
Django Simple Elasticsearch Documentation

Release 1.0.0

James Addison

December 20, 2016

1	Documentation	3
2	Features	5
3	Installation	7
4	License	9
5	Todo	11
6	Jump to a section	13
6.1	Usage	13
6.2	Contributing	15
6.3	Credits	17
6.4	History	17
6.5	1.0.0 (2016-12-20)	17
6.6	0.9.16 (2015-04-24)	17
6.7	0.9.15 (2015-01-31)	18
6.8	0.9.14 (2015-01-31)	18
6.9	0.9.13 (2015-01-30)	18
6.10	0.9.12 (2014-12-17)	18
6.11	0.9.11 (2014-12-08)	18
6.12	0.9.10 (2014-12-04)	18
6.13	0.9.9 (2014-11-24)	18
6.14	0.9.8 (2014-11-23)	18
6.15	0.9.7 (2014-11-16)	19
6.16	0.9.6 (2014-11-16)	19
6.17	0.9.5 (2014-11-15)	19
6.18	0.9.2 (2014-11-12)	19
6.19	0.9.1 (2014-11-10)	19
6.20	0.9.0 (2014-11-10)	19
6.21	0.5.0 (2014-03-05)	19
6.22	Indices and tables	19

This package provides a simple method of creating Elasticsearch indexes for Django models.

Using a version older than 0.9.0? Please be aware that as of v0.9.0, this package has changed in a backwards-incompatible manner. Version 0.5 is deprecated and no longer maintained.

Documentation

Visit the [django-simple-elasticsearch](#) documentation on ReadTheDocs.

Features

- class mixin with a set of `@classmethods` used to handle:
- type mapping definition
- individual object indexing and deletion
- bulk object indexing
- model signal handlers for `pre/post_save` and `pre/post_delete` (optional)
- management command to handle index/type mapping initialization and bulk indexing
- uses Elasticsearch aliases to ease the burden of re-indexing
- small set of Django classes and functions to help deal with Elasticsearch querying
- base search form class to handle input validation, query preparation and response handling
- multi-search processor class to batch multiple Elasticsearch queries via `_msearch`
- ‘get’ shortcut functions
- post index create/rebuild signals available to perform actions after certain stages (ie. add your own percolators)

Installation

At the command line:

```
$ easy_install django-simple-elasticsearch
```

Or:

```
$ pip install django-simple-elasticsearch
```

License

django-simple-elasticsearch is licensed as free software under the BSD license.

Todo

- Review search classes - simplify functionality where possible. This may cause breaking changes.
- Tests. Write them.
- Documentation. Write it.

Jump to a section

6.1 Usage

For a minimal investment of time, Django Simple Elasticsearch offers a number of perks. Implementing a class with the `ElasticsearchIndexMixin` lets you:

- initialize your Elasticsearch indices and mappings via the included `es_manage` management command
- perform Elasticsearch bulk indexing via the same `es_manage` management command
- perform Elasticsearch bulk indexing as well as individual index/delete requests on demand in your code
- connect the available `ElasticsearchIndexMixin` save and delete handlers to Django's available model signals (ie `post_save`, `post_delete`)

Let's look at an example implementation of `ElasticsearchIndexMixin`. Here's a couple of blog-related Models in a `models.py` file:

```
class Blog(models.Model):
    name = models.CharField(max_length=50)
    description = models.TextField()

class BlogPost(models.Model):
    blog = models.ForeignKey(Blog)
    slug = models.SlugField()
    title = models.CharField(max_length=50)
    body = models.TextField()
    created_at = models.DateTimeField(auto_now_add=True)
```

To start with `simple_elasticsearch`, you'll need to tell it that the `BlogPost` class implements the `ElasticsearchIndexMixin` mixin, so in your `settings.py` set the `ELASTICSEARCH_TYPE_CLASSES` setting:

```
ELASTICSEARCH_TYPE_CLASSES = [
    'blog.models.BlogPost'
]
```

If you do not add this setting, everything will still work except for the `es_manage` command - it won't know what indices to create, type mappings to set or what objects to index. As you add additional `ElasticsearchIndexMixin`-based index handlers, add them to this list.

All right, let's add in `ElasticsearchIndexMixin` to the `BlogPost` model. Only pertinent changes from the above `models.py` are shown:

```
from simple_elasticsearch.mixins import ElasticsearchIndexMixin

...

class BlogPost(models.Model, ElasticsearchIndexMixin):
    blog = models.ForeignKey(Blog)
    slug = models.SlugField()
    title = models.CharField(max_length=50)
    body = models.TextField()
    created_at = models.DateTimeField(auto_now_add=True)

    @classmethod
    def get_queryset(cls):
        return BlogPost.objects.all().select_related('blog')

    @classmethod
    def get_index_name(cls):
        return 'blog'

    @classmethod
    def get_type_name(cls):
        return 'posts'

    @classmethod
    def get_type_mapping(cls):
        return {
            "properties": {
                "created_at": {
                    "type": "date",
                    "format": "dateOptionalTime"
                },
                "title": {
                    "type": "string"
                },
                "body": {
                    "type": "string"
                },
                "slug": {
                    "type": "string"
                },
                "blog": {
                    "properties": {
                        "id": {
                            "type": "long"
                        },
                        "name": {
                            "type": "string"
                        },
                        "description": {
                            "type": "string"
                        }
                    }
                }
            }
        }

    @classmethod
    def get_document(cls, obj):
```

```

return {
    'created_at': obj.created_at,
    'title': obj.title,
    'body': obj.body,
    'slug': obj.slug,
    'blog': {
        'id': obj.blog.id,
        'name': obj.blog.name,
        'description': obj.blog.description,
    }
}

```

With this mixin implementation, you can now use the `es_manage` management command to bulk reindex all `BlogPost` items. Note that there are additional `@classmethods` you can override to customize functionality. Sane defaults have been provided for these - see the source for details.

Of course, our `BlogPost` implementation doesn't ensure that your Elasticsearch index is updated every time you save or delete - for this, you can use the `ElasticsearchIndexMixin` built-in save and delete handlers.

```

from django.db.models.signals import post_save, pre_delete

...

post_save.connect(BlogPost.save_handler, sender=BlogPost)
pre_delete.connect(BlogPost.delete_handler, sender=BlogPost)

```

Awesome - Django's magic is applied.

TODO:

- add examples for more complex data situations
- add examples of using `ElasticsearchForm` to derive your own search forms
- add examples of using `es_manage` management command options
- add examples/scenarios when to use `post_indices_create` and `post_indices_rebuild` signals (ie. adding percolators to new indices)

6.2 Contributing

Contributions are welcome, and they are greatly appreciated! Every little bit helps, and credit will always be given.

You can contribute in many ways:

6.2.1 Types of Contributions

Report Bugs

Report bugs at <https://github.com/jaddison/django-simple-elasticsearch/issues>.

If you are reporting a bug, please include:

- Your operating system name and version.
- Any details about your local setup that might be helpful in troubleshooting.
- Detailed steps to reproduce the bug.

Fix Bugs

Look through the GitHub issues for bugs. Anything tagged with “bug” is open to whoever wants to implement it.

Implement Features

Look through the GitHub issues for features. Anything tagged with “feature” is open to whoever wants to implement it.

Write Documentation

Django Simple Elasticsearch could always use more documentation, whether as part of the official Django Simple Elasticsearch docs, in docstrings, or even on the web in blog posts, articles, and such.

Submit Feedback

The best way to send feedback is to file an issue at <https://github.com/jaddison/django-simple-elasticsearch/issues>.

If you are proposing a feature:

- Explain in detail how it would work.
- Keep the scope as narrow as possible, to make it easier to implement.
- Remember that this is a volunteer-driven project, and that contributions are welcome :)

6.2.2 Get Started!

Ready to contribute? Here’s how to set up *django-simple-elasticsearch* for local development.

1. Fork the *django-simple-elasticsearch* repo on GitHub.
2. Clone your fork locally:

```
$ git clone git@github.com:your_name_here/django-simple-elasticsearch.git
```

3. Install your local copy into a virtualenv. Assuming you have virtualenvwrapper installed, this is how you set up your fork for local development:

```
$ mkvirtualenv django-simple-elasticsearch
$ cd django-simple-elasticsearch/
$ python setup.py develop
```

4. Create a branch for local development:

```
$ git checkout -b name-of-your-bugfix-or-feature
```

Now you can make your changes locally.

5. When you’re done making changes, check that your changes pass flake8 and the tests, including testing other Python versions with tox:

```
$ flake8 simple_elasticsearch
$ python setup.py test
$ tox
```

To get flake8 and tox, just pip install them into your virtualenv.

6. Commit your changes and push your branch to GitHub:

```
$ git add .  
$ git commit -m "Your detailed description of your changes."  
$ git push origin name-of-your-bugfix-or-feature
```

7. Submit a pull request through the GitHub website.

6.2.3 Pull Request Guidelines

Before you submit a pull request, check that it meets these guidelines:

1. The pull request should include tests.
2. If the pull request adds functionality, the docs should be updated. Put your new functionality into a function with a docstring, and add the feature to the list in README.rst.
3. The pull request should work for Python 2.6, 2.7, 3.3, and 3.4, and for PyPy. Check https://travis-ci.org/jaddison/django-simple-elasticsearch/pull_requests and make sure that the tests pass for all supported Python versions.

6.3 Credits

6.3.1 Development Lead

- James Addison <code@scottisheyes.com>

6.3.2 Contributors

None yet. Why not be the first?

6.4 History

6.5 1.0.0 (2016-12-20)

- Updated 0.9.x codebase version to 1.0.0.
- Reversed decision on the deprecation of the 0.9.x codebase - it will be maintained in this new 1.x branch, although new functionality will mostly occur on newer releases.
- Adding cleanup command to remove unaliased indices.
- Added ELASTICSEARCH_DELETE_OLD_INDEXES setting to auto-remove after a rebuild.
- Thanks to Github user @jimjkelly for the index removal inspiration.

6.6 0.9.16 (2015-04-24)

- Addressing Django 1.8 warnings.

6.7 0.9.15 (2015-01-31)

- BUGFIX: Merging pull request from @key that addresses Python 3 support (management command now works).

6.8 0.9.14 (2015-01-31)

- BUGFIX: Adding in missing *signals.py* file.

6.9 0.9.13 (2015-01-30)

- Added in new *post_indices_create* and *post_indices_rebuild* signals to allow users to run actions at various points during the index creation and bulk indexing processes.

6.10 0.9.12 (2014-12-17)

- fixed an issue where per-item request parameters were being added to the bulk data request JSON incorrectly. Tests updated.

6.11 0.9.11 (2014-12-08)

- added warning if Django's `DEBUG=True` (causes out of memory errors on constrained systems due to Django query caching)
- added index setting modification on rebuilding indices to remove replicas, lucene segment merging and disabling the refresh interval - restoring the original settings afterwards.

6.12 0.9.10 (2014-12-04)

- added *page* and *page_size* validation in *add_search()*

6.13 0.9.9 (2014-11-24)

- Renamed search form related classes - more breaking changes. Added in support for Django's pagination classes (internal hack).

6.14 0.9.8 (2014-11-23)

- Revamped search form related classes - includes breaking changes.

6.15 0.9.7 (2014-11-16)

- Python3 supported mentioned in PyPi categorization; new testcases added. Minor interface change (added *@classmethod*).

6.16 0.9.6 (2014-11-16)

- Python 3.3+ support, modified (no new) testcases.

6.17 0.9.5 (2014-11-15)

- Added in tox support, initial set of test cases and verified travis-ci is working.

6.18 0.9.2 (2014-11-12)

- Fixed broken management command.

6.19 0.9.1 (2014-11-10)

- Added missing management command module.

6.20 0.9.0 (2014-11-10)

- In what will become version 1.0, this 0.9.x codebase is a revamp of the original codebase (v0.5.x). Completely breaking over previous versions.

6.21 0.5.0 (2014-03-05)

Final release in 0.x codebase - this old codebase is now unmaintained.

6.22 Indices and tables

- [genindex](#)
- [modindex](#)
- [search](#)