

---

# Lumberjack Documentation

*Release 0.1-dev*

**CERN**

July 31, 2015



<b>1</b>	<b>Installation</b>	<b>3</b>
<b>2</b>	<b>Quickstart</b>	<b>5</b>
2.1	Usage . . . . .	5
2.2	Configuration . . . . .	7
2.3	Advanced Usage . . . . .	10
2.4	Post-processors . . . . .	11
2.5	API . . . . .	12
2.6	Backend Modules . . . . .	13
2.7	Developers . . . . .	15
	<b>Python Module Index</b>	<b>31</b>



Lumberjack is a library which connects the Python logging framework to an Elasticsearch backend.

It has a number of useful features, including:

- Asynchronous logging by default
- Automatic time-rotating indices
- Elasticsearch mapping management



---

## Installation

---

Lumberjack can be installed with:

```
$ python setup.py install
```





---

## Quickstart

---

To use Lumberjack, you should have an Elasticsearch cluster already running.

You should instantiate the main Lumberjack class once in your application, specifying a list of Elasticsearch nodes to connect to:

```
from lumberjack import Lumberjack
lj = Lumberjack(hosts=[{'host': 'localhost', 'port': 9200}])
```

for an Elasticsearch node running locally on port 9200.

Once you have your Lumberjack object, you can then create log handlers to attach to loggers from logging:

```
from logging import getLogger, INFO

a_logger = getLogger(__name__)
a_logger.setLevel(INFO)

handler = lj.get_handler()

a_logger.addHandler(handler)
```

Then, you can log events to `a_logger` (or any of its children in the logger hierarchy) and they will be stored in Elasticsearch:

```
a_logger.error('Oh no! Something bad happened...')
```

Isn't that easy!

Note: you might need to wait for the periodic flush of the log queue before the entry actually appears in Elasticsearch. By default this is every 30 seconds.

## 2.1 Usage

Since you're using Elasticsearch to store log data, you probably don't want to just store messages as strings. Good! Lumberjack supports (and encourages) logging dicts to be stored directly in Elasticsearch as JSON:

```
a_logger.info({'message': 'User did something.',
              'username': 'rory',
              'uid': 23})
```

Note: the 'type' of the document in Elasticsearch is determined by the logger's name. For `a_logger` that's `__name__`. This means that it's a good idea to register several sub-loggers, one for each type of event being logged, and then attach the handler to the parent for all of them.

## 2.1.1 Schemas

Lumberjack also abstracts away creating mappings and templates for your data.

Before you log any events, it's a good idea to tell Elasticsearch what kind of data to expect in them. To do this you use the `register_schema` method:

Schemas correspond closely with mappings in Elasticsearch, but are processed by Lumberjack to include some sensible defaults. An example schema might be:

```
{
  '_source': True,
  'properties': {
    'ip_address': {'type': 'ip'},
    'user_id': {'type': 'long'},
    'username': {'type': 'string'}
  }
}
```

This method should be called once per schema to register; it's probably a good idea to call it at the same time as attaching your handler to a `logging.Logger` object.

## 2.1.2 A complete example

So now we've covered an introduction, basic usage and how to use schemas, let's put it all together.

Suppose we are writing a web app, and we want to log logins and searches:

In the general initialisation we put:

```
from logging import getLogger, INFO
from lumberjack import Lumberjack

lj = Lumberjack(hosts=[{'host': 'localhost', 'port': 9200}])

# Register a schema for logins
lj.register_schema('appname.login',
  {
    'properties': {
      'ip_address': {'type': 'ip'},
      'username': {'type': 'string'},
      'successful': {'type': 'boolean'},
      'attempt_number': {'type': 'short'}
    }
  })

# Register a schema for searches
lj.register_schema('appname.search',
  {
    'properties': {
      'ip_address': {'type': 'ip'},
      'query': {
        'type': 'string',
        'index': 'analyzed'
      }
    }
  })

logger = getLogger('appname')
```

```

handler = lj.get_handler()

logger.addHandler(handler)
logger.setLevel(INFO)

```

In the login handling function we put:

```

from logging import getLogger

logger = getLogger('appname.login')

log_event = {'ip_address': some_ip_address_variable,
             'username': the_username,
             'successful': login_ok,
             'attempt_number': attempt_number}
logger.info(log_event)

```

And in the search handling function we put:

```

from logging import getLogger

logger = getLogger('appname.search')

log_event = {'ip_address': some_ip_address_variable,
             'query': query_string}
logger.info(log_event)

```

Now we've integrated elasticsearch logging into our web application.

### 2.1.3 Next steps

At the minimum, you should read the index prefixes bit of the Configuration section.

## 2.2 Configuration

The default config is included at the end of this file for reference.

Lumberjack is configured using a dict of config options passed to the Lumberjack object on instantiation. A safely-modifiable copy of the default config dict can be obtained by calling `lumberjack.get_default_config()`:

```

from lumberjack import Lumberjack, get_default_config
lj = Lumberjack(hosts=[...], config=get_default_config())

```

With the exception of `index_prefix`, these defaults should be sensible for production. You should change `index_prefix` to something different for each of your applications.

The config returned is a deep copy of the default one, so it can be modified as you like without breaking subsequent calls to `lumberjack.get_default_config()`:

```

from lumberjack import Lumberjack, get_default_config

my_config = get_default_config()
my_config['index_prefix'] = 'a-special-prefix-'

lj = Lumberjack(hosts=[...], config=my_config)

```

### 2.2.1 The index prefix

This configures the prefix for the created elasticsearch indices. Indices are created with a constant prefix and a time-based suffix, so might be named *generic-logging-2014.10* for log entries from an unconfigured Lumberjack instance in October 2014.

### 2.2.2 The default mapping

This contains the basis for generating mappings in Elasticsearch. Its values are overridden by the values in the *schema* dict passed to `lumberjack.Lumberjack.register_schema()`. It contains keys like `_source`, `_all`, and `_ttl`.

Note that special processing happens to the `properties` key: instead of being overwritten by the schema's `properties` key, the fields provided to Elasticsearch are the union of the two, with the schema's fields taking precedence.

### 2.2.3 Default properties for types

When a field is given a particular type in the schema, Lumberjack automatically adds some properties to the field. For example, for `string` type fields, Lumberjack disables analysis on them. (The reason for this is that while analysis is a powerful Elasticsearch feature when dealing with natural language documents, for log data it makes little sense.)

### 2.2.4 The interval

This is the (maximum) amount of time to wait between flushes of the log event queue to Elasticsearch. It is an integer or floating-point value in seconds.

### 2.2.5 The maximum queue length

This is the maximum length the queue can grow to before a flush is triggered automatically.

### 2.2.6 The fallback log file

If Lumberjack experiences an error when indexing to Elasticsearch, it will fall back to dumping JSON to the file given in this variable.

### 2.2.7 The default config

```
# -*- coding: utf-8 -*-
#
# This file is part of Lumberjack.
# Copyright 2014 CERN.
#
# Lumberjack is free software: you can redistribute it and/or modify it under
# the terms of the GNU Lesser General Public License as published by the Free
# Software Foundation, either version 3 of the License, or (at your option) any
# later version.
#
# Lumberjack is distributed in the hope that it will be useful, but WITHOUT ANY
# WARRANTY; without even the implied warranty of MERCHANTABILITY or FITNESS FOR
```

```

# A PARTICULAR PURPOSE. See the GNU Lesser General Public License for more
# details.
#
# You should have received a copy of the GNU Lesser General Public License
# along with Lumberjack. If not, see <http://www.gnu.org/licenses/>.

"""The default configuration for Lumberjack."""

from __future__ import absolute_import
from copy import deepcopy

DEFAULT_CONFIG = {
    'default_mapping': {
        '_source': {'enabled': True},
        '_all': {'enabled': False},
        '_ttl': {'enabled': True},
        'properties': {
            'message': {
                'type': 'string'
            },
            '@timestamp': {
                'type': 'date',
            },
            'level': {
                'type': 'integer'
            }
        }
    },
    'default_type_properties': {
        'string': {
            'index': 'not_analyzed',
            'norms': {'enabled': False}
        },
        'date': {
            'format': 'dateOptionalTime'
        }
    },
    'index_prefix': 'generic-logging-',
    'interval': 30,
    'max_queue_length': None,
    'fallback_log_file': '/tmp/lumberjack_fallback.log'
}

def get_default_config():
    """Get a copy of the default config.

    The copy can be modified in-place without affecting the default config
    itself.

    """
    return deepcopy(DEFAULT_CONFIG)

```

## 2.3 Advanced Usage

This details some more advanced usages of Lumberjack. You don't need to read this just to get up and running, but it might be handy when tweaking your cluster later on.

### 2.3.1 Custom Elasticsearch objects

You can also create an `elasticsearch.Elasticsearch` object yourself and pass it in. This is useful if you want to do some customisation of the Elasticsearch connection, for example to connect using Thrift:

```
from elasticsearch import Elasticsearch
from lumberjack import Lumberjack

es = Elasticsearch(connection_class=ThriftConnection)
lj = Lumberjack(elasticsearch=es)
```

### 2.3.2 Index suffixes

Indices created by Lumberjack are named using a constant prefix and a time-based suffix, so might be named `generic-logging-2014.10` for log entries from an unconfigured Lumberjack instance in October 2014. (For details about the prefix, see the Configuration section of the documentation.)

When calling the `lumberjack.Lumberjack.get_handler()` method, you can specify a 'suffix format'. When the handler receives a new log event, it determines the suffix by formatting the time of the event with this string using `time.strftime()`. In this way you can alter the time-period your indices span:

```
from lumberjack import Lumberjack
import logging

lj = Lumberjack(hosts=[...])

day_handler = lj.get_handler(suffix_format='%Y.%m.%d')

my_logger = logging.getLogger(__name__)
my_logger.setLevel(logging.INFO)
my_logger.addHandler(day_handler)
```

In this example a log event that happened on 2014-10-07 would be stored in a different index to one that happened on 2014-10-08. This is useful if you end up with lots of indices that are too small, or too few indices that are too big. (Both of these cases are inefficient.)

In fact you can even customise your handlers based on the frequency of various different events that Lumberjack is attached to:

```
from lumberjack import Lumberjack
import logging

lj = Lumberjack(hosts=[...])

day_handler = lj.get_handler(suffix_format='%Y.%m.%d')
month_handler = lj.get_handler(suffix_format='%Y.%m') # The default

high_volume_logger = logging.getLogger('pageviews')
high_volume_logger.setLevel(logging.INFO)
high_volume_logger.addHandler(day_handler)
```

```
low_volume_logger = logging.getLogger('logins.by.admin')
low_volume_logger.setLevel(logging.INFO)
low_volume_logger.addHandler(month_handler)
```

## 2.4 Post-processors

Post-processors are arbitrary functions which are applied to log entries directly before they go into Elasticsearch. They can be specified easily in the call to `logging.Logger.log()` and family:

```
from logging import getLogger
my_logger = getLogger(__name__)
# Assume we set up Lumberjack and attach it somewhere in the heirarchy of
# this logger.

my_logger.info({'a': 'message'}, {'postprocessors': [some_postprocessor]})
```

In this example, the post-processor `some_postprocessor` will be applied to the logged document immediately before sending to Elasticsearch.

A post-processor is simply a function which is passed the document, and should return a modified version of it. For example, a postprocessor might add the hostname of the machine currently running the program:

```
import socket
def hostname(doc):
    doc['hostname'] = socket.gethostname()

my_logger.info({'a': 'message'}, {'postprocessors': [hostname]})
```

On a system with hostname `load-balancer-01`, this would result in the following document being sent to Elasticsearch:

```
{
  'a': 'message',
  '@timestamp': 1438353254000, # Timestamp added automatically
  'level': 20, # Log level (logging.INFO) added automatically
  'hostname': 'load-balancer-01'
}
```

### 2.4.1 Included post-processors

Lumberjack provides some postprocessors out-of-the-box for you to use, which you can find in `lumberjack.postprocessors`.

#### GeoIP

This post-processor will perform a GeoIP lookup on a field containing an IP address, and include the results in the document:

```
from lumberjack.postprocessors import geoip
my_geoip = geoip(field='ip')

my_logger.info({
  'a': 'message',
```

```
'ip': '128.141.43.1'  
}, {'postprocessors': [my_geoip]})
```

This will result in the following document being sent to Elasticsearch:

```
{  
  'a': 'message',  
  '@timestamp': 1438353254000, # Timestamp added automatically  
  'level': 20, # Log level (logging.INFO) added automatically  
  'ip': '128.141.43.1',  
  'geoip': {  
    'country_code': 'CH',  
    'location': {'lat': 46.1956, 'lon': 6.1481}  
  }  
}
```

## 2.5 API

Module for integrating Elasticsearch into Python's logging library.

**class** `lumberjack.Lumberjack` (*hosts=None, elasticsearch=None, config=None*)

This is the initialisation point for using the lumberjack library.

In the intended use-case, this class is instantiated once and creates handlers for use with Python's logging module.

For each type of log you want to store, you should provide a schema. If you don't, nothing bad will happen, but it makes your cluster rather space-inefficient by default.

You should provide either a list of Elasticsearch hosts, or an already-instantiated `Elasticsearch` object from `elasticsearch-py`.

### Parameters

- **hosts** – A list of Elasticsearch nodes to connect to, in the form `[{'host': '127.0.0.1', 'port': 9200}]`. This is passed directly to `elasticsearch.Elasticsearch`.
- **elasticsearch** – As an alternative to `hosts`, an already-instantiated `elasticsearch.Elasticsearch` object, perhaps with custom transports etc.
- **config** – A configuration for Lumberjack. See the Configuration section in the docs for details.

**get\_handler** (*suffix\_format='%Y.%m'*)

Spawn a new logging handler.

You should use this method to get a `logging.Handler` object to attach to a `logging.Logger` object.

**Note** It is almost definitely unwise to set the formatter of this handler yourself. The integrated formatter prepares documents ready to be inserted into Elasticsearch.

**Parameters** `suffix_format` – The time format string to use as the suffix for the indices. By default your indices will be called, e.g., `generic-logging-2014.09`.

**register\_schema** (*logger, schema*)

Register a new log entry schema.



It is a good idea to register a ‘schema’ for every logger that you attach a handler to. This helps Elasticsearch store the data you provide optimally.

**Note** This method will block until the mapping is registered with Elasticsearch, so you should do it in your initialisation.

#### Parameters

- **logger** – The name of the logger this schema will apply to.
- **schema** – The schema to be used.

#### `trigger_flush()`

Manually trigger a flush of the log queue.

**Note** This is not guaranteed to flush immediately; it merely cancels the wait before the next flush in the `ActionQueue` thread.

#### `lumberjack.get_default_config()`

Get a copy of the default config.

The copy can be modified in-place without affecting the default config itself.

## 2.6 Backend Modules

You shouldn’t need to actually have anything to do with these modules; interaction with the main Lumberjack object should be sufficient.

### 2.6.1 Actions

Provide the `ActionQueue` class.

**class** `lumberjack.actions.ActionQueue` (*elasticsearch, config*)

Hold a queue of actions and a thread to bulk-perform them.

This is instantiated automatically by the `lumberjack.Lumberjack` object. It will keep a queue of indexing actions to be performed in Elasticsearch, and perform them bulk (‘flush’) when one of three things happens:

1. It has waited `interval` seconds without flushing, or
2. The length of the queue has exceeded `max_queue_length`, or
3. A flush is triggered manually.

**Note** You should not need to instantiate, or even interact with, this yourself. It is intended to be wrapped by `lumberjack.Lumberjack`. If you do, for some reason, use this yourself, it is a subclass of `threading.Thread`, so you should call its `start()` method after initialisation.

#### Parameters

- **elasticsearch** – The `elasticsearch.Elasticsearch` object on which to perform the bulk indexing.
- **config** – The Lumberjack config. See the Configuration section in the docs for details.

#### `last_exception`

The last exception raised in the `ActionQueue` thread.

**queue\_index** (*suffix, doc\_type, body, postprocessors=None*)

Queue a new document to be added to Elasticsearch.

If the queue becomes longer than `self.max_queue_length` then a flush is automatically triggered.

#### Parameters

- **suffix** – The suffix of the index into which we should index the document.
- **doc\_type** – The Elasticsearch type of the document to be indexed. Usually this should correspond to a registered schema in Lumberjack.
- **body** – The actual document contents, as a dict.
- **postprocessors** – Any post-processing functions to be run on the document before indexing.

**run** ()

The main method for the ActionQueue thread.

Called by the `start()` method. Not to be called directly.

**trigger\_flush** ()

Manually trigger a flush of the queue.

This is to be called from the main thread, and fires an interrupt in the timeout of the main loop. As such it is not guaranteed to immediately trigger a flush, only to skip the countdown to the next one. This means the flush will happen the next time this thread gets switched to by the Python interpreter.

## 2.6.2 Schemas

Provides SchemaManager class.

**class** `lumberjack.schemas.SchemaManager` (*elasticsearch, config*)

Manage the ‘schemas’ for different types of log data.

A detailed description of schemas is given in the documentation for `lumberjack.Lumberjack.register_schema`.

This class manages a list of schemas registered and ensures that they are processed and passed into Elasticsearch as appropriate.

#### Parameters

- **elasticsearch** – The `elasticsearch.Elasticsearch` object to register mappings and templates with.
- **config** – The Lumberjack config. See the Configuration section in the docs for details.

**register\_schema** (*logger, schema*)

Take a new schema and add it to the roster.

This also automatically parses the schema into a mapping and adds it into the appropriate index template in Elasticsearch.

#### Parameters

- **logger** – The name of the logger which the log data will be emitted on.
- **schema** – The schema data to be processed into a mapping.

## 2.6.3 Handler

Provide classes to fit into the Python logging framework.

**class** `lumberjack.handler.ElasticsearchFormatter` (*fmt=None, datefmt=None*)  
 Formatter which prepares logs for insertion into Elasticsearch.

**format** (*record*)

Add some metadata and deals with string logs.

It adds a `@timestamp` field and a `level` field. `level` contains the loglevel as an integer.

Log data should be in a `dict`, but to be compatible with the generic Python logging recommendations, it can also format log data received as a string. In this case, a `dict` is returned containing a single `message` field, whose data is the string message.

**Parameters** `record` – The `logging.LogRecord` object to be formatted.

**class** `lumberjack.handler.ElasticsearchHandler` (*action\_queue, suffix\_format='%Y.%m'*)  
 Elasticsearch-specific subclass of `logging.LogHandler`.

**Parameters**

- **action\_queue** – A `lumberjack.ActionQueue` object to which the formatted log entries are passed.
- **suffix\_format** – The format from which to generate the time-based index suffixes for Elasticsearch. `strftime()` format.

**emit** (*record*)

Format the log and pass it to an `ElasticsearchContext` instance.

Generates the appropriate index time-suffix based on `self.suffix_format`.

**Parameters** `record` – The `logging.LogRecord` object to format and index.

## 2.7 Developers

### 2.7.1 Contributing

Bug reports, feature requests and other contributions are welcome.

### 2.7.2 Licence

Lumberjack is licensed under the GNU LGPL, which is a set of permissions additional to the GNU GPL.

#### The GNU GPL

GNU GENERAL PUBLIC LICENSE

Version 3, 29 June 2007

Copyright (C) 2007 Free Software Foundation, Inc. <<http://fsf.org/>>  
 Everyone is permitted to copy and distribute verbatim copies  
 of this license document, but changing it is not allowed.

Preamble

The GNU General Public License is a free, copyleft license for software and other kinds of works.

The licenses for most software and other practical works are designed to take away your freedom to share and change the works. By contrast, the GNU General Public License is intended to guarantee your freedom to share and change all versions of a program--to make sure it remains free software for all its users. We, the Free Software Foundation, use the GNU General Public License for most of our software; it applies also to any other work released this way by its authors. You can apply it to your programs, too.

When we speak of free software, we are referring to freedom, not price. Our General Public Licenses are designed to make sure that you have the freedom to distribute copies of free software (and charge for them if you wish), that you receive source code or can get it if you want it, that you can change the software or use pieces of it in new free programs, and that you know you can do these things.

To protect your rights, we need to prevent others from denying you these rights or asking you to surrender the rights. Therefore, you have certain responsibilities if you distribute copies of the software, or if you modify it: responsibilities to respect the freedom of others.

For example, if you distribute copies of such a program, whether gratis or for a fee, you must pass on to the recipients the same freedoms that you received. You must make sure that they, too, receive or can get the source code. And you must show them these terms so they know their rights.

Developers that use the GNU GPL protect your rights with two steps: (1) assert copyright on the software, and (2) offer you this License giving you legal permission to copy, distribute and/or modify it.

For the developers' and authors' protection, the GPL clearly explains that there is no warranty for this free software. For both users' and authors' sake, the GPL requires that modified versions be marked as changed, so that their problems will not be attributed erroneously to authors of previous versions.

Some devices are designed to deny users access to install or run modified versions of the software inside them, although the manufacturer can do so. This is fundamentally incompatible with the aim of protecting users' freedom to change the software. The systematic pattern of such abuse occurs in the area of products for individuals to use, which is precisely where it is most unacceptable. Therefore, we have designed this version of the GPL to prohibit the practice for those products. If such problems arise substantially in other domains, we stand ready to extend this provision to those domains in future versions of the GPL, as needed to protect the freedom of users.

Finally, every program is threatened constantly by software patents. States should not allow patents to restrict development and use of software on general-purpose computers, but in those that do, we wish to avoid the special danger that patents applied to a free program could make it effectively proprietary. To prevent this, the GPL assures that patents cannot be used to render the program non-free.

The precise terms and conditions for copying, distribution and modification follow.

#### TERMS AND CONDITIONS

##### 0. Definitions.

"This License" refers to version 3 of the GNU General Public License.

"Copyright" also means copyright-like laws that apply to other kinds of works, such as semiconductor masks.

"The Program" refers to any copyrightable work licensed under this License. Each licensee is addressed as "you". "Licensees" and "recipients" may be individuals or organizations.

To "modify" a work means to copy from or adapt all or part of the work in a fashion requiring copyright permission, other than the making of an exact copy. The resulting work is called a "modified version" of the earlier work or a work "based on" the earlier work.

A "covered work" means either the unmodified Program or a work based on the Program.

To "propagate" a work means to do anything with it that, without permission, would make you directly or secondarily liable for infringement under applicable copyright law, except executing it on a computer or modifying a private copy. Propagation includes copying, distribution (with or without modification), making available to the public, and in some countries other activities as well.

To "convey" a work means any kind of propagation that enables other parties to make or receive copies. Mere interaction with a user through a computer network, with no transfer of a copy, is not conveying.

An interactive user interface displays "Appropriate Legal Notices" to the extent that it includes a convenient and prominently visible feature that (1) displays an appropriate copyright notice, and (2) tells the user that there is no warranty for the work (except to the extent that warranties are provided), that licensees may convey the work under this License, and how to view a copy of this License. If the interface presents a list of user commands or options, such as a menu, a prominent item in the list meets this criterion.

##### 1. Source Code.

The "source code" for a work means the preferred form of the work for making modifications to it. "Object code" means any non-source form of a work.

A "Standard Interface" means an interface that either is an official standard defined by a recognized standards body, or, in the case of interfaces specified for a particular programming language, one that is widely used among developers working in that language.

The "System Libraries" of an executable work include anything, other than the work as a whole, that (a) is included in the normal form of packaging a Major Component, but which is not part of that Major

Component, and (b) serves only to enable use of the work with that Major Component, or to implement a Standard Interface for which an implementation is available to the public in source code form. A "Major Component", in this context, means a major essential component (kernel, window system, and so on) of the specific operating system (if any) on which the executable work runs, or a compiler used to produce the work, or an object code interpreter used to run it.

The "Corresponding Source" for a work in object code form means all the source code needed to generate, install, and (for an executable work) run the object code and to modify the work, including scripts to control those activities. However, it does not include the work's System Libraries, or general-purpose tools or generally available free programs which are used unmodified in performing those activities but which are not part of the work. For example, Corresponding Source includes interface definition files associated with source files for the work, and the source code for shared libraries and dynamically linked subprograms that the work is specifically designed to require, such as by intimate data communication or control flow between those subprograms and other parts of the work.

The Corresponding Source need not include anything that users can regenerate automatically from other parts of the Corresponding Source.

The Corresponding Source for a work in source code form is that same work.

## 2. Basic Permissions.

All rights granted under this License are granted for the term of copyright on the Program, and are irrevocable provided the stated conditions are met. This License explicitly affirms your unlimited permission to run the unmodified Program. The output from running a covered work is covered by this License only if the output, given its content, constitutes a covered work. This License acknowledges your rights of fair use or other equivalent, as provided by copyright law.

You may make, run and propagate covered works that you do not convey, without conditions so long as your license otherwise remains in force. You may convey covered works to others for the sole purpose of having them make modifications exclusively for you, or provide you with facilities for running those works, provided that you comply with the terms of this License in conveying all material for which you do not control copyright. Those thus making or running the covered works for you must do so exclusively on your behalf, under your direction and control, on terms that prohibit them from making any copies of your copyrighted material outside their relationship with you.

Conveying under any other circumstances is permitted solely under the conditions stated below. Sublicensing is not allowed; section 10 makes it unnecessary.

## 3. Protecting Users' Legal Rights From Anti-Circumvention Law.

No covered work shall be deemed part of an effective technological measure under any applicable law fulfilling obligations under article 11 of the WIPO copyright treaty adopted on 20 December 1996, or

similar laws prohibiting or restricting circumvention of such measures.

When you convey a covered work, you waive any legal power to forbid circumvention of technological measures to the extent such circumvention is effected by exercising rights under this License with respect to the covered work, and you disclaim any intention to limit operation or modification of the work as a means of enforcing, against the work's users, your or third parties' legal rights to forbid circumvention of technological measures.

#### 4. Conveying Verbatim Copies.

You may convey verbatim copies of the Program's source code as you receive it, in any medium, provided that you conspicuously and appropriately publish on each copy an appropriate copyright notice; keep intact all notices stating that this License and any non-permissive terms added in accord with section 7 apply to the code; keep intact all notices of the absence of any warranty; and give all recipients a copy of this License along with the Program.

You may charge any price or no price for each copy that you convey, and you may offer support or warranty protection for a fee.

#### 5. Conveying Modified Source Versions.

You may convey a work based on the Program, or the modifications to produce it from the Program, in the form of source code under the terms of section 4, provided that you also meet all of these conditions:

- a) The work must carry prominent notices stating that you modified it, and giving a relevant date.
- b) The work must carry prominent notices stating that it is released under this License and any conditions added under section 7. This requirement modifies the requirement in section 4 to "keep intact all notices".
- c) You must license the entire work, as a whole, under this License to anyone who comes into possession of a copy. This License will therefore apply, along with any applicable section 7 additional terms, to the whole of the work, and all its parts, regardless of how they are packaged. This License gives no permission to license the work in any other way, but it does not invalidate such permission if you have separately received it.
- d) If the work has interactive user interfaces, each must display Appropriate Legal Notices; however, if the Program has interactive interfaces that do not display Appropriate Legal Notices, your work need not make them do so.

A compilation of a covered work with other separate and independent works, which are not by their nature extensions of the covered work, and which are not combined with it such as to form a larger program, in or on a volume of a storage or distribution medium, is called an "aggregate" if the compilation and its resulting copyright are not used to limit the access or legal rights of the compilation's users beyond what the individual works permit. Inclusion of a covered work

in an aggregate does not cause this License to apply to the other parts of the aggregate.

#### 6. Conveying Non-Source Forms.

You may convey a covered work in object code form under the terms of sections 4 and 5, provided that you also convey the machine-readable Corresponding Source under the terms of this License, in one of these ways:

- a) Convey the object code in, or embodied in, a physical product (including a physical distribution medium), accompanied by the Corresponding Source fixed on a durable physical medium customarily used for software interchange.
- b) Convey the object code in, or embodied in, a physical product (including a physical distribution medium), accompanied by a written offer, valid for at least three years and valid for as long as you offer spare parts or customer support for that product model, to give anyone who possesses the object code either (1) a copy of the Corresponding Source for all the software in the product that is covered by this License, on a durable physical medium customarily used for software interchange, for a price no more than your reasonable cost of physically performing this conveying of source, or (2) access to copy the Corresponding Source from a network server at no charge.
- c) Convey individual copies of the object code with a copy of the written offer to provide the Corresponding Source. This alternative is allowed only occasionally and noncommercially, and only if you received the object code with such an offer, in accord with subsection 6b.
- d) Convey the object code by offering access from a designated place (gratis or for a charge), and offer equivalent access to the Corresponding Source in the same way through the same place at no further charge. You need not require recipients to copy the Corresponding Source along with the object code. If the place to copy the object code is a network server, the Corresponding Source may be on a different server (operated by you or a third party) that supports equivalent copying facilities, provided you maintain clear directions next to the object code saying where to find the Corresponding Source. Regardless of what server hosts the Corresponding Source, you remain obligated to ensure that it is available for as long as needed to satisfy these requirements.
- e) Convey the object code using peer-to-peer transmission, provided you inform other peers where the object code and Corresponding Source of the work are being offered to the general public at no charge under subsection 6d.

A separable portion of the object code, whose source code is excluded from the Corresponding Source as a System Library, need not be included in conveying the object code work.

A "User Product" is either (1) a "consumer product", which means any tangible personal property which is normally used for personal, family, or household purposes, or (2) anything designed or sold for incorporation



into a dwelling. In determining whether a product is a consumer product, doubtful cases shall be resolved in favor of coverage. For a particular product received by a particular user, "normally used" refers to a typical or common use of that class of product, regardless of the status of the particular user or of the way in which the particular user actually uses, or expects or is expected to use, the product. A product is a consumer product regardless of whether the product has substantial commercial, industrial or non-consumer uses, unless such uses represent the only significant mode of use of the product.

"Installation Information" for a User Product means any methods, procedures, authorization keys, or other information required to install and execute modified versions of a covered work in that User Product from a modified version of its Corresponding Source. The information must suffice to ensure that the continued functioning of the modified object code is in no case prevented or interfered with solely because modification has been made.

If you convey an object code work under this section in, or with, or specifically for use in, a User Product, and the conveying occurs as part of a transaction in which the right of possession and use of the User Product is transferred to the recipient in perpetuity or for a fixed term (regardless of how the transaction is characterized), the Corresponding Source conveyed under this section must be accompanied by the Installation Information. But this requirement does not apply if neither you nor any third party retains the ability to install modified object code on the User Product (for example, the work has been installed in ROM).

The requirement to provide Installation Information does not include a requirement to continue to provide support service, warranty, or updates for a work that has been modified or installed by the recipient, or for the User Product in which it has been modified or installed. Access to a network may be denied when the modification itself materially and adversely affects the operation of the network or violates the rules and protocols for communication across the network.

Corresponding Source conveyed, and Installation Information provided, in accord with this section must be in a format that is publicly documented (and with an implementation available to the public in source code form), and must require no special password or key for unpacking, reading or copying.

#### 7. Additional Terms.

"Additional permissions" are terms that supplement the terms of this License by making exceptions from one or more of its conditions. Additional permissions that are applicable to the entire Program shall be treated as though they were included in this License, to the extent that they are valid under applicable law. If additional permissions apply only to part of the Program, that part may be used separately under those permissions, but the entire Program remains governed by this License without regard to the additional permissions.

When you convey a copy of a covered work, you may at your option remove any additional permissions from that copy, or from any part of it. (Additional permissions may be written to require their own removal in certain cases when you modify the work.) You may place

additional permissions on material, added by you to a covered work, for which you have or can give appropriate copyright permission.

Notwithstanding any other provision of this License, for material you add to a covered work, you may (if authorized by the copyright holders of that material) supplement the terms of this License with terms:

- a) Disclaiming warranty or limiting liability differently from the terms of sections 15 and 16 of this License; or
- b) Requiring preservation of specified reasonable legal notices or author attributions in that material or in the Appropriate Legal Notices displayed by works containing it; or
- c) Prohibiting misrepresentation of the origin of that material, or requiring that modified versions of such material be marked in reasonable ways as different from the original version; or
- d) Limiting the use for publicity purposes of names of licensors or authors of the material; or
- e) Declining to grant rights under trademark law for use of some trade names, trademarks, or service marks; or
- f) Requiring indemnification of licensors and authors of that material by anyone who conveys the material (or modified versions of it) with contractual assumptions of liability to the recipient, for any liability that these contractual assumptions directly impose on those licensors and authors.

All other non-permissive additional terms are considered "further restrictions" within the meaning of section 10. If the Program as you received it, or any part of it, contains a notice stating that it is governed by this License along with a term that is a further restriction, you may remove that term. If a license document contains a further restriction but permits relicensing or conveying under this License, you may add to a covered work material governed by the terms of that license document, provided that the further restriction does not survive such relicensing or conveying.

If you add terms to a covered work in accord with this section, you must place, in the relevant source files, a statement of the additional terms that apply to those files, or a notice indicating where to find the applicable terms.

Additional terms, permissive or non-permissive, may be stated in the form of a separately written license, or stated as exceptions; the above requirements apply either way.

## 8. Termination.

You may not propagate or modify a covered work except as expressly provided under this License. Any attempt otherwise to propagate or modify it is void, and will automatically terminate your rights under this License (including any patent licenses granted under the third paragraph of section 11).

However, if you cease all violation of this License, then your

license from a particular copyright holder is reinstated (a) provisionally, unless and until the copyright holder explicitly and finally terminates your license, and (b) permanently, if the copyright holder fails to notify you of the violation by some reasonable means prior to 60 days after the cessation.

Moreover, your license from a particular copyright holder is reinstated permanently if the copyright holder notifies you of the violation by some reasonable means, this is the first time you have received notice of violation of this License (for any work) from that copyright holder, and you cure the violation prior to 30 days after your receipt of the notice.

Termination of your rights under this section does not terminate the licenses of parties who have received copies or rights from you under this License. If your rights have been terminated and not permanently reinstated, you do not qualify to receive new licenses for the same material under section 10.

#### 9. Acceptance Not Required for Having Copies.

You are not required to accept this License in order to receive or run a copy of the Program. Ancillary propagation of a covered work occurring solely as a consequence of using peer-to-peer transmission to receive a copy likewise does not require acceptance. However, nothing other than this License grants you permission to propagate or modify any covered work. These actions infringe copyright if you do not accept this License. Therefore, by modifying or propagating a covered work, you indicate your acceptance of this License to do so.

#### 10. Automatic Licensing of Downstream Recipients.

Each time you convey a covered work, the recipient automatically receives a license from the original licensors, to run, modify and propagate that work, subject to this License. You are not responsible for enforcing compliance by third parties with this License.

An "entity transaction" is a transaction transferring control of an organization, or substantially all assets of one, or subdividing an organization, or merging organizations. If propagation of a covered work results from an entity transaction, each party to that transaction who receives a copy of the work also receives whatever licenses to the work the party's predecessor in interest had or could give under the previous paragraph, plus a right to possession of the Corresponding Source of the work from the predecessor in interest, if the predecessor has it or can get it with reasonable efforts.

You may not impose any further restrictions on the exercise of the rights granted or affirmed under this License. For example, you may not impose a license fee, royalty, or other charge for exercise of rights granted under this License, and you may not initiate litigation (including a cross-claim or counterclaim in a lawsuit) alleging that any patent claim is infringed by making, using, selling, offering for sale, or importing the Program or any portion of it.

#### 11. Patents.

A "contributor" is a copyright holder who authorizes use under this

License of the Program or a work on which the Program is based. The work thus licensed is called the contributor's "contributor version".

A contributor's "essential patent claims" are all patent claims owned or controlled by the contributor, whether already acquired or hereafter acquired, that would be infringed by some manner, permitted by this License, of making, using, or selling its contributor version, but do not include claims that would be infringed only as a consequence of further modification of the contributor version. For purposes of this definition, "control" includes the right to grant patent sublicenses in a manner consistent with the requirements of this License.

Each contributor grants you a non-exclusive, worldwide, royalty-free patent license under the contributor's essential patent claims, to make, use, sell, offer for sale, import and otherwise run, modify and propagate the contents of its contributor version.

In the following three paragraphs, a "patent license" is any express agreement or commitment, however denominated, not to enforce a patent (such as an express permission to practice a patent or covenant not to sue for patent infringement). To "grant" such a patent license to a party means to make such an agreement or commitment not to enforce a patent against the party.

If you convey a covered work, knowingly relying on a patent license, and the Corresponding Source of the work is not available for anyone to copy, free of charge and under the terms of this License, through a publicly available network server or other readily accessible means, then you must either (1) cause the Corresponding Source to be so available, or (2) arrange to deprive yourself of the benefit of the patent license for this particular work, or (3) arrange, in a manner consistent with the requirements of this License, to extend the patent license to downstream recipients. "Knowingly relying" means you have actual knowledge that, but for the patent license, your conveying the covered work in a country, or your recipient's use of the covered work in a country, would infringe one or more identifiable patents in that country that you have reason to believe are valid.

If, pursuant to or in connection with a single transaction or arrangement, you convey, or propagate by procuring conveyance of, a covered work, and grant a patent license to some of the parties receiving the covered work authorizing them to use, propagate, modify or convey a specific copy of the covered work, then the patent license you grant is automatically extended to all recipients of the covered work and works based on it.

A patent license is "discriminatory" if it does not include within the scope of its coverage, prohibits the exercise of, or is conditioned on the non-exercise of one or more of the rights that are specifically granted under this License. You may not convey a covered work if you are a party to an arrangement with a third party that is in the business of distributing software, under which you make payment to the third party based on the extent of your activity of conveying the work, and under which the third party grants, to any of the parties who would receive the covered work from you, a discriminatory patent license (a) in connection with copies of the covered work conveyed by you (or copies made from those copies), or (b) primarily

for and in connection with specific products or compilations that contain the covered work, unless you entered into that arrangement, or that patent license was granted, prior to 28 March 2007.

Nothing in this License shall be construed as excluding or limiting any implied license or other defenses to infringement that may otherwise be available to you under applicable patent law.

12. No Surrender of Others' Freedom.

If conditions are imposed on you (whether by court order, agreement or otherwise) that contradict the conditions of this License, they do not excuse you from the conditions of this License. If you cannot convey a covered work so as to satisfy simultaneously your obligations under this License and any other pertinent obligations, then as a consequence you may not convey it at all. For example, if you agree to terms that obligate you to collect a royalty for further conveying from those to whom you convey the Program, the only way you could satisfy both those terms and this License would be to refrain entirely from conveying the Program.

13. Use with the GNU Affero General Public License.

Notwithstanding any other provision of this License, you have permission to link or combine any covered work with a work licensed under version 3 of the GNU Affero General Public License into a single combined work, and to convey the resulting work. The terms of this License will continue to apply to the part which is the covered work, but the special requirements of the GNU Affero General Public License, section 13, concerning interaction through a network will apply to the combination as such.

14. Revised Versions of this License.

The Free Software Foundation may publish revised and/or new versions of the GNU General Public License from time to time. Such new versions will be similar in spirit to the present version, but may differ in detail to address new problems or concerns.

Each version is given a distinguishing version number. If the Program specifies that a certain numbered version of the GNU General Public License "or any later version" applies to it, you have the option of following the terms and conditions either of that numbered version or of any later version published by the Free Software Foundation. If the Program does not specify a version number of the GNU General Public License, you may choose any version ever published by the Free Software Foundation.

If the Program specifies that a proxy can decide which future versions of the GNU General Public License can be used, that proxy's public statement of acceptance of a version permanently authorizes you to choose that version for the Program.

Later license versions may give you additional or different permissions. However, no additional obligations are imposed on any author or copyright holder as a result of your choosing to follow a later version.

15. Disclaimer of Warranty.

THERE IS NO WARRANTY FOR THE PROGRAM, TO THE EXTENT PERMITTED BY APPLICABLE LAW. EXCEPT WHEN OTHERWISE STATED IN WRITING THE COPYRIGHT HOLDERS AND/OR OTHER PARTIES PROVIDE THE PROGRAM "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESSED OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE. THE ENTIRE RISK AS TO THE QUALITY AND PERFORMANCE OF THE PROGRAM IS WITH YOU. SHOULD THE PROGRAM PROVE DEFECTIVE, YOU ASSUME THE COST OF ALL NECESSARY SERVICING, REPAIR OR CORRECTION.

16. Limitation of Liability.

IN NO EVENT UNLESS REQUIRED BY APPLICABLE LAW OR AGREED TO IN WRITING WILL ANY COPYRIGHT HOLDER, OR ANY OTHER PARTY WHO MODIFIES AND/OR CONVEYS THE PROGRAM AS PERMITTED ABOVE, BE LIABLE TO YOU FOR DAMAGES, INCLUDING ANY GENERAL, SPECIAL, INCIDENTAL OR CONSEQUENTIAL DAMAGES ARISING OUT OF THE USE OR INABILITY TO USE THE PROGRAM (INCLUDING BUT NOT LIMITED TO LOSS OF DATA OR DATA BEING RENDERED INACCURATE OR LOSSES SUSTAINED BY YOU OR THIRD PARTIES OR A FAILURE OF THE PROGRAM TO OPERATE WITH ANY OTHER PROGRAMS), EVEN IF SUCH HOLDER OR OTHER PARTY HAS BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.

17. Interpretation of Sections 15 and 16.

If the disclaimer of warranty and limitation of liability provided above cannot be given local legal effect according to their terms, reviewing courts shall apply local law that most closely approximates an absolute waiver of all civil liability in connection with the Program, unless a warranty or assumption of liability accompanies a copy of the Program in return for a fee.

END OF TERMS AND CONDITIONS

How to Apply These Terms to Your New Programs

If you develop a new program, and you want it to be of the greatest possible use to the public, the best way to achieve this is to make it free software which everyone can redistribute and change under these terms.

To do so, attach the following notices to the program. It is safest to attach them to the start of each source file to most effectively state the exclusion of warranty; and each file should have at least the "copyright" line and a pointer to where the full notice is found.

```
{one line to give the program's name and a brief idea of what it does.}
Copyright (C) {year} {name of author}
```

This program is free software: you can redistribute it and/or modify it under the terms of the GNU General Public License as published by the Free Software Foundation, either version 3 of the License, or (at your option) any later version.

This program is distributed in the hope that it will be useful, but WITHOUT ANY WARRANTY; without even the implied warranty of MERCHANTABILITY or FITNESS FOR A PARTICULAR PURPOSE. See the GNU General Public License for more details.

You should have received a copy of the GNU General Public License along with this program. If not, see <<http://www.gnu.org/licenses/>>.

Also add information on how to contact you by electronic and paper mail.

If the program does terminal interaction, make it output a short notice like this when it starts in an interactive mode:

```
{project} Copyright (C) {year} {fullname}
This program comes with ABSOLUTELY NO WARRANTY; for details type `show w'.
This is free software, and you are welcome to redistribute it
under certain conditions; type `show c' for details.
```

The hypothetical commands `show w' and `show c' should show the appropriate parts of the General Public License. Of course, your program's commands might be different; for a GUI interface, you would use an "about box".

You should also get your employer (if you work as a programmer) or school, if any, to sign a "copyright disclaimer" for the program, if necessary. For more information on this, and how to apply and follow the GNU GPL, see <http://www.gnu.org/licenses/>.

The GNU General Public License does not permit incorporating your program into proprietary programs. If your program is a subroutine library, you may consider it more useful to permit linking proprietary applications with the library. If this is what you want to do, use the GNU Lesser General Public License instead of this License. But first, please read <http://www.gnu.org/philosophy/why-not-lgpl.html>.

## The GNU LGPL

GNU LESSER GENERAL PUBLIC LICENSE  
Version 3, 29 June 2007

Copyright (C) 2007 Free Software Foundation, Inc. <http://fsf.org/>  
Everyone is permitted to copy and distribute verbatim copies  
of this license document, but changing it is not allowed.

This version of the GNU Lesser General Public License incorporates the terms and conditions of version 3 of the GNU General Public License, supplemented by the additional permissions listed below.

### 0. Additional Definitions.

As used herein, "this License" refers to version 3 of the GNU Lesser General Public License, and the "GNU GPL" refers to version 3 of the GNU General Public License.

"The Library" refers to a covered work governed by this License, other than an Application or a Combined Work as defined below.

An "Application" is any work that makes use of an interface provided by the Library, but which is not otherwise based on the Library. Defining a subclass of a class defined by the Library is deemed a mode of using an interface provided by the Library.

A "Combined Work" is a work produced by combining or linking an Application with the Library. The particular version of the Library

with which the Combined Work was made is also called the "Linked Version".

The "Minimal Corresponding Source" for a Combined Work means the Corresponding Source for the Combined Work, excluding any source code for portions of the Combined Work that, considered in isolation, are based on the Application, and not on the Linked Version.

The "Corresponding Application Code" for a Combined Work means the object code and/or source code for the Application, including any data and utility programs needed for reproducing the Combined Work from the Application, but excluding the System Libraries of the Combined Work.

### 1. Exception to Section 3 of the GNU GPL.

You may convey a covered work under sections 3 and 4 of this License without being bound by section 3 of the GNU GPL.

### 2. Conveying Modified Versions.

If you modify a copy of the Library, and, in your modifications, a facility refers to a function or data to be supplied by an Application that uses the facility (other than as an argument passed when the facility is invoked), then you may convey a copy of the modified version:

- a) under this License, provided that you make a good faith effort to ensure that, in the event an Application does not supply the function or data, the facility still operates, and performs whatever part of its purpose remains meaningful, or
- b) under the GNU GPL, with none of the additional permissions of this License applicable to that copy.

### 3. Object Code Incorporating Material from Library Header Files.

The object code form of an Application may incorporate material from a header file that is part of the Library. You may convey such object code under terms of your choice, provided that, if the incorporated material is not limited to numerical parameters, data structure layouts and accessors, or small macros, inline functions and templates (ten or fewer lines in length), you do both of the following:

- a) Give prominent notice with each copy of the object code that the Library is used in it and that the Library and its use are covered by this License.
- b) Accompany the object code with a copy of the GNU GPL and this license document.

### 4. Combined Works.

You may convey a Combined Work under terms of your choice that, taken together, effectively do not restrict modification of the portions of the Library contained in the Combined Work and reverse engineering for debugging such modifications, if you also do each of the following:



- a) Give prominent notice with each copy of the Combined Work that the Library is used in it and that the Library and its use are covered by this License.
- b) Accompany the Combined Work with a copy of the GNU GPL and this license document.
- c) For a Combined Work that displays copyright notices during execution, include the copyright notice for the Library among these notices, as well as a reference directing the user to the copies of the GNU GPL and this license document.
- d) Do one of the following:
  - 0) Convey the Minimal Corresponding Source under the terms of this License, and the Corresponding Application Code in a form suitable for, and under terms that permit, the user to recombine or relink the Application with a modified version of the Linked Version to produce a modified Combined Work, in the manner specified by section 6 of the GNU GPL for conveying Corresponding Source.
  - 1) Use a suitable shared library mechanism for linking with the Library. A suitable mechanism is one that (a) uses at run time a copy of the Library already present on the user's computer system, and (b) will operate properly with a modified version of the Library that is interface-compatible with the Linked Version.
- e) Provide Installation Information, but only if you would otherwise be required to provide such information under section 6 of the GNU GPL, and only to the extent that such information is necessary to install and execute a modified version of the Combined Work produced by recombining or relinking the Application with a modified version of the Linked Version. (If you use option 4d0, the Installation Information must accompany the Minimal Corresponding Source and Corresponding Application Code. If you use option 4d1, you must provide the Installation Information in the manner specified by section 6 of the GNU GPL for conveying Corresponding Source.)

#### 5. Combined Libraries.

You may place library facilities that are a work based on the Library side by side in a single library together with other library facilities that are not Applications and are not covered by this License, and convey such a combined library under terms of your choice, if you do both of the following:

- a) Accompany the combined library with a copy of the same work based on the Library, uncombined with any other library facilities, conveyed under the terms of this License.
- b) Give prominent notice with the combined library that part of it is a work based on the Library, and explaining where to find the accompanying uncombined form of the same work.

#### 6. Revised Versions of the GNU Lesser General Public License.

The Free Software Foundation may publish revised and/or new versions of the GNU Lesser General Public License from time to time. Such new versions will be similar in spirit to the present version, but may differ in detail to address new problems or concerns.

Each version is given a distinguishing version number. If the Library as you received it specifies that a certain numbered version of the GNU Lesser General Public License "or any later version" applies to it, you have the option of following the terms and conditions either of that published version or of any later version published by the Free Software Foundation. If the Library as you received it does not specify a version number of the GNU Lesser General Public License, you may choose any version of the GNU Lesser General Public License ever published by the Free Software Foundation.

If the Library as you received it specifies that a proxy can decide whether future versions of the GNU Lesser General Public License shall apply, that proxy's public statement of acceptance of any version is permanent authorization for you to choose that version for the Library.

|

`lumberjack`, [12](#)  
`lumberjack.actions`, [13](#)  
`lumberjack.handler`, [15](#)  
`lumberjack.schemas`, [14](#)



## A

ActionQueue (class in lumberjack.actions), 13

## E

ElasticsearchFormatter (class in lumberjack.handler), 15

ElasticsearchHandler (class in lumberjack.handler), 15

emit() (lumberjack.handler.ElasticsearchHandler method), 15

## F

format() (lumberjack.handler.ElasticsearchFormatter method), 15

## G

get\_default\_config() (in module lumberjack), 13

get\_handler() (lumberjack.Lumberjack method), 12

## L

last\_exception (lumberjack.actions.ActionQueue attribute), 13

Lumberjack (class in lumberjack), 12

lumberjack (module), 12

lumberjack.actions (module), 13

lumberjack.handler (module), 15

lumberjack.schemas (module), 14

## Q

queue\_index() (lumberjack.actions.ActionQueue method), 13

## R

register\_schema() (lumberjack.Lumberjack method), 12

register\_schema() (lumberjack.schemas.SchemaManager method), 14

run() (lumberjack.actions.ActionQueue method), 14

## S

SchemaManager (class in lumberjack.schemas), 14

## T

trigger\_flush() (lumberjack.actions.ActionQueue method), 14

trigger\_flush() (lumberjack.Lumberjack method), 13