.4 Nodes

class elasticsearch.client.**NodesClient**(*client*)

hot_threads(**args*, ***kwargs*)
An API allowing to get the current hot threads on each node in the cluster. http://www.elasticsearch.org/guide/en/elasticsearch/reference/current/cluster-nodes-hot-threads.html

Parameters

- **node_id** – A comma-separated list of node IDs or names to limit the returned information; use *_local* to return information from the node you're connecting to, leave empty to get information from all nodes

- **type** – The type to sample (default: cpu)

- **interval** – The interval for the second sampling of threads

- **snapshots** – Number of samples of thread stacktrace (default: 10)

- **threads** – Specify the number of threads to provide information for (default: 3)

info(**args*, ***kwargs*)
The cluster nodes info API allows to retrieve one or more (or all) of the cluster nodes information. http://www.elasticsearch.org/guide/en/elasticsearch/reference/current/cluster-nodes-info.html

Parameters

- **node_id** – A comma-separated list of node IDs or names to limit the returned information; use *_local* to return information from the node you're connecting to, leave empty to get information from all nodes

- **metric** – A comma-separated list of metrics you wish returned. Leave empty to return all. Choices are "settings", "os", "process", "jvm", "thread_pool", "network", "transport", "http", "plugin"

- **flat_settings** – Return settings in flat format (default: false)

- **human** – Whether to return time and byte values in human-readable format., default False

shutdown(**args*, ***kwargs*)
The nodes shutdown API allows to shutdown one or more (or all) nodes in the cluster. http://www.elasticsearch.org/guide/en/elasticsearch/reference/current/cluster-nodes-shutdown.html

Parameters

- **node_id** – A comma-separated list of node IDs or names to perform the operation on; use *_local* to perform the operation on the node you're connected to, leave empty to perform the operation on all nodes

- **delay** – Set the delay for the operation (default: 1s)

- **exit** – Exit the JVM as well (default: true)

stats(**args*, ***kwargs*)
> The cluster nodes stats API allows to retrieve one or more (or all) of the cluster nodes statistics.
> http://www.elasticsearch.org/guide/en/elasticsearch/reference/current/cluster-nodes-stats.html

> **Parameters**
>
> - **node_id** – A comma-separated list of node IDs or names to limit the returned information; use *_local* to return information from the node you're connecting to, leave empty to get information from all nodes
> - **metric** – Limit the information returned to the specified metrics. Possible options are: "_all", "breaker", "fs", "http", "indices", "jvm", "network", "os", "process", "thread_pool", "transport"
> - **index_metric** – Limit the information returned for *indices* metric to the specific index metrics. Isn't used if *indices* (or *all*) metric isn't specified. Possible options are: "_all", "completion", "docs", "fielddata", "filter_cache", "flush", "get", "id_cache", "indexing", "merge", "percolate", "refresh", "search", "segments", "store", "warmer"
> - **completion_fields** – A comma-separated list of fields for *fielddata* and *suggest* index metric (supports wildcards)
> - **fielddata_fields** – A comma-separated list of fields for *fielddata* index metric (supports wildcards)
> - **fields** – A comma-separated list of fields for *fielddata* and *completion* index metric (supports wildcards)
> - **groups** – A comma-separated list of search groups for *search* index metric
> - **human** – Whether to return time and byte values in human-readable format., default False
> - **level** – Return indices stats aggregated at node, index or shard level, default 'node'
> - **types** – A comma-separated list of document types for the *indexing* index metric

4.1.5 Cat

class elasticsearch.client.**CatClient**(*client*)

aliases(**args*, ***kwargs*)
> http://www.elasticsearch.org/guide/en/elasticsearch/reference/master/cat-aliases.html

> **Parameters**
>
> - **name** – A comma-separated list of alias names to return
> - **h** – Comma-separated list of column names to display
> - **help** – Return help information, default False
> - **local** – Return local information, do not retrieve the state from master node (default: false)
> - **master_timeout** – Explicit operation timeout for connection to master node
> - **v** – Verbose mode. Display column headers, default False

allocation(**args*, ***kwargs*)
> Allocation provides a snapshot of how shards have located around the cluster and the state of disk usage.
> http://www.elasticsearch.org/guide/en/elasticsearch/reference/master/cat-allocation.html

> **Parameters**

- **node_id** – A comma-separated list of node IDs or names to limit the returned information
- **bytes** – The unit in which to display byte values
- **h** – Comma-separated list of column names to display
- **help** – Return help information, default False
- **local** – Return local information, do not retrieve the state from master node (default: false)
- **master_timeout** – Explicit operation timeout for connection to master node
- **v** – Verbose mode. Display column headers, default False

count(**args*, ***kwargs*)
> Count provides quick access to the document count of the entire cluster, or individual indices.
> http://www.elasticsearch.org/guide/en/elasticsearch/reference/master/cat-count.html

> **Parameters**
>
> - **index** – A comma-separated list of index names to limit the returned information
> - **h** – Comma-separated list of column names to display
> - **help** – Return help information, default False
> - **local** – Return local information, do not retrieve the state from master node (default: false)
> - **master_timeout** – Explicit operation timeout for connection to master node
> - **v** – Verbose mode. Display column headers, default False

fielddata(**args*, ***kwargs*)
> Shows information about currently loaded fielddata on a per-node basis.
> http://www.elasticsearch.org/guide/en/elasticsearch/reference/master/cat-fielddata.html

> **Parameters**
>
> - **fields** – A comma-separated list of fields to return the fielddata size
> - **bytes** – The unit in which to display byte values
> - **fields** – A comma-separated list of fields to return the fielddata size
> - **h** – Comma-separated list of column names to display
> - **help** – Return help information (default: 'false')
> - **local** – Return local information, do not retrieve the state from master node (default: false)
> - **master_timeout** – Explicit operation timeout for connection to master node
> - **v** – Verbose mode. Display column headers (default: 'false')

health(**args*, ***kwargs*)
> health is a terse, one-line representation of the same information from `health()` API
> http://www.elasticsearch.org/guide/en/elasticsearch/reference/master/cat-health.html

> **Parameters**
>
> - **h** – Comma-separated list of column names to display
> - **help** – Return help information, default False
> - **local** – Return local information, do not retrieve the state from master node (default: false)
> - **master_timeout** – Explicit operation timeout for connection to master node
> - **ts** – Set to false to disable timestamping, default True

- **v** – Verbose mode. Display column headers, default False

help(*args*, **kwargs*)

A simple help for the cat api. http://www.elasticsearch.org/guide/en/elasticsearch/reference/master/cat.html

> **Parameters help** – Return help information, default False

indices(*args*, **kwargs*)

The indices command provides a cross-section of each index. http://www.elasticsearch.org/guide/en/elasticsearch/reference/master/cat-indices.html

> **Parameters**
>
> - **index** – A comma-separated list of index names to limit the returned information
> - **bytes** – The unit in which to display byte values
> - **h** – Comma-separated list of column names to display
> - **help** – Return help information, default False
> - **local** – Return local information, do not retrieve the state from master node (default: false)
> - **master_timeout** – Explicit operation timeout for connection to master node
> - **pri** – Set to true to return stats only for primary shards, default False
> - **v** – Verbose mode. Display column headers, default False

master(*args*, **kwargs*)

Displays the master's node ID, bound IP address, and node name. http://www.elasticsearch.org/guide/en/elasticsearch/reference/master/cat-master.html

> **Parameters**
>
> - **h** – Comma-separated list of column names to display
> - **help** – Return help information, default False
> - **local** – Return local information, do not retrieve the state from master node (default: false)
> - **master_timeout** – Explicit operation timeout for connection to master node
> - **v** – Verbose mode. Display column headers, default False

nodes(*args*, **kwargs*)

The nodes command shows the cluster topology. http://www.elasticsearch.org/guide/en/elasticsearch/reference/master/cat-nodes.html

> **Parameters**
>
> - **h** – Comma-separated list of column names to display
> - **help** – Return help information, default False
> - **local** – Return local information, do not retrieve the state from master node (default: false)
> - **master_timeout** – Explicit operation timeout for connection to master node
> - **v** – Verbose mode. Display column headers, default False

pending_tasks(*args*, **kwargs*)

pending_tasks provides the same information as the `pending_tasks()` API in a convenient tabular format. http://www.elasticsearch.org/guide/en/elasticsearch/reference/master/cat-pending-tasks.html

> **Parameters**
>
> - **h** – Comma-separated list of column names to display

- **help** – Return help information, default False
- **local** – Return local information, do not retrieve the state from master node (default: false)
- **master_timeout** – Explicit operation timeout for connection to master node
- **v** – Verbose mode. Display column headers, default False

plugins(**args*, ***kwargs*)
  http://www.elasticsearch.org/guide/en/elasticsearch/reference/current/cat-plugins.html

  Parameters

  - **h** – Comma-separated list of column names to display
  - **help** – Return help information, default False
  - **local** – Return local information, do not retrieve the state from master node (default: false)
  - **master_timeout** – Explicit operation timeout for connection to master node
  - **v** – Verbose mode. Display column headers, default False

recovery(**args*, ***kwargs*)
  recovery is a view of shard replication. http://www.elasticsearch.org/guide/en/elasticsearch/reference/master/cat-recovery.html

  Parameters

  - **index** – A comma-separated list of index names to limit the returned information
  - **bytes** – The unit in which to display byte values
  - **h** – Comma-separated list of column names to display
  - **help** – Return help information, default False
  - **local** – Return local information, do not retrieve the state from master node (default: false)
  - **master_timeout** – Explicit operation timeout for connection to master node
  - **v** – Verbose mode. Display column headers, default False

segments(**args*, ***kwargs*)
  The segments command is the detailed view of Lucene segments per index.

  Parameters

  - **index** – A comma-separated list of index names to limit the returned information
  - **h** – Comma-separated list of column names to display
  - **help** – Return help information, default False
  - **local** – Return local information, do not retrieve the state from master node (default: false)
  - **master_timeout** – Explicit operation timeout for connection to master node
  - **v** – Verbose mode. Display column headers, default False

shards(**args*, ***kwargs*)
  The shards command is the detailed view of what nodes contain which shards.
  http://www.elasticsearch.org/guide/en/elasticsearch/reference/master/cat-shards.html

  Parameters

  - **index** – A comma-separated list of index names to limit the returned information
  - **h** – Comma-separated list of column names to display

- **help** – Return help information, default False

- **local** – Return local information, do not retrieve the state from master node (default: false)

- **master_timeout** – Explicit operation timeout for connection to master node

- **v** – Verbose mode. Display column headers, default False

thread_pool(**args*, ***kwargs*)
    Get information about thread pools. http://www.elasticsearch.org/guide/en/elasticsearch/reference/master/cat-thread-pool.html

    Parameters

- **full_id** – Enables displaying the complete node ids (default: 'false')

- **h** – Comma-separated list of column names to display

- **help** – Return help information (default: 'false')

- **local** – Return local information, do not retrieve the state from master node (default: false)

- **master_timeout** – Explicit operation timeout for connection to master node

- **v** – Verbose mode. Display column headers (default: 'false')

Snapshot —

class elasticsearch.client.**SnapshotClient**(*client*)

create(**args*, ***kwargs*)
    Create a snapshot in repository http://www.elasticsearch.org/guide/en/elasticsearch/reference/master/modules-snapshots.html

    Parameters

- **repository** – A repository name

- **snapshot** – A snapshot name

- **body** – The snapshot definition

- **master_timeout** – Explicit operation timeout for connection to master node

- **wait_for_completion** – Should this request wait until the operation has completed before returning, default False

create_repository(**args*, ***kwargs*)
    Registers a shared file system repository. http://www.elasticsearch.org/guide/en/elasticsearch/reference/master/modules-snapshots.html

    Parameters

- **repository** – A repository name

- **body** – The repository definition

- **master_timeout** – Explicit operation timeout for connection to master node

- **timeout** – Explicit operation timeout

delete(**args*, ***kwargs*)
    Deletes a snapshot from a repository. http://www.elasticsearch.org/guide/en/elasticsearch/reference/master/modules-snapshots.html

    Parameters

- **repository** – A repository name

- **snapshot** – A snapshot name

- **master_timeout** – Explicit operation timeout for connection to master node

delete_repository(**args*, ***kwargs*)
  Removes a shared file system repository. http://www.elasticsearch.org/guide/en/elasticsearch/reference/master/modules-snapshots.html

  Parameters

  - **repository** – A comma-separated list of repository names

  - **master_timeout** – Explicit operation timeout for connection to master node

  - **timeout** – Explicit operation timeout

get(**args*, ***kwargs*)
  Retrieve information about a snapshot. http://www.elasticsearch.org/guide/en/elasticsearch/reference/master/modules-snapshots.html

  Parameters

  - **repository** – A comma-separated list of repository names

  - **snapshot** – A comma-separated list of snapshot names

  - **master_timeout** – Explicit operation timeout for connection to master node

get_repository(**args*, ***kwargs*)
  Return information about registered repositories. http://www.elasticsearch.org/guide/en/elasticsearch/reference/master/modules-snapshots.html

  Parameters

  - **repository** – A comma-separated list of repository names

  - **master_timeout** – Explicit operation timeout for connection to master node

  - **local** – Return local information, do not retrieve the state from master node (default: false)

restore(**args*, ***kwargs*)
  Restore a snapshot. http://www.elasticsearch.org/guide/en/elasticsearch/reference/master/modules-snapshots.html

  Parameters

  - **repository** – A repository name

  - **snapshot** – A snapshot name

  - **body** – Details of what to restore

  - **master_timeout** – Explicit operation timeout for connection to master node

  - **wait_for_completion** – Should this request wait until the operation has completed before returning, default False

status(**args*, ***kwargs*)
  http://www.elasticsearch.org/guide/en/elasticsearch/reference/current/modules-snapshots.html

  Parameters

  - **repository** – A repository name

  - **snapshot** – A comma-separated list of snapshot names

> • **master_timeout** – Explicit operation timeout for connection to master node

4.2 Exceptions

class `elasticsearch.`**`ImproperlyConfigured`**
    Exception raised when the config passed to the client is inconsistent or invalid.

class `elasticsearch.`**`ElasticsearchException`**
    Base class for all exceptions raised by this package's operations (doesn't apply to `ImproperlyConfigured`).

class `elasticsearch.`**`SerializationError`**(*ElasticsearchException*)
    Data passed in failed to serialize properly in the `Serializer` being used.

class `elasticsearch.`**`TransportError`**(*ElasticsearchException*)
    Exception raised when ES returns a non-OK (>=400) HTTP status code. Or when an actual connection error happens; in that case the `status_code` will be set to `'N/A'`.

    `error`
        A string error message.

    `info`
        Dict of returned error info from ES, where available.

    `status_code`
        The HTTP status code of the response that precipitated the error or `'N/A'` if not applicable.

class `elasticsearch.`**`NotFoundError`**(*TransportError*)
    Exception representing a 404 status code.

class `elasticsearch.`**`ConflictError`**(*TransportError*)
    Exception representing a 409 status code.

class `elasticsearch.`**`RequestError`**(*TransportError*)
    Exception representing a 400 status code.

class `elasticsearch.`**`ConnectionError`**(*TransportError*)
    Error raised when there was an exception while talking to ES. Original exception from the underlying `Connection` implementation is available as `.info`.

4.3 Connection Layer API

All of the classes reponsible for handling the connection to the Elasticsearch cluster. The default subclasses used can be overriden by passing parameters to the `Elasticsearch` class. All of the arguments to the client will be passed on to `Transport`, `ConnectionPool` and `Connection`.

For example if you wanted to use your own implementation of the `ConnectionSelector` class you can just pass in the *selector_class* parameter.

4.3.1 Transport

class elasticsearch.**Transport**(*hosts,* *connection_class=Urllib3HttpConnection,* *connection_pool_class=ConnectionPool,* *nodes_to_host_callback=construct_hosts_list,* *sniff_on_start=False,* *sniffer_timeout=None,* *sniff_on_connection_fail=False,* *serializer=JSONSerializer(),* *max_retries=3,* *** kwargs*)

Encapsulation of transport-related to logic. Handles instantiation of the individual connections as well as creating a connection pool to hold them.

Main interface is the *perform_request* method.

> **Parameters**
>
> - **hosts** – list of dictionaries, each containing keyword arguments to create a *connection_class* instance
>
> - **connection_class** – subclass of `Connection` to use
>
> - **connection_pool_class** – subclass of `ConnectionPool` to use
>
> - **host_info_callback** – callback responsible for taking the node information from */_cluser/nodes*, along with already extracted information, and producing a list of arguments (same as *hosts* parameter)
>
> - **sniff_on_start** – flag indicating whether to obtain a list of nodes from the cluser at startup time
>
> - **sniffer_timeout** – number of seconds between automatic sniffs
>
> - **sniff_on_connection_fail** – flag controlling if connection failure triggers a sniff
>
> - **sniff_timeout** – timeout used for the sniff request - it should be a fast api call and we are talking potentially to more nodes so we want to fail quickly.
>
> - **serializer** – serializer instance
>
> - **serializers** – optional dict of serializer instances that will be used for deserializing data coming from the server. (key is the mimetype)
>
> - **default_mimetype** – when no mimetype is specified by the server response assume this mimetype, defaults to *'application/json'*
>
> - **max_retries** – maximum number of retries before an exception is propagated
>
> - **send_get_body_as** – for GET requests with body this option allows you to specify an alternate way of execution for environments that don't support passing bodies with GET requests. If you set this to 'POST' a POST method will be used instead, if to 'source' then the body will be serialized and passed as a query parameter *source*.

Any extra keyword arguments will be passed to the *connection_class* when creating and instance unless overriden by that connection's options provided as part of the hosts parameter.

add_connection(*host*)

> Create a new `Connection` instance and add it to the pool.
>
> > **Parameters host** – kwargs that will be used to create the instance

get_connection()

> Retreive a `Connection` instance from the `ConnectionPool` instance.

mark_dead(*connection*)

> Mark a connection as dead (failed) in the connection pool. If sniffing on failure is enabled this will initiate the sniffing process.

> **Parameters connection** – instance of `Connection` that failed

perform_request(*method*, *url*, *params=None*, *body=None*)

> Perform the actual request. Retrieve a connection from the connection pool, pass all the information to it's perform_request method and return the data.
>
> If an exception was raised, mark the connection as failed and retry (up to *max_retries* times).
>
> If the operation was succesful and the connection used was previously marked as dead, mark it as live, resetting it's failure count.
>
> **Parameters**
>
> - **method** – HTTP method to use
>
> - **url** – absolute url (without host) to target
>
> - **params** – dictionary of query parameters, will be handed over to the underlying `Connection` class for serialization
>
> - **body** – body of the request, will be serializes using serializer and passed to the connection

set_connections(*hosts*)

> Instantiate all the connections and crate new connection pool to hold them. Tries to identify unchanged hosts and re-use existing `Connection` instances.
>
> **Parameters hosts** – same as *__init__*

sniff_hosts()

> Obtain a list of nodes from the cluster and create a new connection pool using the information retrieved.
>
> To extract the node connection parameters use the *nodes_to_host_callback*.

4.3.2 Connection Pool

class elasticsearch.**ConnectionPool**(*connections*, *dead_timeout=60*, *selector_class=RoundRobinSelector*, *randomize_hosts=True*, ***kwargs*)

> Container holding the `Connection` instances, managing the selection process (via a `ConnectionSelector`) and dead connections.
>
> It's only interactions are with the `Transport` class that drives all the actions within *ConnectionPool*.
>
> Initially connections are stored on the class as a list and, along with the connection options, get passed to the *ConnectionSelector* instance for future reference.
>
> Upon each request the *Transport* will ask for a *Connection* via the *get_connection* method. If the connection fails (it's *perform_request* raises a *ConnectionError*) it will be marked as dead (via *mark_dead*) and put on a timeout (if it fails N times in a row the timeout is exponentially longer - the formula is *default_timeout * 2 ** (fail_count - 1)*). When the timeout is over the connection will be resurrected and returned to the live pool. A connection that has been peviously marked as dead and succeedes will be marked as live (it's fail count will be deleted).
>
> **Parameters**
>
> - **connections** – list of tuples containing the `Connection` instance and it's options
>
> - **dead_timeout** – number of seconds a connection should be retired for after a failure, increases on consecutive failures
>
> - **timeout_cutoff** – number of consecutive failures after which the timeout doesn't increase
>
> - **selector_class** – `ConnectionSelector` subclass to use

- **randomize_hosts** – shuffle the list of connections upon arrival to avoid dog piling effect across processes

get_connection()
> Return a connection from the pool using the *ConnectionSelector* instance.
>
> It tries to resurrect eligible connections, forces a resurrection when no connections are availible and passes the list of live connections to the selector instance to choose from.
>
> Returns a connection instance and it's current fail count.

mark_dead(*connection*, *now=None*)
> Mark the connection as dead (failed). Remove it from the live pool and put it on a timeout.
>
> > **Parameters connection** – the failed instance

mark_live(*connection*)
> Mark connection as healthy after a resurrection. Resets the fail counter for the connection.
>
> > **Parameters connection** – the connection to redeem

resurrect(*force=False*)
> Attempt to resurrect a connection from the dead pool. It will try to locate one (not all) eligible (it's timeout is over) connection to return to th live pool.
>
> > **Parameters force** – resurrect a connection even if there is none eligible (used when we have no live connections)

4.3.3 Connection Selector

class elasticsearch.**ConnectionSelector**(*opts*)
> Simple class used to select a connection from a list of currently live connection instances. In init time it is passed a dictionary containing all the connections' options which it can then use during the selection process. When the *select* method is called it is given a list of *currently* live connections to choose from.
>
> The options dictionary is the one that has been passed to `Transport` as *hosts* param and the same that is used to construct the Connection object itself. When the Connection was created from information retrieved from the cluster via the sniffing process it will be the dictionary returned by the *host_info_callback*.
>
> Example of where this would be useful is a zone-aware selector that would only select connections from it's own zones and only fall back to other connections where there would be none in it's zones.
>
> > **Parameters opts** – dictionary of connection instances and their options

select(*connections*)
> Select a connection from the given list.
>
> > **Parameters connections** – list of live connections to choose from

4.3.4 Connection

class elasticsearch.**Connection**(*host='localhost'*, *port=9200*, *url_prefix=''*, *timeout=10*, ***kwargs*)
> Class responsible for maintaining a connection to an Elasticsearch node. It holds persistent connection pool to it and it's main interface (*perform_request*) is thread-safe.
>
> Also responsible for logging.
>
> > **Parameters**
> >
> > - **host** – hostname of the node (default: localhost)

- **port** – port to use (default: 9200)

- **url_prefix** – optional url prefix for elasticsearch

- **timeout** – default timeout in seconds (default: 10)

log_request_fail(*method*, *full_url*, *body*, *duration*, *status_code=None*, *exception=None*)
> Log an unsuccessful API call.

log_request_success(*method*, *full_url*, *path*, *body*, *status_code*, *response*, *duration*)
> Log a successful API call.

4.4 Transport classes

List of transport classes that can be used, simply import your choice and pass it to the constructor of `Elasticsearch` as *connection_class*. Note that Thrift and Memcached protocols are experimental and require a plugin to be installed in your cluster as well as additional dependencies (*thrift==0.9* and *pylibmc==1.2*).

For example to use the thrift connection just import it and use it. The connection classes are aware of their respective default ports (9500 for thrift) so there is no need to specify them unless modified:

```python
from elasticsearch import Elasticsearch, ThriftConnection
es = Elasticsearch(connection_class=ThriftConnection)
```

4.4.1 Connection

class elasticsearch.connection.**Connection**(*host='localhost'*, *port=9200*, *url_prefix=''*, *timeout=10*, ***kwargs*)
> Class responsible for maintaining a connection to an Elasticsearch node. It holds persistent connection pool to it and it's main interface (*perform_request*) is thread-safe.
>
> Also responsible for logging.
>
> > **Parameters**
> >
> > - **host** – hostname of the node (default: localhost)
> >
> > - **port** – port to use (default: 9200)
> >
> > - **url_prefix** – optional url prefix for elasticsearch
> >
> > - **timeout** – default timeout in seconds (default: 10)

4.4.2 Urllib3HttpConnection

class elasticsearch.connection.**Urllib3HttpConnection**(*host='localhost'*, *port=9200*, *http_auth=None*, *use_ssl=False*, *maxsize=10*, ***kwargs*)
> Default connection class using the *urllib3* library and the http protocol.
>
> > **Parameters**
> >
> > - **http_auth** – optional http auth information as either ':' separated string or a tuple
> >
> > - **use_ssl** – use ssl for the connection if *True*
> >
> > - **maxsize** – the maximum number of connections which will be kept open to this host.

4.4.3 RequestsHttpConnection

class `elasticsearch.connection.`**`RequestsHttpConnection`**(*host='localhost',   port=9200, http_auth=None, use_ssl=False, **kwargs*)

> Connection using the *requests* library.
>
> > **Parameters**
> >
> > > • **http_auth** – optional http auth information as either ':' separated string or a tuple
> > >
> > > • **use_ssl** – use ssl for the connection if *True*

4.4.4 ThriftConnection

class `elasticsearch.connection.`**`ThriftConnection`**(*host='localhost',   port=9500, framed_transport=False, use_ssl=False, **kwargs*)

> Connection using the *thrift* protocol to communicate with elasticsearch.
>
> See https://github.com/elasticsearch/elasticsearch-transport-thrift for additional info.
>
> > **Parameters framed_transport** – use *TTransport.TFramedTransport* instead of
>
> *TTransport.TBufferedTransport*

4.4.5 MemcachedConnection

class `elasticsearch.connection.`**`MemcachedConnection`**(*host='localhost',   port=11211, **kwargs*)

> Client using the *pylibmc* python library to communicate with elasticsearch using the memcached protocol. Requires plugin in the cluster.
>
> See https://github.com/elasticsearch/elasticsearch-transport-memcached for more details.

4.5 Helpers

Collection of simple helper functions that abstract some specifics or the raw API.

`elasticsearch.helpers.`**`streaming_bulk`**(*client, actions, chunk_size=500, raise_on_error=False, expand_action_callback=<function expand_action at 0x7f90f1ed56e0>, **kwargs*)

> Streaming bulk consumes actions from the iterable passed in and yields results per action. For non-streaming usecases use `bulk()` which is a wrapper around streaming bulk that returns summary information about the bulk operation once the entire input is consumed and sent.
>
> This function expects the action to be in the format as returned by `search()`, for example:

```
{
    '_index': 'index-name',
    '_type': 'document',
    '_id': 42,
    '_parent': 5,
    '_ttl': '1d',
    '_source': {
        ...
```

```
    }
}
```

Alternatively, if *_source* is not present, it will pop all metadata fields from the doc and use the rest as the document data.

Alternative actions (*_op_type* field defaults to *index*) can be sent as well:

```
{
    '_op_type': 'delete',
    '_index': 'index-name',
    '_type': 'document',
    '_id': 42,
}
{
    '_op_type': 'update',
    '_index': 'index-name',
    '_type': 'document',
    '_id': 42,
    'doc': {'question': 'The life, universe and everything.'}
}
```

Parameters

- **client** – instance of `Elasticsearch` to use
- **actions** – iterable containing the actions to be executed
- **chunk_size** – number of docs in one chunk sent to es (default: 500)
- **raise_on_error** – raise *BulkIndexError* containing errors (as *.errors* from the execution of the last chunk)
- **expand_action_callback** – callback executed on each action passed in, should return a tuple containing the action line and the data line (*None* if data line should be omitted).

elasticsearch.helpers.**bulk**(*client*, *actions*, *stats_only=False*, ***kwargs*)

Helper for the `bulk()` api that provides a more human friendly interface - it consumes an iterator of actions and sends them to elasticsearch in chunks. It returns a tuple with summary information - number of successfully executed actions and either list of errors or number of errors if *stats_only* is set to *True*.

See `streaming_bulk()` for more information and accepted formats.

Parameters

- **client** – instance of `Elasticsearch` to use
- **actions** – iterator containing the actions
- **stats_only** – if *True* only report number of successful/failed operations instead of just number of successful and a list of error responses

Any additional keyword arguments will be passed to `streaming_bulk()` which is used to execute the operation.

elasticsearch.helpers.**scan**(*client*, *query=None*, *scroll='5m'*, ***kwargs*)

Simple abstraction on top of the `scroll()` api - a simple iterator that yields all hits as returned by underlining scroll requests.

Parameters

- **client** – instance of `Elasticsearch` to use

- **query** – body for the `search()` api

- **scroll** – Specify how long a consistent view of the index should be maintained for scrolled search

Any additional keyword arguments will be passed to the initial `search()` call.

elasticsearch.helpers.**reindex**(*client*,   *source_index*,   *target_index*,   *target_client=None*,
chunk_size=500, *scroll='5m'*)
Reindex all documents from one index to another, potentially (if *target_client* is specified) on a different cluster.

Note: This helper doesn't transfer mappings, just the data.

Parameters

- **client** – instance of `Elasticsearch` to use (for read if *target_client* is specified as well)

- **source_index** – index (or list of indices) to read documents from

- **target_index** – name of the index in the target cluster to populate

- **target_client** – optional, is specified will be used for writing (thus enabling reindex between clusters)

- **chunk_size** – number of docs in one chunk sent to es (default: 500)

- **scroll** – Specify how long a consistent view of the index should be maintained for scrolled search

4.6 Changelog

4.6.1 1.1.1 (2014-07-04)

Bugfix release fixing escaping issues with *request_timeout*.

4.6.2 1.1.0 (2014-07-02)

Compatibility with newest Elasticsearch APIs.

- Test helpers - *ElasticsearchTestCase* and *get_test_client* for use in your tests

- Python 3.2 compatibility

- Use `simplejson` if installed instead of stdlib json library

- Introducing a global *request_timeout* parameter for per-call timeout

- Bug fixes

4.6.3 1.0.0 (2014-02-11)

Elasticsearch 1.0 compatibility. See 0.4.X releases (and 0.4 branch) for code compatible with 0.90 elasticsearch.

- major breaking change - compatible with 1.0 elasticsearch releases only!

- Add an option to change the timeout used for sniff requests (*sniff_timeout*).

- empty responses from the server are now returned as empty strings instead of None

---

- *get_alias* now has *name* as another optional parameter due to issue #4539 in es repo. Note that the order of params have changed so if you are not using keyword arguments this is a breaking change.

4.6.4 0.4.4 (2013-12-23)

- *helpers.bulk_index* renamed to *helpers.bulk* (alias put in place for backwards compatibility, to be removed in future versions)

- Added *helpers.streaming_bulk* to consume an iterator and yield results per operation

- *helpers.bulk* and *helpers.streaming_bulk* are no longer limitted to just index operations.

- unicode body (for *incices.analyze* for example) is now handled correctly

- changed *perform_request* on *Connection* classes to return headers as well. This is a backwards incompatible change for people who have developed their own connection class.

- changed deserialization mechanics. Users who provided their own serializer that didn't extend *JSONSerializer* need to specify a *mimetype* class attribute.

- minor bug fixes

4.6.5 0.4.3 (2013-10-22)

- Fixes to *helpers.bulk_index*, better error handling

- More benevolent *hosts* argument parsing for *Elasticsearch*

- *requests* no longer required (nor recommended) for install

4.6.6 0.4.2 (2013-10-08)

- *ignore* param acceted by all APIs

- Fixes to *helpers.bulk_index*

4.6.7 0.4.1 (2013-09-24)

Initial release.

License

Copyright 2013 Elasticsearch

Licensed under the Apache License, Version 2.0 (the "License"); you may not use this file except in compliance with the License. You may obtain a copy of the License at

http://www.apache.org/licenses/LICENSE-2.0

Unless required by applicable law or agreed to in writing, software distributed under the License is distributed on an "AS IS" BASIS, WITHOUT WARRANTIES OR CONDITIONS OF ANY KIND, either express or implied. See the License for the specific language governing permissions and limitations under the License.

Indices and tables

- *genindex*
- *modindex*
- *search*

e