
collective.elasticsearch Documentation

Release 2.0.0

Nathan Van Gheem

Dec 21, 2018

Contents

1	Overview	1
2	Compatibility	3
3	State	5
4	Celery support	7
4.1	Installation	7
4.2	Configuration	9
4.3	Changelog	10
5	Indices and tables	13

CHAPTER 1

Overview

This package aims to index all fields the portal_catalog indexes and allows you to delete the *Title*, *Description* and *SearchableText* indexes which can provide significant improvement to performance and RAM usage.

Then, ElasticSearch queries are ONLY used when Title, Description and SearchableText text are in the query. Otherwise, the plone's default catalog will be used. This is because Plone's default catalog is faster on normal queries than using ElasticSearch.

CHAPTER 2

Compatibility

Only unit tested with Plone 5 with Dexterity types and archetypes.

It should also work with Plone 4.3 and Plone 5.1.

Deployed with Elasticsearch 2.4.3

CHAPTER 3

State

Support for all index column types is done EXCEPT for the DateRecurringIndex index column type. If you are doing a full text search along with a query that contains a DateRecurringIndex column, it will not work.

This package comes with Celery support where all indexing operations will be pushed into celery to be run asynchronously.

Please see instructions for `collective.celery` to see how this works.

Contents:

4.1 Installation

4.1.1 `collective.elasticsearch`

To install `collective.elasticsearch` into the global Python environment (or a workingenv), using a traditional Zope 2 instance, you can do this:

- When you're reading this you have probably already run `easy_install collective.elasticsearch`. Find out how to install `setuptools` (and `EasyInstall`) here: <http://peak.telecommunity.com/DevCenter/EasyInstall>
- If you are using Zope 2.9 (not 2.10), get `pythonproducts` and install it via:

```
python setup.py install --home /path/to/instance
```

into your Zope instance.

- Create a file called `collective.elasticsearch-configure.zcml` in the `/path/to/instance/etc/package-includes` directory. The file should only contain this:

```
<include package="collective.elasticsearch" />
```

Alternatively, if you are using `zc.buildout` and the `plone.recipe.zope2instance` recipe to manage your project, you can do this:

- Add `collective.elasticsearch` to the list of eggs to install, e.g.:

```
[buildout]
...
eggs =
    ...
    collective.elasticsearch
```

- Tell the `plone.recipe.zope2instance` recipe to install a ZCML slug:

```
[instance]
recipe = plone.recipe.zope2instance
...
zcml =
    collective.elasticsearch
```

- Re-run buildout, e.g. with:

```
$ ./bin/buildout
```

You can skip the ZCML slug if you are going to explicitly include the package from another package's `configure.zcml` file.

4.1.2 elasticsearch

Less than 5 minutes:

- Download & install Java
- Download & install Elastic Search
- `bin/elasticsearch`

Step by Step for Ubuntu:

- `add-apt-repository ppa:webupd8team/java`
- `apt-get update`
- `apt-get install git curl oracle-java7-installer`
- `curl -O https://download.elastic.co/elasticsearch/release/org/elasticsearch/distribution/tar/elasticsearch/2.4.3/elasticsearch-2.4.3.tar.gz`
- `tar xfvz elasticsearch-2.4.3.tar.gz`
- `cd elasticsearch`
- `bin/elasticsearch`

Step by Step for CentOS/RedHat:

- `yum -y install java-1.8.0-openjdk.x86_64`
- `alternatives --auto java`
- `curl -O https://download.elastic.co/elasticsearch/release/org/elasticsearch/distribution/tar/elasticsearch/2.4.3/elasticsearch-2.4.3.tar.gz`
- `tar xfvz elasticsearch-2.4.3.tar.gz`
- `cd elasticsearch`
- `bin/elasticsearch`

Does it work?

- curl `http://localhost:9200/`
- Do you see the Hudsucker Proxy reference? “You Know, for Search”

4.2 Configuration

4.2.1 Basic configuration

- Goto Control Panel
- Add “Elastic Search” in Add-on Products
- Click “Elastic Search” in “Add-on Configuration”
- Enable
- Click “Convert Catalog”
- Click “Rebuild Catalog”

4.2.2 Changing the index used for elasticsearch

The index used for elasticsearch is the path to the portal_catalog by default. So you don’t have anything to do if you have several plone site on the same instance (the plone site id would be different).

However, if you want to use the same elasticsearch instance with several plone instance, you may end up having conflicts. In that case, you may want to manually set the index used by adding the following code to the `__init__.py` file of your module:

```
from Products.CMFPlone.CatalogTool import CatalogTool
from collective.elasticsearch.es import CUSTOM_INDEX_NAME_ATTR

setattr(CatalogTool, CUSTOM_INDEX_NAME_ATTR, "my_elasticsearch_custom_index")
```

4.2.3 Adding custom index which are not in the catalog

An adapter is used to define the mapping between the index and the elasticsearch properties. You can override the `_default_mapping` attribute to add your own indexes:

```
<adapter
    factory=".mapping.MyMappingAdapter"
    provides="collective.elasticsearch.interfaces.IMappingProvider"
    for="zope.interface.Interface
        collective.elasticsearch.interfaces.IElasticSearchCatalog"
    layer=".layers.MyLayer" />
```

```
@implementer(IMappingProvider)
class MyMappingAdapter(object):

    _default_mapping = {
        'SearchableText': {'store': False, 'type': 'string', 'index': 'analyzed'},
        'Title': {'store': False, 'type': 'string', 'index': 'analyzed'},
```

(continues on next page)

(continued from previous page)

```
'Description': {'store': False, 'type': 'string', 'index': 'analyzed'},
'MyOwnIndex': {'store': False, 'type': 'string', 'index': 'analyzed'},
}
```

4.2.4 Changing the settings of the index

If you want to customize your elasticsearch index, you can override the `get_index_creation_body` method on the `MappingAdapter`:

```
@implementer(IMappingProvider)
class MyMappingAdapter(object):

    def get_index_creation_body(self):
        return {
            "settings" : {
                "number_of_shards": 1,
                "number_of_replicas": 0
            }
        }
```

4.2.5 Changing the query made to elasticsearch

The query generation is handled by another adapter:

```
<adapter
    factory=".query.QueryAssembler"
    provides=".interfaces.IQueryAssembler"
    for="zope.interface.Interface
        .interfaces.IElasticSearchCatalog" />
```

You will have to override the `__call__` method to change the query. Look at the original adapter to have a better idea on what you need to change.

4.3 Changelog

4.3.1 2.0.5 (2018-12-20)

- Fix sort order parsing and implementation [vangheem]

4.3.2 2.0.4 (2018-12-17)

- Fix date queries to work with *min:max* as well as *minmax* [vangheem]

4.3.3 2.0.3 (2018-12-04)

- Add `IReindexActive` to request as a flag for other code [lucid-0]

4.3.4 2.0.2 (2018-11-27)

- Fix error upgrading collective.elasticsearch [vangheem]
- Fix error causing “Server Status” on @@elastic-controlpanel to be empty. [fulv]

4.3.5 2.0.1 (2018-01-05)

- Prevent critical error when by chance query value is None. [thomasdesvenain]
- Minor code cleanup: readability, pep8, 80 cols, zca decorators. [jensens]
- Fix date criteria: ‘minmax’ instead of ‘min:max’ + string to date conversion [ebrehault]

4.3.6 2.0.0a6 (2017-03-29)

- Gracefully handle upgrades in the settings interface so it doesn’t break for people upgrading. [vangheem]

4.3.7 2.0.0a5 (2017-03-29)

- Running indexing as admin as it is possible to initiate reindex or index on an object that you do not have permissions for [vangheem]

4.3.8 2.0.0a4 (2017-03-27)

- released

4.3.9 2.0.0a3 (2017-03-27)

- Add a method to set the body of the request during index creation. [Gagaro]
- Fixed get brain in lazy list with negative indexes. [thomasdesvenain]
- The list of indexes that forces es search is configurable. [thomasdesvenain]
- Works under Plone 4.3. [thomasdesvenain]
- Works with archetypes contents. [thomasdesvenain]

4.3.10 2.0.0a2 (2016-07-19)

- We can pass a custom results factory and custom query parameters to IElasticSearchCatalog.search() method. So we can use it as a public interface for custom needs. [thomasdesvenain]
- Prevent from unindex before reindex when uid is unchanged, for instance at rename. Use a set for to-remove list. [thomasdesvenain]
- Fix indexing when removing the Title and Description indexes from Plone [vangheem]

4.3.11 2.0.0a1 (2016-06-06)

- upgrade to elasticsearch 2.x [vangheem]

4.3.12 1.0.1a4 (2016-05-22)

- provide better search query [vangheem]

4.3.13 1.0.1a3 (2016-03-22)

- make sure to get alias definition right [vangheem]

4.3.14 1.0.1a2 (2016-03-18)

- create index as an alias so you can potentially work on an existing alias without needing downtime [vangheem]

4.3.15 1.0.1a1 (2016-02-25)

- change default sorting to descending. Closes: <https://github.com/collective/collective.elasticsearch/issues/12> [neilferreira]

4.3.16 1.0.0a1 (2016-02-25)

- Initial release

CHAPTER 5

Indices and tables

- `genindex`
- `modindex`
- `search`