

---

# **structlog Documentation**

***Release***

**Author**

**May 15, 2017**



---

## Contents

---

<b>1</b>	<b>Easier Logging</b>	<b>3</b>
<b>2</b>	<b>Data Binding</b>	<b>5</b>
<b>3</b>	<b>Powerful Pipelines</b>	<b>7</b>
<b>4</b>	<b>Formatting</b>	<b>9</b>
<b>5</b>	<b>Output</b>	<b>11</b>
<b>6</b>	<b>User's Guide</b>	<b>13</b>
6.1	Basics . . . . .	13
6.2	Integration with Existing Systems . . . . .	29
6.3	Advanced Topics . . . . .	37
<b>7</b>	<b>API Reference</b>	<b>39</b>
7.1	API Reference . . . . .	39
<b>8</b>	<b>Project Information</b>	<b>53</b>
8.1	Backward Compatibility . . . . .	53
8.2	How To Contribute . . . . .	53
8.3	License and Hall of Fame . . . . .	56
8.4	Authors . . . . .	56
8.5	Changelog . . . . .	56
<b>9</b>	<b>Indices and tables</b>	<b>63</b>
	<b>Python Module Index</b>	<b>65</b>



Release v17.2.0 (*What's new?*).

`structlog` makes logging in Python less painful and more powerful by adding structure to your log entries.

It's up to you whether you want `structlog` to take care about the **output** of your log entries or whether you prefer to **forward** them to an existing logging system like the standard library's `logging` module. *No monkey patching* involved in either case.



# CHAPTER 1

---

## Easier Logging

---

You can stop writing prose and start thinking in terms of an event that happens in the context of key/value pairs:

```
>>> from structlog import get_logger
>>> log = get_logger()
>>> log.info("key_value_logging", out_of_the_box=True, effort=0)
2016-04-20 16:20.13 key_value_logging          effort=0 out_of_the_box=True
```

Each log entry is a meaningful dictionary instead of an opaque string now!





## CHAPTER 2

---

### Data Binding

---

Since log entries are dictionaries, you can start binding and re-binding key/value pairs to your loggers to ensure they are present in every following logging call:

```
>>> log = log.bind(user="anonymous", some_key=23)
>>> log = log.bind(user="hynek", another_key=42)
>>> log.info("user.logged_in", happy=True)
2016-04-20 16:20.13 user.logged_in          another_key=42 happy=True some_
↳key=23 user='hynek'
```



## CHAPTER 3

---

### Powerful Pipelines

---

Each log entry goes through a **processor pipeline** that is just a chain of functions that receive a dictionary and return a new dictionary that gets fed into the next function. That allows for simple but powerful data manipulation:

```
def timestamper(logger, log_method, event_dict):  
    """Add a timestamp to each log entry."""  
    event_dict["timestamp"] = time.time()  
    return event_dict
```

There are plenty of processors for most common tasks coming with `structlog`:

- Collectors of **call stack information** (“How did this log entry happen?”),
- ...and **exceptions** (“What happened?”).
- Unicode encoders/decoders.
- Flexible **timestamping**.



`structlog` is completely flexible about *how* the resulting log entry is emitted. Since each log entry is a dictionary, it can be formatted to **any** format:

- A colorful key/value format for *local development*,
- **JSON** for easy parsing,
- or some standard format you have parsers for like `nginx` or `Apache httpd`.

Internally, formatters are processors whose return value (usually a string) is passed into loggers that are responsible for the output of your message. `structlog` comes with multiple useful formatters out of-the-box.



`structlog` is also very flexible with the final output of your log entries:

- A **built-in** lightweight printer like in the examples above. Easy to use and fast.
- Use the **standard library**'s or **Twisted**'s logging modules for compatibility. In this case `structlog` works like a wrapper that formats a string and passes them off into existing systems that won't ever know that `structlog` even exists. Or the other way round: `structlog` comes with a logging formatter that allows for processing third party log records.
- Don't format it to a string at all! `structlog` passes you a dictionary and you can do with it whatever you want. Reported uses cases are sending them out via network or saving them in a database.





## Basics

### Why...

#### ...Structured Logging?

I believe the widespread use of format strings in logging is based on two presumptions:

- The first level consumer of a log message is a human.
- The programmer knows what information is needed to debug an issue.

I believe these presumptions are **no longer correct** in server side software.

—Paul Querna

Structured logging means that you don't write hard-to-parse and hard-to-keep-consistent prose in your logs but that you log *events* that happen in a *context* instead.

#### ...structlog?

Because it's easy and you don't have to replace your underlying logger (but you can!) – you just add structure to your log entries and format them to strings before they hit your real loggers.

`structlog` supports you with accepting key/value pairs as arguments, building your context as you go (e.g. if a user logs in, you bind their user name to your current logger) and log events when they happen (i.e. the user does something log-worthy):

```
>>> from structlog import get_logger
>>> log = get_logger()
>>> log = log.bind(user='anonymous', some_key=23)
>>> log = log.bind(user='hynek', another_key=42)
>>> log.info('user.logged_in', happy=True)
2016-04-20 16:20.13 user.logged_in          another_key=42 happy=True some_
key=23 user='hynek'
```

This ability to bind key/values pairs to a logger frees you from using conditionals, closures, or boilerplate methods to log out all relevant data.

Additionally, `structlog` offers you a flexible way to *filter* and *modify* your log entries using so called *processors* before the entry is passed to your real logger. The possibilities include *logging in JSON*, adding arbitrary meta data like *timestamps*, counting events as metrics, or *dropping log entries* caused by your monitoring system.

`structlog` is also flexible enough to allow transparent *thread local* storage for your context if you don't like the idea of local bindings as in the example above.

## Getting Started

### Installation

`structlog` can be easily installed using:

```
$ pip install structlog
```

If you'd like colorful output in development (you know you do!), install using:

```
$ pip install structlog[dev]
```

### Your First Log Entry

A lot of effort went into making `structlog` accessible without reading pages of documentation. And indeed, the simplest possible usage looks like this:

```
>>> import structlog
>>> log = structlog.get_logger()
>>> log.msg("greeted", whom="world", more_than_a_string=[1, 2, 3])
2016-09-17 10:13.45 greeted                                more_than_a_string=[1, 2, 3] whom=
↳ 'world'
```

Here, `structlog` takes full advantage of its hopefully useful default settings:

- Output is sent to `standard out` instead of exploding into the user's face. Yes, that seems a rather controversial attitude towards logging.
- All keywords are formatted using `structlog.dev.ConsoleRenderer`. That in turn uses `repr()` to serialize all values to strings. Thus, it's easy to add support for logging of your own objects<sup>0</sup>.
- If you have `colorama` installed, it's even rendered in nice *colors*.

It should be noted that even in most complex logging setups the example would still look just like that thanks to *Configuration*.

---

**Note:** For brevity and to enable doctests, all further examples in `structlog`'s documentation use the more simplistic `structlog.processors.KeyValueRenderer()` without timestamps.

---

There you go, structured logging! However, this alone wouldn't warrant its own package. After all, there's even a *recipe* on structured logging for the standard library. So let's go a step further.

---

<sup>0</sup> In production, you're more likely to use `JSONRenderer` that can also be customized using a `__structlog__` method so you don't have to change your `repr` methods to something they weren't originally intended for.

## Building a Context

Imagine a hypothetical web application that wants to log out all relevant data with just the API from above:

```
from structlog import get_logger

log = get_logger()

def view(request):
    user_agent = request.get('HTTP_USER_AGENT', 'UNKNOWN')
    peer_ip = request.client_addr
    if something:
        log.msg('something', user_agent=user_agent, peer_ip=peer_ip)
        return 'something'
    elif something_else:
        log.msg('something_else', user_agent=user_agent, peer_ip=peer_ip)
        return 'something_else'
    else:
        log.msg('else', user_agent=user_agent, peer_ip=peer_ip)
        return 'else'
```

The calls themselves are nice and straight to the point, however you're repeating yourself all over the place. At this point, you'll be tempted to write a closure like

```
def log_closure(event):
    log.msg(event, user_agent=user_agent, peer_ip=peer_ip)
```

inside of the view. Problem solved? Not quite. What if the parameters are introduced step by step? Do you really want to have a logging closure in each of your views?

Let's have a look at a better approach:

```
from structlog import get_logger

logger = get_logger()

def view(request):
    log = logger.bind(
        user_agent=request.get('HTTP_USER_AGENT', 'UNKNOWN'),
        peer_ip=request.client_addr,
    )
    foo = request.get('foo')
    if foo:
        log = log.bind(foo=foo)
    if something:
        log.msg('something')
        return 'something'
    elif something_else:
        log.msg('something_else')
        return 'something_else'
    else:
        log.msg('else')
        return 'else'
```

Suddenly your logger becomes your closure!

For structlog, a log entry is just a dictionary called *event dict[ionary]*:

- You can pre-build a part of the dictionary step by step. These pre-saved values are called the *context*.
- As soon as an *event* happens – which is a dictionary too – it is merged together with the *context* to an *event dict* and logged out.
- If you don’t like the concept of pre-building a context: just don’t! Convenient key-value-based logging is great to have on its own.
- To keep as much order of the keys as possible, an `collections.OrderedDict` is used for the context by default.
- The recommended way of binding values is the one in these examples: creating new loggers with a new context. If you’re okay with giving up immutable local state for convenience, you can also use *thread/greenlet local storage* for the context.

## structlog and Standard Library’s logging

structlog’s primary application isn’t printing though. Instead, it’s intended to wrap your *existing* loggers and **add structure** and *incremental context building* to them. For that, structlog is *completely* agnostic of your underlying logger – you can use it with any logger you like.

The most prominent example of such an ‘existing logger’ is without doubt the logging module in the standard library. To make this common case as simple as possible, structlog comes with some tools to help you:

```
>>> import logging
>>> logging.basicConfig()
>>> from structlog import get_logger, configure
>>> from structlog.stdlib import LoggerFactory
>>> configure(logger_factory=LoggerFactory())
>>> log = get_logger()
>>> log.warn('it works!', difficulty='easy')
WARNING:structlog...:difficulty='easy' event='it works!'
```

In other words, you tell structlog that you would like to use the standard library logger factory and keep calling `get_logger()` like before.

Since structlog is mainly used together with standard library’s logging, there’s *more* goodness to make it as fast and convenient as possible.

## Liked what you saw?

Now you’re all set for the rest of the user’s guide. If you want to see more code, make sure to check out the *Examples*!

## Loggers

### Bound Loggers

The center of structlog is the immutable log wrapper *BoundLogger*.

All it does is:

- Keep a *context dictionary* and a *logger* that it’s wrapping,
- recreate itself with (optional) *additional* context data (the `bind()` and `new()` methods),
- recreate itself with *less* data (`unbind()`),

- and finally relay *all* other method calls to the wrapped logger<sup>0</sup> after processing the log entry with the configured chain of *processors*.

You won't be instantiating it yourself though. For that there is the `structlog.wrap_logger()` function (or the convenience function `structlog.get_logger()` we'll discuss in a minute):

```
>>> from structlog import wrap_logger
>>> class PrintLogger(object):
...     def msg(self, message):
...         print(message)
>>> def proc(logger, method_name, event_dict):
...     print('I got called with', event_dict)
...     return repr(event_dict)
>>> log = wrap_logger(PrintLogger(), processors=[proc], context_class=dict)
>>> log2 = log.bind(x=42)
>>> log == log2
False
>>> log.msg('hello world')
I got called with {'event': 'hello world'}
{'event': 'hello world'}
>>> log2.msg('hello world')
I got called with {'x': 42, 'event': 'hello world'}
{'x': 42, 'event': 'hello world'}
>>> log3 = log2.unbind('x')
>>> log == log3
True
>>> log3.msg('nothing bound anymore', foo='but you can structure the event too')
I got called with {'foo': 'but you can structure the event too', 'event': 'nothing_
↳bound anymore'}
{'foo': 'but you can structure the event too', 'event': 'nothing bound anymore'}
```

As you can see, it accepts one mandatory and a few optional arguments:

**logger** The one and only positional argument is the logger that you want to wrap and to which the log entries will be proxied. If you wish to use a *configured logger factory*, set it to *None*.

**processors** A list of callables that can *filter*, *mutate*, and *format* the log entry before it gets passed to the wrapped logger.

Default is `[format_exc_info(), KeyValueRenderer]`.

**context\_class** The class to save your context in. Particularly useful for *thread local context storage*.

Default is `collections.OrderedDict`.

Additionally, the following arguments are allowed too:

**wrapper\_class** A class to use instead of *BoundLogger* for wrapping. This is useful if you want to sub-class *BoundLogger* and add custom logging methods. *BoundLogger*'s *bind*/*new* methods are sub-classing friendly so you won't have to re-implement them. Please refer to the *related example* for how this may look.

**initial\_values** The values that new wrapped loggers are automatically constructed with. Useful for example if you want to have the module name as part of the context.

---

**Note:** Free your mind from the preconception that log entries have to be serialized to strings eventually. All *structlog* cares about is a *dictionary* of *keys* and *values*. What happens to it depends on the logger you wrap and your processors alone.

---

<sup>0</sup> Since this is slightly magicy, *structlog* comes with concrete loggers for the *Standard Library Logging* and *Twisted* that offer you explicit APIs for the supported logging methods but behave identically like the generic *BoundLogger* otherwise.

This gives you the power to log directly to databases, log aggregation servers, web services, and whatnot.

---

### Printing and Testing

To save you the hassle of using standard library logging for simple standard out logging, `structlog` ships a `PrintLogger` that can log into arbitrary files – including standard out (which is the default if no file is passed into the constructor):

```
>>> from structlog import PrintLogger
>>> PrintLogger().info('hello world!')
hello world!
```

It's handy for both examples and in combination with tools like `runit` or `stdout/stderr-forwarding`.

Additionally – mostly for unit testing – `structlog` also ships with a logger that just returns whatever it gets passed into it: `ReturnLogger`.

```
>>> from structlog import ReturnLogger
>>> ReturnLogger().msg(42) == 42
True
>>> obj = ['hi']
>>> ReturnLogger().msg(obj) is obj
True
>>> ReturnLogger().msg('hello', when='again')
(('hello',), {'when': 'again'})
```

## Configuration

### Global Defaults

To make logging as unintrusive and straight-forward to use as possible, `structlog` comes with a plethora of configuration options and convenience functions. Let me start at the end and introduce you to the ultimate convenience function that relies purely on configuration: `structlog.get_logger()` (and its Twisted-friendly alias `structlog.getLogger()`).

The goal is to reduce your per-file logging boilerplate to:

```
from structlog import get_logger
logger = get_logger()
```

while still giving you the full power via configuration.

To achieve that you'll have to call `structlog.configure()` on app initialization (of course, only if you're not content with the defaults). The *example* from the previous chapter could thus have been written as following:

```
>>> configure(processors=[proc], context_class=dict)
>>> log = wrap_logger(PrintLogger())
>>> log.msg('hello world')
I got called with {'event': 'hello world'}
{'event': 'hello world'}
```

In fact, it could even be written like

```
>>> configure(processors=[proc], context_class=dict)
>>> log = get_logger()
>>> log.msg('hello world')
I got called with {'event': 'hello world'}
{'event': 'hello world'}
```

because `PrintLogger` is the default `LoggerFactory` used (see [Logger Factories](#)).

`structlog` tries to behave in the least surprising way when it comes to handling defaults and configuration:

1. Arguments passed to `structlog.wrap_logger()` *always* take the highest precedence over configuration. That means that you can overwrite whatever you've configured for each logger respectively.
2. If you leave them on `None`, `structlog` will check whether you've configured default values using `structlog.configure()` and uses them if so.
3. If you haven't configured or passed anything at all, the default fallback values are used which means `collections.OrderedDict` for context and [`StackInfoRenderer`, `format_exc_info()`, `KeyValueRenderer`] for the processor chain, and `False` for `cache_logger_on_first_use`.

If necessary, you can always reset your global configuration back to default values using `structlog.reset_defaults()`. That can be handy in tests.

---

**Note:** Since you will call `structlog.wrap_logger()` (or one of the `get_logger()` functions) most likely at import time and thus before you had a chance to configure `structlog`, they return a **proxy** that returns a correct wrapped logger on first `bind()/new()`.

Therefore, you must not call `new()` or `bind()` in module scope! Use `get_logger()`'s `initial_values` to achieve pre-populated contexts.

To enable you to log with the module-global logger, it will create a temporary `BoundLogger` and relay the log calls to it on *each call*. Therefore if you have nothing to bind but intend to do lots of log calls in a function, it makes sense performance-wise to create a local logger by calling `bind()` or `new()` without any parameters. See also [Performance](#).

---

## Logger Factories

To make `structlog.get_logger()` work, one needs one more option that hasn't been discussed yet: `logger_factory`.

It is a callable that returns the logger that gets wrapped and returned. In the simplest case, it's a function that returns a logger – or just a class. But you can also pass in an instance of a class with a `__call__` method for more complicated setups.

New in version 0.4.0: `structlog.get_logger()` can optionally take positional parameters.

These will be passed to the logger factories. For example, if you use `run structlog.get_logger('a name')` and configure `structlog` to use the standard library `LoggerFactory` which has support for positional parameters, the returned logger will have the name `'a name'`.

When writing custom logger factories, they should always accept positional parameters even if they don't use them. That makes sure that loggers are interchangeable.

For the common cases of standard library logging and Twisted logging, `structlog` comes with two factories built right in:

- `structlog.stdlib.LoggerFactory`
- `structlog.twisted.LoggerFactory`

So all it takes to use structlog with standard library logging is this:

```
>>> from structlog import get_logger, configure
>>> from structlog.stdlib import LoggerFactory
>>> configure(logger_factory=LoggerFactory())
>>> log = get_logger()
>>> log.critical('this is too easy!')
event='this is too easy!'
```

By using structlog's `structlog.stdlib.LoggerFactory`, it is also ensured that variables like function names and line numbers are expanded correctly in your log format.

The *Twisted example* shows how easy it is for Twisted.

---

**Note:** `LoggerFactory()`-style factories always need to get passed as *instances* like in the examples above. While neither allows for customization using parameters yet, they may do so in the future.

---

Calling `structlog.get_logger()` without configuration gives you a perfectly useful `structlog.PrintLogger` with the default values explained above. I don't believe silent loggers are a sensible default.

## Where to Configure

The best place to perform your configuration varies with applications and frameworks. Ideally as late as possible but *before* non-framework (i.e. your) code is executed. If you use standard library's logging, it makes sense to configure them next to each other.

**Django** Django has to date unfortunately no concept of an application assembler or “app is done” hooks. Therefore the bottom of your `settings.py` will have to do.

**Flask** See [Logging Application Errors](#).

**Pyramid** [Application constructor](#).

**Twisted** The [plugin definition](#) is the best place. If your app is not a plugin, put it into your `tac` file (and then [learn](#) about plugins).

If you have no choice but *have* to configure on import time in module-global scope, or can't rule out for other reasons that that your `structlog.configure()` gets called more than once, structlog offers `structlog.configure_once()` that raises a warning if structlog has been configured before (no matter whether using `structlog.configure()` or `configure_once()`) but doesn't change anything.

## Thread Local Context

### Immutability

You should call some functions with some arguments.

—David Reid

The behavior of copying itself, adding new values, and returning the result is useful for applications that keep somehow their own context using classes or closures. Twisted is a *fine example* for that. Another possible approach is passing wrapped loggers around or log only within your view where you gather errors and events using return codes and exceptions. If you are willing to do that, you should stick to it because [immutable state](#) is a very good thing<sup>0</sup>. Sooner or later, global state and mutable data lead to unpleasant surprises.

---

<sup>0</sup> In the spirit of Python's 'consenting adults', structlog doesn't enforce the immutability with technical means. However, if you don't meddle with undocumented data, the objects can be safely considered immutable.



However, in the case of conventional web development, we realize that passing loggers around seems rather cumbersome, intrusive, and generally against the mainstream culture. And since it's more important that people actually *use* structlog than to be pure and snobby, structlog contains a dirty but convenient trick: thread local context storage which you may already know from [Flask](#):

Thread local storage makes your logger's context global but *only within the current thread*<sup>0</sup>. In the case of web frameworks this usually means that your context becomes global to the current request.

The following explanations may sound a bit confusing at first but the [Flask example](#) illustrates how simple and elegant this works in practice.

## Wrapped Dicts

In order to make your context thread local, structlog ships with a function that can wrap any dict-like class to make it usable for thread local storage: `structlog.threadlocal.wrap_dict()`.

Within one thread, every instance of the returned class will have a *common* instance of the wrapped dict-like class:

```
>>> from structlog.threadlocal import wrap_dict
>>> WrappedDictClass = wrap_dict(dict)
>>> d1 = WrappedDictClass({'a': 1})
>>> d2 = WrappedDictClass({'b': 2})
>>> d3 = WrappedDictClass()
>>> d3['c'] = 3
>>> d1 is d3
False
>>> d1 == d2 == d3 == WrappedDictClass()
True
>>> d3
<WrappedDict-...({'a': 1, 'b': 2, 'c': 3})>
```

Then use an instance of the generated class as the context class:

```
configure(context_class=WrappedDictClass())
```

**Note: Remember:** the instance of the class *doesn't* matter. Only the class *type* matters because *all* instances of one class *share* the *same* data.

`structlog.threadlocal.wrap_dict()` returns always a completely *new* wrapped class:

```
>>> from structlog.threadlocal import wrap_dict
>>> WrappedDictClass = wrap_dict(dict)
>>> AnotherWrappedDictClass = wrap_dict(dict)
>>> WrappedDictClass() != AnotherWrappedDictClass()
True
>>> WrappedDictClass.__name__
WrappedDict-41e8382d-bee5-430e-ad7d-133c844695cc
>>> AnotherWrappedDictClass.__name__
WrappedDict-e0fc330e-e5eb-42ee-bcec-ffd7bd09ad09
```

In order to be able to bind values temporarily to a logger, `structlog.threadlocal` comes with a context manager: `tmp_bind()`:

<sup>0</sup> Special care has been taken to detect and support greenlets properly.

```
>>> log.bind(x=42)
<BoundLogger(context=<WrappedDict-...({'x': 42})>, ...)>
>>> log.msg("event!")
x=42 event='event!'
>>> with tmp_bind(log, x=23, y="foo") as tmp_log:
...     tmp_log.msg("another event!")
x=23 y='foo' event='another event!'
>>> log.msg("one last event!")
x=42 event='one last event!'
```

The state before the `with` statement is saved and restored once it's left.

If you want to detach a logger from thread local data, there's `structlog.threadlocal.as_immutable()`.

### Downsides & Caveats

The convenience of having a thread local context comes at a price though:

#### Warning:

- If you can't rule out that your application re-uses threads, you *must* remember to **initialize your thread local context** at the start of each request using `new()` (instead of `bind()`). Otherwise you may start a new request with the context still filled with data from the request before.
- **Don't** stop assigning the results of your `bind()`s and `new()`s!

#### Do:

```
log = log.new(y=23)
log = log.bind(x=42)
```

#### Don't:

```
log.new(y=23)
log.bind(x=42)
```

Although the state is saved in a global data structure, you still need the global wrapped logger produce a real bound logger. Otherwise each log call will result in an instantiation of a temporary BoundLogger. See [Configuration](#) for more details.

The general sentiment against thread locals is that they're hard to test. In this case we feel like this is an acceptable trade-off. You can easily write deterministic tests using a call-capturing processor if you use the API properly (cf. warning above).

This big red box is also what separates immutable local from mutable global data.

### Processors

The true power of structlog lies in its *combinable log processors*. A log processor is a regular callable, i.e. a function or an instance of a class with a `__call__()` method.

### Chains

The *processor chain* is a list of processors. Each processors receives three positional arguments:

**logger** Your wrapped logger object. For example `logging.Logger`.

**method\_name** The name of the wrapped method. If you called `log.warn('foo')`, it will be `"warn"`.

**event\_dict** Current context together with the current event. If the context was `{'a': 42}` and the event is `"foo"`, the initial `event_dict` will be `{'a': 42, 'event': 'foo'}`.

The return value of each processor is passed on to the next one as `event_dict` until finally the return value of the last processor gets passed into the wrapped logging method.

## Examples

If you set up your logger like:

```
from structlog import PrintLogger, wrap_logger
wrapped_logger = PrintLogger()
logger = wrap_logger(wrapped_logger, processors=[f1, f2, f3, f4])
log = logger.new(x=42)
```

and call `log.msg('some_event', y=23)`, it results in the following call chain:

```
wrapped_logger.msg(
    f4(wrapped_logger, 'msg',
        f3(wrapped_logger, 'msg',
            f2(wrapped_logger, 'msg',
                f1(wrapped_logger, 'msg', {'event': 'some_event', 'x': 42, 'y': 23})
            )
        )
    )
)
```

In this case, `f4` has to make sure it returns something `wrapped_logger.msg` can handle (see [Adapting and Rendering](#)).

The simplest modification a processor can make is adding new values to the `event_dict`. Parsing human-readable timestamps is tedious, not so **UNIX timestamps** – let's add one to each log entry!

```
import calendar
import time

def timestamper(logger, log_method, event_dict):
    event_dict["timestamp"] = calendar.timegm(time.gmtime())
    return event_dict
```

Please note, that `structlog` comes with such an processor built in: *TimeStamper*.

## Filtering

If a processor raises `structlog.DropEvent`, the event is silently dropped.

Therefore, the following processor drops every entry:

```
from structlog import DropEvent

def dropper(logger, method_name, event_dict):
    raise DropEvent
```

But we can do better than that! How about dropping only log entries that are marked as coming from a certain peer (e.g. monitoring)?

```
from structlog import DropEvent

class ConditionalDropper(object):
    def __init__(self, peer_to_ignore):
        self._peer_to_ignore = peer_to_ignore

    def __call__(self, logger, method_name, event_dict):
        """
        >>> cd = ConditionalDropper('127.0.0.1')
        >>> cd(None, None, {'event': 'foo', 'peer': '10.0.0.1'})
        {'peer': '10.0.0.1', 'event': 'foo'}
        >>> cd(None, None, {'event': 'foo', 'peer': '127.0.0.1'})
        Traceback (most recent call last):
        ...
        DropEvent
        """
        if event_dict.get('peer') == self._peer_to_ignore:
            raise DropEvent
        else:
            return event_dict
```

## Adapting and Rendering

An important role is played by the *last* processor because its duty is to adapt the `event_dict` into something the underlying logging method understands. With that, it's also the *only* processor that needs to know anything about the underlying system.

It can return one of three types:

- A string that is passed as the first (and only) positional argument to the underlying logger.
- A tuple of (`args`, `kwargs`) that are passed as `log_method(*args, **kwargs)`.
- A dictionary which is passed as `log_method(**kwargs)`.

Therefore `return 'hello world'` is a shortcut for `return (('hello world',), {})` (the example in [Chains](#) assumes this shortcut has been taken).

This should give you enough power to use `structlog` with any logging system while writing agnostic processors that operate on dictionaries.

Changed in version 14.0.0: Allow final processor to return a *dict*.

## Examples

The probably most useful formatter for string based loggers is [JSONRenderer](#). Advanced log aggregation and analysis tools like [logstash](#) offer features like telling them “this is JSON, deal with it” instead of fiddling with regular expressions.

More examples can be found in the [examples](#) chapter. For a list of shipped processors, check out the [API documentation](#).

## Third Party Packages

Since processors are self-contained callables, it's easy to write your own and to share them in separate packages. The following processor packages are known to be currently available on PyPI:

- `structlog-prett`: Processors for prettier output – a code syntax highlighter, JSON and XML prettifiers, a multi-line string printer, and a numeric value rounder.

Please feel free to submit a pull request to extend this list with *your* package!

## Examples

This chapter is intended to give you a taste of realistic usage of `structlog`.

### Flask and Thread Local Data

In the simplest case, you bind a unique request ID to every incoming request so you can easily see which log entries belong to which request.

```
import logging
import sys
import uuid

import flask
import structlog

from some_module import some_function

logger = structlog.get_logger()
app = flask.Flask(__name__)

@app.route("/login", methods=["POST", "GET"])
def some_route():
    log = logger.new(
        request_id=str(uuid.uuid4()),
    )
    # do something
    # ...
    log.info("user logged in", user="test-user")
    # gives you:
    # event='user logged in' request_id='ffcdc44f-b952-4b5f-95e6-0f1f3a9ee5fd' user=
    ↪ 'test-user'
    # ...
    some_function()
    # ...
    return "logged in!"

if __name__ == "__main__":
    logging.basicConfig(
        format="% (message) s",
        stream=sys.stdout,
        level=logging.INFO,
    )
    structlog.configure(
```

```
processors=[
    structlog.processors.KeyValueRenderer(
        key_order=["event", "request_id"],
    ),
],
context_class=structlog.threadlocal.wrap_dict(dict),
logger_factory=structlog.stdlib.LoggerFactory(),
)
app.run()
```

some\_module.py

```
from structlog import get_logger

logger = get_logger()

def some_function():
    # later then:
    logger.error('user did something', something='shot_in_foot')
    # gives you:
    # event='user did something' request_id='ffcdc44f-b952-4b5f-95e6-0f1f3a9ee5fd'
    ↪ something='shot_in_foot'
```

While wrapped loggers are *immutable* by default, this example demonstrates how to circumvent that using a thread local dict implementation for context data for convenience (hence the requirement for using *new()* for re-initializing the logger).

Please note that `structlog.stdlib.LoggerFactory` is a totally magic-free class that just deduces the name of the caller's module and does a `logging.getLogger()` with it. It's used by `structlog.get_logger()` to rid you of logging boilerplate in application code. If you prefer to name your standard library loggers explicitly, a positional argument to `get_logger()` gets passed to the factory and used as the name.

## Twisted, and Logging Out Objects

If you prefer to log less but with more context in each entry, you can bind everything important to your logger and log it out with each log entry.

```
import sys
import uuid

import structlog
import twisted

from twisted.internet import protocol, reactor

logger = structlog.getLogger()

class Counter(object):
    i = 0

    def inc(self):
        self.i += 1

    def __repr__(self):
```

```

        return str(self.i)

class Echo(protocol.Protocol):
    def connectionMade(self):
        self._counter = Counter()
        self._log = logger.new(
            connection_id=str(uuid.uuid4()),
            peer=self.transport.getPeer().host,
            count=self._counter,
        )

    def dataReceived(self, data):
        self._counter.inc()
        log = self._log.bind(data=data)
        self.transport.write(data)
        log.msg('echoed data!')

if __name__ == "__main__":
    structlog.configure(
        processors=[structlog.twisted.EventAdapter()],
        logger_factory=structlog.twisted.LoggerFactory(),
    )
    twisted.python.log.startLogging(sys.stderr)
    reactor.listenTCP(1234, protocol.Factory.forProtocol(Echo))
    reactor.run()

```

gives you something like:

```

... peer='127.0.0.1' connection_id='1c6c0cb5-...' count=1 data='123\n' event='echoed_
↳data!'
... peer='127.0.0.1' connection_id='1c6c0cb5-...' count=2 data='456\n' event='echoed_
↳data!'
... peer='127.0.0.1' connection_id='1c6c0cb5-...' count=3 data='foo\n' event='echoed_
↳data!'
... peer='10.10.0.1' connection_id='85234511-...' count=1 data='cba\n' event='echoed_
↳data!'
... peer='127.0.0.1' connection_id='1c6c0cb5-...' count=4 data='bar\n' event='echoed_
↳data!'

```

Since Twisted’s logging system is a bit peculiar, structlog ships with an *adapter* so it keeps behaving like you’d expect it to behave.

I’d also like to point out the Counter class that doesn’t do anything spectacular but gets bound *once* per connection to the logger and since its repr is the number itself, it’s logged out correctly for each event. This shows off the strength of keeping a dict of objects for context instead of passing around serialized strings.

## Processors

*Processors* are a both simple and powerful feature of structlog.

So you want timestamps as part of the structure of the log entry, censor passwords, filter out log entries below your log level before they even get rendered, and get your output as JSON for convenient parsing? Here you go:

```

>>> import datetime, logging, sys
>>> from structlog import wrap_logger

```

```

>>> from structlog.processors import JSONRenderer
>>> from structlog.stdlib import filter_by_level
>>> logging.basicConfig(stream=sys.stdout, format='%(message)s')
>>> def add_timestamp(_, __, event_dict):
...     event_dict['timestamp'] = datetime.datetime.utcnow()
...     return event_dict
>>> def censor_password(_, __, event_dict):
...     pw = event_dict.get('password')
...     if pw:
...         event_dict['password'] = '*CENSORED*'
...     return event_dict
>>> log = wrap_logger(
...     logging.getLogger(__name__),
...     processors=[
...         filter_by_level,
...         add_timestamp,
...         censor_password,
...         JSONRenderer(indent=1, sort_keys=True)
...     ]
... )
>>> log.info('something.filtered')
>>> log.warning('something.not_filtered', password='secret')
{
  "event": "something.not_filtered",
  "password": "*CENSORED*",
  "timestamp": "datetime.datetime(..., ..., ..., ..., ...)"
}

```

structlog comes with many handy processors build right in – for a list of shipped processors, check out the [API documentation](#).

## Development

To make development a more pleasurable experience, structlog comes with the `structlog.dev` module.

The highlight is `structlog.dev.ConsoleRenderer` that offers nicely aligned and colorful console output while in development:

To use it, just add it as a renderer to your processor chain. It will recognize logger names, log levels, time stamps, stack infos, and tracebacks as produced by structlog's processors and render them in special ways.

Suggested configuration for standard library logging:

```

import structlog

structlog.configure(
    processors=[
        structlog.stdlib.add_logger_name,
        structlog.stdlib.add_log_level,
        structlog.stdlib.PositionalArgumentsFormatter(),
        structlog.processors.TimeStamper(fmt="%Y-%m-%d %H:%M.%S"),
        structlog.processors.StackInfoRenderer(),
        structlog.processors.format_exc_info,
        structlog.dev.ConsoleRenderer() # <===
    ],
    context_class=dict,
    logger_factory=structlog.stdlib.LoggerFactory(),

```



```

2015-12-19 19:13.45 [debug    ] debugging is hard      [some_logger] a_list=['1', '2']
2015-12-19 19:13.45 [info     ] informative!        [some_logger] some_key='some value'
2015-12-19 19:13.45 [warning  ] uhuh!              [some_logger]
2015-12-19 19:13.45 [error   ] omg                [some_logger] a_dict={'a': 42, 'b': 'foo'}
2015-12-19 19:13.45 [critical] wtf                [some_logger] what=SomeClass(x=1, y='z')
2015-12-19 19:13.45 [error   ] poor me            [another_logger]

Stack (most recent call last):
  File "t.py", line 43, in <module>
    l2.exception("poor me", stack_info=True)

=====

Traceback (most recent call last):
  File "t.py", line 41, in <module>
    0/0
ZeroDivisionError: division by zero
2015-12-19 19:13.45 [info     ] all better now      [some_logger]
Stack (most recent call last):
  File "t.py", line 44, in <module>
    l.info("all better now", stack_info=True)

```

```

wrapper_class=structlog.stdlib.BindLogger,
cache_logger_on_first_use=True,
)

```

## Integration with Existing Systems

structlog can be used immediately with *any* existing logger. However it comes with special wrappers for the Python standard library and Twisted that are optimized for their respective underlying loggers and contain less magic.

### Standard Library Logging

Ideally, structlog should be able to be used as a drop-in replacement for standard library’s `logging` by wrapping it. In other words, you should be able to replace your call to `logging.getLogger()` by a call to `structlog.get_logger()` and things should keep working as before (if structlog is configured right, see *Suggested Configurations* below).

If you run into incompatibilities, it is a *bug* so please take the time to [report it](#)! If you’re a heavy `logging` user, your [help](#) to ensure a better compatibility would be highly appreciated!

### Concrete Bound Logger

To make structlog’s behavior less magic, it ships with a standard library-specific wrapper class that has an explicit API instead of improvising: `structlog.stdlib.BindLogger`. It behaves exactly like the generic `structlog.BindLogger` except:

- it’s slightly faster due to less overhead,
- has an explicit API that mirrors the log methods of standard library’s `logging.Logger`,
- hence causing less cryptic error messages if you get method names wrong.

## Processors

structlog comes with a few standard library-specific processors:

**`render_to_log_kwargs()`**: Renders the event dictionary into keyword arguments for `logging.log()` that attaches everything except the *event* field to the *extra* argument. This is useful if you want to render your log entries entirely within `logging`.

**`filter_by_level()`**: Checks the log entry's log level against the configuration of standard library's logging. Log entries below the threshold get silently dropped. Put it at the beginning of your processing chain to avoid expensive operations happen in the first place.

**`add_logger_name()`**: Adds the name of the logger to the event dictionary under the key `logger`.

**`add_log_level()`**: Adds the log level to the event dictionary under the key `level`.

**`PositionalArgumentsFormatter`**: This processes and formats positional arguments (if any) passed to log methods in the same way the logging module would do, e.g. `logger.info("Hello, %s", name)`.

structlog also comes with `ProcessorFormatter` which is a `logging.Formatter` that enables you to format non-structlog log entries using structlog renderers *and* multiplex structlog's output with different renderers (see below for an example).

## Suggested Configurations

Depending *where* you'd like to do your formatting, you can take one of three approaches:

### Rendering Using logging-based Formatters

```
import structlog

structlog.configure(
    processors=[
        structlog.stdlib.filter_by_level,
        structlog.stdlib.add_logger_name,
        structlog.stdlib.add_log_level,
        structlog.stdlib.PositionalArgumentsFormatter(),
        structlog.processors.StackInfoRenderer(),
        structlog.processors.format_exc_info,
        structlog.processors.UnicodeDecoder(),
        structlog.stdlib.render_to_log_kwargs,
    ],
    context_class=dict,
    logger_factory=structlog.stdlib.LoggerFactory(),
    wrapper_class=structlog.stdlib.BoundLogger,
    cache_logger_on_first_use=True,
)
```

Now you have the event dict available within each log record. If you want all your log entries (i.e. also those not from your app/structlog) to be formatted as JSON, you can use the [python-json-logger library](#):

```
import logging
import sys

from pythonjsonlogger import jsonlogger

handler = logging.StreamHandler(sys.stdout)
```

```
handler.setFormatter(jsonlogger.JsonFormatter())
root_logger = logging.getLogger()
root_logger.addHandler(handler)
```

Now both structlog and logging will emit JSON logs:

```
>>> structlog.get_logger("test").warning("hello")
{"message": "hello", "logger": "test", "level": "warning"}

>>> logging.getLogger("test").warning("hello")
{"message": "hello"}
```

## Rendering Using structlog-based Formatters Within logging

You can leave rendering for later and yet use structlog's renderers for it:

```
import logging.config
import structlog

timestamp = structlog.processors.TimeStamper(fmt="%Y-%m-%d %H:%M:%S")
pre_chain = [
    # Add the log level and a timestamp to the event_dict if the log entry
    # is not from structlog.
    structlog.stdlib.add_log_level,
    timestamp,
]

logging.config.dictConfig({
    "version": 1,
    "disable_existing_loggers": False,
    "formatters": {
        "plain": {
            "()": structlog.stdlib.ProcessorFormatter,
            "processor": structlog.dev.ConsoleRenderer(colors=False),
            "foreign_pre_chain": pre_chain,
        },
        "colored": {
            "()": structlog.stdlib.ProcessorFormatter,
            "processor": structlog.dev.ConsoleRenderer(colors=True),
            "foreign_pre_chain": pre_chain,
        },
    },
    "handlers": {
        "default": {
            "level": "DEBUG",
            "class": "logging.StreamHandler",
            "formatter": "colored",
        },
        "file": {
            "level": "DEBUG",
            "class": "logging.handlers.WatchedFileHandler",
            "filename": "test.log",
            "formatter": "plain",
        },
    },
    "loggers": {
```

```

        "": {
            "handlers": ["default", "file"],
            "level": "DEBUG",
            "propagate": True,
        },
    },
}
))
structlog.configure(
    processors=[
        structlog.stdlib.add_log_level,
        structlog.stdlib.PositionalArgumentsFormatter(),
        structlog.processors.StackInfoRenderer(),
        structlog.processors.format_exc_info,
        structlog.stdlib.ProcessorFormatter.wrap_for_formatter,
    ],
    context_class=dict,
    logger_factory=structlog.stdlib.LoggerFactory(),
    wrapper_class=structlog.stdlib.BoundLogger,
    cache_logger_on_first_use=True,
)

```

This defines two formatters: one plain and one colored. Both are run for each log entry. Log entries that do not originate from structlog, are additionally pre-processed using a cached `timestamp` and `add_log_level()`.

```

>>> logging.getLogger().warning("bar")
2017-03-06 11:49:27 [warning ] bar

>>> structlog.get_logger("structlog").warning("foo", x=42)
2017-03-06 11:49:32 [warning ] foo                                x=42

>>> print(open("test.log").read())
2017-03-06 11:49:27 [warning ] bar
2017-03-06 11:49:32 [warning ] foo                                x=42

```

(sadly, you have to imagine the colors in the first two outputs)

If you leave `foreign_pre_chain` *None*, formatting will be left to `logging`. Meaning: you can define a format for `ProcessorFormatter` too!

## Rendering Within structlog

A basic configuration to output structured logs in JSON format looks like this:

```

import structlog

structlog.configure(
    processors=[
        structlog.stdlib.filter_by_level,
        structlog.stdlib.add_logger_name,
        structlog.stdlib.add_log_level,
        structlog.stdlib.PositionalArgumentsFormatter(),
        structlog.processors.TimeStamper(fmt="iso"),
        structlog.processors.StackInfoRenderer(),
        structlog.processors.format_exc_info,
        structlog.processors.UnicodeDecoder(),
        structlog.processors.JSONRenderer()
    ]
)

```

```

    ],
    context_class=dict,
    logger_factory=structlog.stdlib.LoggerFactory(),
    wrapper_class=structlog.stdlib.BoundLogger,
    cache_logger_on_first_use=True,
)

```

(if you're still running Python 2, replace `UnicodeDecoder` through `UnicodeEncoder`)

To make your program behave like a proper 12 factor app that outputs only JSON to stdout, configure the logging module like this:

```

import logging
import sys

logging.basicConfig(
    format="% (message) s",
    stream=sys.stdout,
    level=logging.INFO,
)

```

In this case *only* your own logs are formatted as JSON:

```

>>> structlog.get_logger("test").warning("hello")
{"event": "hello", "logger": "test", "level": "warning", "timestamp": "2017-03-
↪06T07:39:09.518720Z"}

>>> logging.getLogger("test").warning("hello")
hello

```

## Twisted

**Warning:** Since `sys.exc_clear()` has been dropped in Python 3, there is currently no way to avoid multiple tracebacks in your log files if using structlog together with Twisted on Python 3.

### Concrete Bound Logger

To make structlog's behavior less magic, it ships with a Twisted-specific wrapper class that has an explicit API instead of improvising: `structlog.twisted.BoundLogger`. It behaves exactly like the generic `structlog.BoundLogger` except:

- it's slightly faster due to less overhead,
- has an explicit API (`msg()` and `err()`),
- hence causing less cryptic error messages if you get method names wrong.

In order to avoid that structlog disturbs your CamelCase harmony, it comes with an alias for `structlog.get_logger()` called `structlog.getLogger()`.

## Processors

structlog comes with two Twisted-specific processors:

**EventAdapter** This is useful if you have an existing Twisted application and just want to wrap your loggers for now. It takes care of transforming your event dictionary into something `twisted.python.log.err` can digest.

For example:

```
def onError(fail):
    failure = fail.trap(MoonExploded)
    log.err(failure, _why='event-that-happend')
```

will still work as expected.

Needs to be put at the end of the processing chain. It formats the event using a renderer that needs to be passed into the constructor:

```
configure(processors=[EventAdapter(KeyValueRenderer())])
```

The drawback of this approach is that Twisted will format your exceptions as multi-line log entries which is painful to parse. Therefore `structlog` comes with:

**JSONRenderer** Goes a step further and circumvents Twisted logger's Exception/Failure handling and renders it itself as JSON strings. That gives you regular and simple-to-parse single-line JSON log entries no matter what happens.

## Bending Foreign Logging To Your Will

`structlog` comes with a wrapper for Twisted's log observers to ensure the rest of your logs are in JSON too: `JSONLogObserverWrapper()`.

What it does is determining whether a log entry has been formatted by `JSONRenderer` and if not, converts the log entry to JSON with `event` being the log message and putting Twisted's `system` into a second key.

So for example:

```
2013-09-15 22:02:18+0200 [-] Log opened.
```

becomes:

```
2013-09-15 22:02:18+0200 [-] {"event": "Log opened.", "system": "-"}
```

There is obviously some redundancy here. Also, I'm presuming that if you write out JSON logs, you're going to let something else parse them which makes the human-readable date entries more trouble than they're worth.

To get a clean log without timestamps and additional system fields (`[-]`), `structlog` comes with `PlainFileLogObserver` that writes only the plain message to a file and `plainJSONStdOutLogger()` that composes it with the aforementioned `JSONLogObserverWrapper()` and gives you a pure JSON log without any timestamps or other noise straight to `standard out`:

```
$ twistd -n --logger structlog.twisted.plainJSONStdOutLogger web
{"event": "Log opened.", "system": "-"}
{"event": "twistd 13.1.0 (python 2.7.3) starting up.", "system": "-"}
{"event": "reactor class: twisted...EPollReactor.", "system": "-"}
{"event": "Site starting on 8080", "system": "-"}
{"event": "Starting factory <twisted.web.server.Site ...>", ...}
...
```

## Suggested Configuration

```
import structlog

structlog.configure(
    processors=[
        structlog.processors.StackInfoRenderer(),
        structlog.twisted.JSONRenderer()
    ],
    context_class=dict,
    logger_factory=structlog.twisted.LoggerFactory(),
    wrapper_class=structlog.twisted.BoundLogger,
    cache_logger_on_first_use=True,
)
```

See also *Logging Best Practices*.

## Logging Best Practices

The best practice for you depends very much on your context. To give you some pointers nevertheless, here are a few scenarios that may be applicable to you.

Pull requests for further interesting approaches as well as refinements and more complete examples are very welcome.

### Common Ideas

Logging is not a new concept and in no way special to Python. Logfiles have existed for decades and there's little reason to reinvent the wheel in our little world.

There are several concepts that are very well-solved in general and especially in heterogeneous environments, using special tooling for Python applications does more harm than good and makes the operations staff build dart board with your pictures.

Therefore let's rely on proven tools as much as possible and do only the absolutely necessary inside of Python<sup>0</sup>. A very nice approach is to simply log to [standard out](#) and let other tools take care of the rest.

### runit

One runner that makes this very easy is the venerable [runit](#) project which made it a part of its design: server processes don't detach but log to standard out instead. There it gets processed by other software – potentially by one of its own tools: [svlogd](#). We use it extensively and it has proven itself extremely robust and capable; check out [this tutorial](#) if you'd like to try it.

If you're not quite convinced and want an overview on running daemons, have a look at [cue's daemon showdown](#) that discusses the most common ones.

### Local Logging

There are basically two common ways to log to local logfiles: writing yourself into files and syslog.

<sup>0</sup> This is obviously a privileged UNIX-centric view but even Windows has tools and means for log management although we won't be able to discuss them here.

### Syslog

The simplest approach to logging is to forward your entries to the `syslogd`. Twisted, uwsgi, and runit support it directly. It will happily add a timestamp and write wherever you tell it in its configuration. You can also log from multiple processes into a single file and use your system's `logrotate` for log rotation.

The only downside is that syslog has some quirks that show itself under high load like rate limits (they can be switched off) and lost log entries.

### runit's svlogd

If you'll choose runit for running your daemons, `svlogd` is a nicer approach. It receives the log entries via a UNIX pipe and acts on them which includes adding of parse-friendly timestamps in `tai64n` as well as filtering and log rotation.

### Centralized Logging

Nowadays you usually don't want your logfiles in compressed archives distributed over dozens – if not thousands – servers. You want them at a single location; parsed and easy to query.

### Syslog (Again!)

The widely deployed syslog implementation `rsyslog` supports remote logging out-of-the-box. Have a look at [this post](#) by Revolution Systems on the how.

Since syslog is such a widespread solution, there are also ways to use it with basically any centralized product.

### Logstash with logstash-forwarder

`Logstash` is a great way to parse, save, and search your logs.

The general modus operandi is that you have log shippers that parse your log files and forward the log entries to your Logstash server and store is in `elasticsearch`. If your log entries consist of a JSON dictionary (and perhaps a `tai64n` timestamp), this is pretty easy and efficient.

If you can't decide on a log shipper, `logstash-forwarder` (formerly known as Lumberjack) works really well. When Logstash's `lumberjack` input is configured to use `codec => "json"`, having `structlog` output JSON is all you need. See the documentation on the [Standard Library Logging](#) for an example configuration.

### Graylog2

`Graylog` goes one step further. It not only supports everything those above do (and then some); you can also log directly JSON entries towards it – optionally even through an AMQP server (like `RabbitMQ`) for better reliability. Additionally, `Graylog's Extended Log Format (GELF)` allows for structured data which makes it an obvious choice to use together with `structlog`.



## Advanced Topics

### Custom Wrappers

structlog comes with a generic bound logger called `structlog.BoundLogger` that can be used to wrap any logger class you fancy. It does so by intercepting unknown method names and proxying them to the wrapped logger.

This works fine, except that it has a performance penalty and the API of `BoundLogger` isn't clear from reading the documentation because large parts depend on the wrapped logger. An additional reason is that you may want to have semantically meaningful log method names that add meta data to log entries as it is fit (see example below).

To solve that, structlog offers you to use an own wrapper class which you can configure using `structlog.configure()`. And to make it easier for you, it comes with the class `structlog.BoundLoggerBase` which takes care of all data binding duties so you just add your log methods if you choose to sub-class it.

### Example

It's much easier to demonstrate with an example:

```
>>> from structlog import BoundLoggerBase, PrintLogger, wrap_logger
>>> class SemanticLogger(BoundLoggerBase):
...     def msg(self, event, **kw):
...         if not 'status' in kw:
...             return self._proxy_to_logger('msg', event, status='ok', **kw)
...         else:
...             return self._proxy_to_logger('msg', event, **kw)
...
...     def user_error(self, event, **kw):
...         self.msg(event, status='user_error', **kw)
>>> log = wrap_logger(PrintLogger(), wrapper_class=SemanticLogger)
>>> log = log.bind(user='fprefect')
>>> log.user_error('user.forgot_towel')
user='fprefect' status='user_error' event='user.forgot_towel'
```

You can observe the following:

- The wrapped logger can be found in the instance variable `structlog.BoundLoggerBase._logger`.
- The helper method `structlog.BoundLoggerBase._proxy_to_logger()` that is a DRY convenience function that runs the processor chain, handles possible `DropEvents` and calls a named function on `_logger`.
- You can run the chain by hand though using `structlog.BoundLoggerBase._process_event()`.

These two methods and one attribute is all you need to write own wrapper classes.

### Performance

structlog's default configuration tries to be as unsurprising and not confusing to new developers as possible. Some of the choices made come with an avoidable performance price tag – although its impact is debatable.

Here are a few hints how to get most out of structlog in production:

1. Use plain *dicts* as context classes. Python is full of them and they are highly optimized:

```
configure(context_class=dict)
```

If you don't use automated parsing (you should!) and need predictable order of your keys for some reason, use the `key_order` argument of `KeyValueRenderer`.

2. Use a specific wrapper class instead of the generic one. `structlog` comes with ones for the *Standard Library Logging* and for *Twisted*:

```
configure(wrapper_class=structlog.stdlib.BoundLogger)
```

*Writing own wrapper classes* is straightforward too.

3. Avoid (frequently) calling log methods on loggers you get back from `structlog.wrap_logger()` and `structlog.get_logger()`. Since those functions are usually called in module scope and thus before you are able to configure them, they return a proxy that assembles the correct logger on demand.

Create a local logger if you expect to log frequently without binding:

```
logger = structlog.get_logger()
def f():
    log = logger.bind()
    for i in range(1000000000):
        log.info('iterated', i=i)
```

4. Set the `cache_logger_on_first_use` option to `True` so the aforementioned on-demand loggers will be assembled only once and cached for future uses:

```
configure(cache_logger_on_first_use=True)
```

This has the only drawback is that later calls on `configure()` don't have any effect on already cached loggers – that shouldn't matter outside of testing though.

5. Use a faster JSON serializer than the standard library. Possible alternatives are among others `simplejson`, `UltraJSON`, or `RapidJSON` (Python 3 only):

```
structlog.processors.JSONRenderer(serializer=rapidjson.dumps)
```

## API Reference

---

**Note:** The examples here use a very simplified configuration using the minimalistic `structlog.processors.KeyValueRenderer` for brevity and to enable doctests. The output is going to be different (nicer!) with default configuration.

---

### structlog Package

`structlog.get_logger(*args, **initial_values)`

Convenience function that returns a logger according to configuration.

```
>>> from structlog import get_logger
>>> log = get_logger(y=23)
>>> log.msg("hello", x=42)
y=23 x=42 event='hello'
```

#### Parameters

- **args** – *Optional* positional arguments that are passed unmodified to the logger factory. Therefore it depends on the factory what they mean.
- **initial\_values** – Values that are used to pre-populate your contexts.

**Return type** A proxy that creates a correctly configured bound logger when necessary.

See [Configuration](#) for details.

If you prefer CamelCase, there's an alias for your reading pleasure: `structlog.get_logger()`.

New in version 0.4.0: `args`

`structlog.getLogger(*args, **initial_values)`  
CamelCase alias for `structlog.get_logger()`.

This function is supposed to be in every source file – we don’t want it to stick out like a sore thumb in frameworks like Twisted or Zope.

`structlog.wrap_logger(logger, processors=None, wrapper_class=None, context_class=None, cache_logger_on_first_use=None, logger_factory_args=None, **initial_values)`

Create a new bound logger for an arbitrary `logger`.

Default values for `processors`, `wrapper_class`, and `context_class` can be set using `configure()`.

If you set an attribute here, `configure()` calls have *no* effect for the *respective* attribute.

In other words: selective overwriting of the defaults while keeping some *is* possible.

#### Parameters

- **initial\_values** – Values that are used to pre-populate your contexts.
- **logger\_factory\_args** (*tuple*) – Values that are passed unmodified as `*logger_factory_args` to the logger factory if not `None`.

**Return type** A proxy that creates a correctly configured bound logger when necessary.

See `configure()` for the meaning of the rest of the arguments.

New in version 0.4.0: `logger_factory_args`

`structlog.configure(processors=None, wrapper_class=None, context_class=None, logger_factory=None, cache_logger_on_first_use=None)`

Configures the **global** defaults.

They are used if `wrap_logger()` has been called without arguments.

Also sets the global class attribute `is_configured` to `True` on first call. Can be called several times, keeping an argument at `None` leaves `is` unchanged from the current setting.

Use `reset_defaults()` to undo your changes.

#### Parameters

- **processors** (*list*) – List of processors.
- **wrapper\_class** (*type*) – Class to use for wrapping loggers instead of `structlog.BoundLogger`. See *Standard Library Logging*, *Twisted*, and *Custom Wrappers*.
- **context\_class** (*type*) – Class to be used for internal context keeping.
- **logger\_factory** (*callable*) – Factory to be called to create a new logger that shall be wrapped.
- **cache\_logger\_on\_first\_use** (*bool*) – `wrap_logger` doesn’t return an actual wrapped logger but a proxy that assembles one when it’s first used. If this option is set to `True`, this assembled logger is cached. See *Performance*.

New in version 0.3.0: `cache_logger_on_first_use`

`structlog.configure_once(*args, **kw)`  
Configures iff structlog isn’t configured yet.

It does *not* matter whether `is` was configured using `configure()` or `configure_once()` before.

Raises a `RuntimeWarning` if repeated configuration is attempted.

`structlog.reset_defaults()`

Resets global default values to builtins.

That means [`StackInfoRenderer`, `format_exc_info()`, `TimeStamper`, `ConsoleRenderer`] for `processors`, `BoundLogger` for `wrapper_class`, `OrderedDict` for `context_class`, `PrintLoggerFactory` for `logger_factory`, and `False` for `cache_logger_on_first_use`.

Also sets the global class attribute `is_configured` to `False`.

**class** `structlog.BoundLogger` (*logger*, *processors*, *context*)

A generic BoundLogger that can wrap anything.

Every unknown method will be passed to the wrapped logger. If that's too much magic for you, try `structlog.stdlib.BoundLogger` or `structlog.twisted.BoundLogger` which also take advantage of knowing the wrapped class which generally results in better performance.

Not intended to be instantiated by yourself. See `wrap_logger()` and `get_logger()`.

**bind** (*\*\*new\_values*)

Return a new logger with *new\_values* added to the existing ones.

**Return type** `self.__class__`

**new** (*\*\*new\_values*)

Clear context and binds *initial\_values* using `bind()`.

Only necessary with dict implementations that keep global state like those wrapped by `structlog.threadlocal.wrap_dict()` when threads are re-used.

**Return type** `self.__class__`

**unbind** (*\*keys*)

Return a new logger with *keys* removed from the context.

**Raises** `KeyError` – If the key is not part of the context.

**Return type** `self.__class__`

**class** `structlog.PrintLogger` (*file=None*)

Print events into a file.

**Parameters** **file** (*file*) – File to print to. (default: stdout)

```
>>> from structlog import PrintLogger
>>> PrintLogger().msg('hello')
hello
```

Useful if you just capture your stdout with tools like `runit` or if you forward your stderr to syslog.

Also very useful for testing and examples since logging is finicky in doctests.

**critical** (*message*)

Print *message*.

**debug** (*message*)

Print *message*.

**err** (*message*)

Print *message*.

**error** (*message*)

Print *message*.

**failure** (*message*)

Print *message*.

**info** (*message*)  
Print *message*.

**log** (*message*)  
Print *message*.

**msg** (*message*)  
Print *message*.

**warning** (*message*)  
Print *message*.

**class** structlog.**PrintLoggerFactory** (*file=None*)  
Produce *PrintLoggers*.

To be used with *structlog.configure()* 's *logger\_factory*.

**Parameters** **file** (*file*) – File to print to. (default: stdout)

Positional arguments are silently ignored.

New in version 0.4.0.

**class** structlog.**ReturnLogger**  
Return the arguments that it's called with.

```
>>> from structlog import ReturnLogger
>>> ReturnLogger().msg('hello')
'hello'
>>> ReturnLogger().msg('hello', when='again')
(('hello',), {'when': 'again'})
```

Useful for unit tests.

Changed in version 0.3.0: Allow for arbitrary arguments and keyword arguments to be passed in.

**critical** (*\*args, \*\*kw*)  
Return tuple of *args*, *kw* or just *args*[0] if only one arg passed

**debug** (*\*args, \*\*kw*)  
Return tuple of *args*, *kw* or just *args*[0] if only one arg passed

**err** (*\*args, \*\*kw*)  
Return tuple of *args*, *kw* or just *args*[0] if only one arg passed

**error** (*\*args, \*\*kw*)  
Return tuple of *args*, *kw* or just *args*[0] if only one arg passed

**failure** (*\*args, \*\*kw*)  
Return tuple of *args*, *kw* or just *args*[0] if only one arg passed

**info** (*\*args, \*\*kw*)  
Return tuple of *args*, *kw* or just *args*[0] if only one arg passed

**log** (*\*args, \*\*kw*)  
Return tuple of *args*, *kw* or just *args*[0] if only one arg passed

**msg** (*\*args, \*\*kw*)  
Return tuple of *args*, *kw* or just *args*[0] if only one arg passed

**warning** (*\*args, \*\*kw*)  
Return tuple of *args*, *kw* or just *args*[0] if only one arg passed

**class** `structlog.ReturnLoggerFactory`

Produce and cache *ReturnLoggers*.

To be used with `structlog.configure()`'s *logger\_factory*.

Positional arguments are silently ignored.

New in version 0.4.0.

**exception** `structlog.DropEvent`

If raised by an processor, the event gets silently dropped.

Derives from `BaseException` because it's technically not an error.

**class** `structlog.BoundLoggerBase` (*logger, processors, context*)

Immutable context carrier.

Doesn't do any actual logging; examples for useful subclasses are:

- the generic *BoundLogger* that can wrap anything,
- `structlog.twisted.BoundLogger`,
- and `structlog.stdlib.BoundLogger`.

See also *Custom Wrappers*.

**`__logger = None`**

Wrapped logger.

---

**Note:** Despite underscore available **read-only** to custom wrapper classes.

See also *Custom Wrappers*.

---

**`__process_event`** (*method\_name, event, event\_kw*)

Combines creates an *event\_dict* and runs the chain.

Call it to combine your *event* and *context* into an *event\_dict* and process using the processor chain.

#### Parameters

- **`method_name`** (*str*) – The name of the logger method. Is passed into the processors.
- **`event`** – The event – usually the first positional argument to a logger.
- **`event_kw`** – Additional event keywords. For example if someone calls `log.msg('foo', bar=42)`, *event* would be `'foo'` and *event\_kw* `{'bar': 42}`.

**Raises** *structlog.DropEvent* if log entry should be dropped.

**Raises** *ValueError* if the final processor doesn't return a string, tuple, or a dict.

**Return type** *tuple* of (*\*args, \*\*kw*)

---

**Note:** Despite underscore available to custom wrapper classes.

See also *Custom Wrappers*.

---

Changed in version 14.0.0: Allow final processor to return a *dict*.

**`__proxy_to_logger`** (*method\_name, event=None, \*\*event\_kw*)

Run processor chain on event & call *method\_name* on wrapped logger.

DRY convenience method that runs `_process_event()`, takes care of handling `structlog.DropEvent`, and finally calls `method_name` on `_logger` with the result.

#### Parameters

- **method\_name** (*str*) – The name of the method that’s going to get called. Technically it should be identical to the method the user called because it also get passed into processors.
- **event** – The event – usually the first positional argument to a logger.
- **event\_kw** – Additional event keywords. For example if someone calls `log.msg('foo', bar=42)`, `event` would be `'foo'` and `event_kw` `{'bar': 42}`.

---

**Note:** Despite underscore available to custom wrapper classes.

See also *Custom Wrappers*.

---

**bind** (*\*\*new\_values*)

Return a new logger with *new\_values* added to the existing ones.

**Return type** `self.__class__`

**new** (*\*\*new\_values*)

Clear context and binds *initial\_values* using `bind()`.

Only necessary with dict implementations that keep global state like those wrapped by `structlog.threadlocal.wrap_dict()` when threads are re-used.

**Return type** `self.__class__`

**unbind** (*\*keys*)

Return a new logger with *keys* removed from the context.

**Raises** **KeyError** – If the key is not part of the context.

**Return type** `self.__class__`

## dev Module

Helpers that make development with structlog more pleasant.

**class** `structlog.dev.ConsoleRenderer` (*pad\_event=30, colors=True, repr\_native\_str=False*)

Render *event\_dict* nicely aligned, possibly in colors, and ordered.

#### Parameters

- **pad\_event** (*int*) – Pad the event to this many characters.
- **colors** (*bool*) – Use colors for a nicer output.
- **repr\_native\_str** (*bool*) – When `True`, `repr()` is also applied to native strings (i.e. unicode on Python 3 and bytes on Python 2). Setting this to `False` is useful if you want to have human-readable non-ASCII output on Python 2. The *event* key is *never* `repr()`-ed.

Requires the `colorama` package if *colors* is `True`.

New in version 16.0.

New in version 16.1: *colors*

New in version 17.1: *repr\_native\_str*



## threadlocal Module

Primitives to keep context global but thread (and greenlet) local.

`structlog.threadlocal.wrap_dict(dict_class)`

Wrap a dict-like class and return the resulting class.

The wrapped class and used to keep global in the current thread.

**Parameters** `dict_class` (*type*) – Class used for keeping context.

**Return type** *type*

`structlog.threadlocal.tmp_bind(logger, **tmp_values)`

Bind *tmp\_values* to *logger* & memorize current state. Rewind afterwards.

```
>>> from structlog import wrap_logger, PrintLogger
>>> from structlog.threadlocal import tmp_bind, wrap_dict
>>> logger = wrap_logger(PrintLogger(), context_class=wrap_dict(dict))
>>> with tmp_bind(logger, x=5) as tmp_logger:
...     logger = logger.bind(y=3)
...     tmp_logger.msg("event")
x=5 y=3 event='event'
>>> logger.msg("event")
event='event'
```

`structlog.threadlocal.as_immutable(logger)`

Extract the context from a thread local logger into an immutable logger.

**Parameters** `logger` (`structlog.BoundLogger`) – A logger with *possibly* thread local state.

**Return type** `BoundLogger` with an immutable context.

## processors Module

Processors useful regardless of the logging framework.

**class** `structlog.processors.JSONRenderer(serializer=<function dumps>, **dumps_kw)`

Render the *event\_dict* using *serializer* (*event\_dict*, *\*\*json\_kw*).

**Parameters**

- `json_kw` (*dict*) – Are passed unmodified to *serializer*.
- `serializer` (*callable*) – A `json.dumps()`-compatible callable that will be used to format the string. This can be used to use alternative JSON encoders like `simplejson` or `RapidJSON` (faster but Python 3-only) (default: `json.dumps()`).

New in version 0.2.0: Support for `__structlog__` serialization method.

New in version 15.4.0: `serializer` parameter.

```
>>> from structlog.processors import JSONRenderer
>>> JSONRenderer(sort_keys=True)(None, None, {'a': 42, 'b': [1, 2, 3]})
'{"a": 42, "b": [1, 2, 3]}'
```

Bound objects are attempted to be serialize using a `__structlog__` method. If none is defined, `repr()` is used:

```
>>> class C1(object):
...     def __structlog__(self):
...         return ['C1!']
...     def __repr__(self):
...         return '__structlog__ took precedence'
>>> class C2(object):
...     def __repr__(self):
...         return 'No __structlog__, so this is used.'
>>> from structlog.processors import JSONRenderer
>>> JSONRenderer(sort_keys=True)(None, None, {'c1': C1(), 'c2': C2()})
'{"c1": ["C1!"], "c2": "No __structlog__, so this is used."}'
```

Please note that additionally to strings, you can also return any type the standard library JSON module knows about – like in this example a list.

**class** structlog.processors.**KeyValueRenderer** (*sort\_keys=False*, *key\_order=None*,  
*drop\_missing=False*, *repr\_native\_str=True*)  
 Render *event\_dict* as a list of *Key=repr(Value)* pairs.

#### Parameters

- **sort\_keys** (*bool*) – Whether to sort keys when formatting.
- **key\_order** (*list*) – List of keys that should be rendered in this exact order. Missing keys will be rendered as *None*, extra keys depending on *sort\_keys* and the dict class.
- **drop\_missing** (*bool*) – When *True*, extra keys in *key\_order* will be dropped rather than rendered as *None*.
- **repr\_native\_str** (*bool*) – When *True*, *repr()* is also applied to native strings (i.e. unicode on Python 3 and bytes on Python 2). Setting this to *False* is useful if you want to have human-readable non-ASCII output on Python 2.

New in version 0.2.0: *key\_order*

New in version 16.1.0: *drop\_missing*

New in version 17.1.0: *repr\_native\_str*

```
>>> from structlog.processors import KeyValueRenderer
>>> KeyValueRenderer(sort_keys=True)(None, None, {'a': 42, 'b': [1, 2, 3]})
'a=42 b=[1, 2, 3]'
>>> KeyValueRenderer(key_order=['b', 'a'])(None, None,
...                                         {'a': 42, 'b': [1, 2, 3]})
'b=[1, 2, 3] a=42'
```

**class** structlog.processors.**UnicodeDecoder** (*encoding='utf-8'*, *errors='replace'*)  
 Decode byte string values in *event\_dict*.

#### Parameters

- **encoding** (*str*) – Encoding to decode from (default: *'utf-8'*).
- **errors** (*str*) – How to cope with encoding errors (default: *'replace'*).

Useful if you're running Python 3 as otherwise *b"abc"* will be rendered as *'b"abc"'*.

Just put it in the processor chain before the renderer.

New in version 15.4.0.

**class** structlog.processors.**UnicodeEncoder** (*encoding='utf-8'*, *errors='backslashreplace'*)  
 Encode unicode values in *event\_dict*.

**Parameters**

- **encoding** (*str*) – Encoding to encode to (default: 'utf-8').
- **errors** (*str*) – How to cope with encoding errors (default 'backslashreplace').

Useful if you're running Python 2 as otherwise `u"abc"` will be rendered as `'u"abc"'`.

Just put it in the processor chain before the renderer.

`structlog.processors.format_exc_info(logger, name, event_dict)`

Replace an `exc_info` field by an `exception` string field:

If `event_dict` contains the key `exc_info`, there are two possible behaviors:

- If the value is a tuple, render it into the key `exception`.
- If the value is an Exception *and* you're running Python 3, render it into the key `exception`.
- If the value true but no tuple, obtain `exc_info` ourselves and render that.

If there is no `exc_info` key, the `event_dict` is not touched. This behavior is analogue to the one of the `stdlib's` logging.

```
>>> from structlog.processors import format_exc_info
>>> try:
...     raise ValueError
... except ValueError:
...     format_exc_info(None, None, {'exc_info': True})
{'exception': 'Traceback (most recent call last):...
```

**class** `structlog.processors.StackInfoRenderer`

Add stack information with key `stack` if `stack_info` is true.

Useful when you want to attach a stack dump to a log entry without involving an exception.

It works analogously to the `stack_info` argument of the Python 3 standard library logging but works on both 2 and 3.

New in version 0.4.0.

**class** `structlog.processors.ExceptionPrettyPrinter` (*file=None*)

Pretty print exceptions and remove them from the `event_dict`.

**Parameters** **file** (*file*) – Target file for output (default: `sys.stdout`).

This processor is mostly for development and testing so you can read exceptions properly formatted.

It behaves like `format_exc_info()` except it removes the exception data from the event dictionary after printing it.

It's tolerant to having `format_exc_info` in front of itself in the processor chain but doesn't require it. In other words, it handles both `exception` as well as `exc_info` keys.

New in version 0.4.0.

Changed in version 16.0.0: Added support for passing exceptions as `exc_info` on Python 3.

**class** `structlog.processors.TimeStamper` (*fnt=None, utc=True*)

Add a timestamp to `event_dict`.

---

**Note:** You should let OS tools take care of timestamping. See also [Logging Best Practices](#).

---

**Parameters**

- **fmt** (*str*) – strftime format string, or "iso" for ISO 8601, or *None* for a UNIX timestamp.
- **utc** (*bool*) – Whether timestamp should be in UTC or local time.
- **key** (*str*) – Target key in *event\_dict* for added timestamps.

```
>>> from structlog.processors import TimeStamper
>>> TimeStamper() (None, None, {})
{'timestamp': 1378994017}
>>> TimeStamper(fmt='iso') (None, None, {})
{'timestamp': '2013-09-12T13:54:26.996778Z'}
>>> TimeStamper(fmt='%Y', key='year') (None, None, {})
{'year': '2013'}
```

## stdlib Module

Processors and helpers specific to the `logging` module from the Python standard library.

See also *structlog's standard library support*.

**class** `structlog.stdlib.BoundLogger` (*logger, processors, context*)

Python Standard Library version of `structlog.BoundLogger`. Works exactly like the generic one except that it takes advantage of knowing the logging methods in advance.

Use it like:

```
structlog.configure(
    wrapper_class=structlog.stdlib.BoundLogger,
)
```

**bind** (*\*\*new\_values*)

Return a new logger with *new\_values* added to the existing ones.

**Return type** *self.\_\_class\_\_*

**critical** (*event=None, \*args, \*\*kw*)

Process event and call `logging.Logger.critical()` with the result.

**debug** (*event=None, \*args, \*\*kw*)

Process event and call `logging.Logger.debug()` with the result.

**error** (*event=None, \*args, \*\*kw*)

Process event and call `logging.Logger.error()` with the result.

**exception** (*event=None, \*args, \*\*kw*)

Process event and call `logging.Logger.error()` with the result, after setting `exc_info` to *True*.

**info** (*event=None, \*args, \*\*kw*)

Process event and call `logging.Logger.info()` with the result.

**log** (*level, event, \*args, \*\*kw*)

Process event and call the appropriate logging method depending on *level*.

**new** (*\*\*new\_values*)

Clear context and binds *initial\_values* using `bind()`.

Only necessary with dict implementations that keep global state like those wrapped by `structlog.threadlocal.wrap_dict()` when threads are re-used.

**Return type** *self.\_\_class\_\_*

**unbind**(*\*keys*)

Return a new logger with *keys* removed from the context.

**Raises** `KeyError` – If the key is not part of the context.

**Return type** `self.__class__`

**warn**(*event=None, \*args, \*\*kw*)

Process event and call `logging.Logger.warning()` with the result.

**warning**(*event=None, \*args, \*\*kw*)

Process event and call `logging.Logger.warning()` with the result.

**class** `structlog.stdlib.LoggerFactory`(*ignore\_frame\_names=None*)

Build a standard library logger when an *instance* is called.

Sets a custom logger using `logging.setLoggerClass()` so variables in log format are expanded properly.

```
>>> from structlog import configure
>>> from structlog.stdlib import LoggerFactory
>>> configure(logger_factory=LoggerFactory())
```

**Parameters** `ignore_frame_names` (*list of str*) – When guessing the name of a logger, skip frames whose names *start* with one of these. For example, in pyramid applications you’ll want to set it to `['venusian', 'pyramid.config']`.

**\_\_call\_\_**(*\*args*)

Deduce the caller’s module name and create a stdlib logger.

If an optional argument is passed, it will be used as the logger name instead of guesswork. This optional argument would be passed from the `structlog.get_logger()` call. For example `structlog.get_logger('foo')` would cause this method to be called with `'foo'` as its first positional argument.

**Return type** `logging.Logger`

Changed in version 0.4.0: Added support for optional positional arguments. Using the first one for naming the constructed logger.

`structlog.stdlib.render_to_log_kwargs`(*wrapped\_logger, method\_name, event\_dict*)

Render *event\_dict* into keyword arguments for `logging.log()`.

The *event* field is translated into *msg* and the rest of the *event\_dict* is added as *extra*.

This allows you to defer formatting to `logging`.

New in version 17.1.0.

`structlog.stdlib.filter_by_level`(*logger, name, event\_dict*)

Check whether logging is configured to accept messages from this log level.

Should be the first processor if stdlib’s filtering by level is used so possibly expensive processors like exception formatters are avoided in the first place.

```
>>> import logging
>>> from structlog.stdlib import filter_by_level
>>> logging.basicConfig(level=logging.WARN)
>>> logger = logging.getLogger()
>>> filter_by_level(logger, 'warn', {})
{}
>>> filter_by_level(logger, 'debug', {})
Traceback (most recent call last):
```

```
...
DropEvent
```

`structlog.stdlib.add_log_level(logger, method_name, event_dict)`

Add the log level to the event dict.

`structlog.stdlib.add_logger_name(logger, method_name, event_dict)`

Add the logger name to the event dict.

**class** `structlog.stdlib.PositionalArgumentsFormatter(remove_positional_args=True)`

Apply stdlib-like string formatting to the *event* key.

If the *positional\_args* key in the event dict is set, it must contain a tuple that is used for formatting (using the *%s* string formatting operator) of the value from the *event* key. This works in the same way as the stdlib handles arguments to the various log methods: if the tuple contains only a single *dict* argument it is used for keyword placeholders in the *event* string, otherwise it will be used for positional placeholders.

*positional\_args* is populated by `structlog.stdlib.BoundLogger` or can be set manually.

The *remove\_positional\_args* flag can be set to *False* to keep the *positional\_args* key in the event dict; by default it will be removed from the event dict after formatting a message.

**class** `structlog.stdlib.ProcessorFormatter(processor, foreign_pre_chain=None, keep_exc_info=False, keep_stack_info=False, *args, **kwargs)`

Call `structlog` processors on `logging.LogRecords`.

This `logging.Formatter` allows to configure `logging` to call *processor* on structlog-borne log entries (origin is determined solely on the fact whether the *msg* field on the `logging.LogRecord` is a dict or not).

This allows for two interesting use cases:

1. You can format non-structlog log entries.
2. You can multiplex log records into multiple `logging.Handlers`.

Please refer to *Standard Library Logging* for examples.

#### Parameters

- **processor** (*callable*) – A structlog processor.
- **foreign\_pre\_chain** – If not *None*, it is used as an iterable of processors that is applied to non-structlog log entries before *processor*. If *None*, formatting is left to `logging`. (default: *None*)
- **keep\_exc\_info** (*bool*) – *exc\_info* on `logging.LogRecords` is added to the *event\_dict* and removed afterwards. Set this to *True* to keep it on the `logging.LogRecord`. (default: *False*)
- **keep\_stack\_info** (*bool*) – Same as *keep\_exc\_info* except for Python 3's *stack\_info*.

#### Return type `str`

New in version 17.1.0.

New in version 17.2.0: *keep\_exc\_info* and *keep\_stack\_info*

**static** `wrap_for_formatter(logger, name, event_dict)`

Wrap *logger*, *name*, and *event\_dict*.

The result is later unpacked by `ProcessorFormatter` when formatting log entries.

Use this static method as the renderer (i.e. final processor) if you want to use `ProcessorFormatter` in your `logging` configuration.

## twisted Module

Processors and tools specific to the `Twisted` networking engine.

See also *structlog's Twisted support*.

**class** structlog.twisted.**BoundLogger** (*logger, processors, context*)

Twisted-specific version of `structlog.BoundLogger`.

Works exactly like the generic one except that it takes advantage of knowing the logging methods in advance.

Use it like:

```
configure(
    wrapper_class=structlog.twisted.BoundLogger,
)
```

**bind** (*\*\*new\_values*)

Return a new logger with *new\_values* added to the existing ones.

**Return type** `self.__class__`

**err** (*event=None, \*\*kw*)

Process event and call `log.err()` with the result.

**msg** (*event=None, \*\*kw*)

Process event and call `log.msg()` with the result.

**new** (*\*\*new\_values*)

Clear context and binds *initial\_values* using `bind()`.

Only necessary with dict implementations that keep global state like those wrapped by `structlog.threadlocal.wrap_dict()` when threads are re-used.

**Return type** `self.__class__`

**unbind** (*\*keys*)

Return a new logger with *keys* removed from the context.

**Raises** `KeyError` – If the key is not part of the context.

**Return type** `self.__class__`

**class** structlog.twisted.**LoggerFactory**

Build a Twisted logger when an *instance* is called.

```
>>> from structlog import configure
>>> from structlog.twisted import LoggerFactory
>>> configure(logger_factory=LoggerFactory())
```

**\_\_call\_\_** (*\*args*)

Positional arguments are silently ignored.

**Rvalue** A new Twisted logger.

Changed in version 0.4.0: Added support for optional positional arguments.

**class** structlog.twisted.**EventAdapter** (*dictRenderer=None*)

Adapt an `event_dict` to Twisted logging system.

Particularly, make a wrapped `twisted.python.log.err` behave as expected.

**Parameters** `dictRenderer` (*callable*) – Renderer that is used for the actual log message.

Please note that structlog comes with a dedicated `JSONRenderer`.

**Must** be the last processor in the chain and requires a `dictRenderer` for the actual formatting as an constructor argument in order to be able to fully support the original behaviors of `log.msg()` and `log.err()`.

**class** `structlog.twisted.JSONRenderer` (*serializer=<function dumps>, \*\*dumps\_kw*)

Behaves like `structlog.processors.JSONRenderer` except that it formats tracebacks and failures itself if called with `err()`.

---

**Note:** This ultimately means that the messages get logged out using `msg()`, and *not* `err()` which renders failures in separate lines.

Therefore it will break your tests that contain assertions using `flushLoggedErrors`.

---

Not an adapter like `EventAdapter` but a real formatter. Nor does it require to be adapted using it.

Use together with a `JSONLogObserverWrapper`-wrapped Twisted logger like `plainJSONStdOutLogger()` for pure-JSON logs.

`structlog.twisted.plainJSONStdOutLogger()`

Return a logger that writes only the message to stdout.

Transforms non-`JSONRenderer` messages to JSON.

Ideal for JSONifying log entries from Twisted plugins and libraries that are outside of your control:

```
$ twistd -n --logger structlog.twisted.plainJSONStdOutLogger web
{"event": "Log opened.", "system": "-"}
{"event": "twistd 13.1.0 (python 2.7.3) starting up.", "system": "-"}
{"event": "reactor class: twisted...EPollReactor.", "system": "-"}
{"event": "Site starting on 8080", "system": "-"}
{"event": "Starting factory <twisted.web.server.Site ...>", ...}
...
```

Composes `PlainFileLogObserver` and `JSONLogObserverWrapper` to a usable logger.

New in version 0.2.0.

`structlog.twisted.JSONLogObserverWrapper` (*observer*)

Wrap a log *observer* and render non-`JSONRenderer` entries to JSON.

**Parameters** `observer` (*ILogObserver*) – Twisted log observer to wrap. For example

`PlainFileObserver` or Twisted's stock `FileLogObserver`

New in version 0.2.0.

**class** `structlog.twisted.PlainFileLogObserver` (*file*)

Write only the the plain message without timestamps or anything else.

Great to just print JSON to stdout where you catch it with something like `runit`.

**Parameters** `file` (*file*) – File to print to.

New in version 0.2.0.



---

## Project Information

---

`structlog` is dual-licensed under [Apache License, version 2](#) and [MIT](#), available from [PyPI](#), the source code can be found on [GitHub](#), the documentation at <http://www.structlog.org/>.

`structlog` targets Python 2.7, 3.4 and newer, and PyPy.

If you need any help, visit us on `#structlog` on [Freenode](#)!

## Backward Compatibility

`structlog` has a very strong backward compatibility policy that is inspired by the one of the [Twisted framework](#).

Put simply, you shouldn't ever be afraid to upgrade `structlog` if you're using its public APIs. If there will ever be need to break compatibility, it will be announced in the [Changelog](#) and raise deprecation warning for a year before it's finally really broken.

**Warning:** You cannot however rely on the default settings and the `structlog.dev` module. They may be adjusted in the future to provide a better experience when starting to use `structlog`. So please make sure to **always** properly configure your applications.

## How To Contribute

First off, thank you for considering contributing to `structlog`! It's people like *you* who make it is such a great tool for everyone.

This document is mainly to help you to get started by codifying tribal knowledge and expectations and make it more accessible to everyone. But don't be afraid to open half-finished PRs and ask questions if something is unclear!

## Workflow

- No contribution is too small! Please submit as many fixes for typos and grammar bloopers as you can!
- Try to limit each pull request to *one* change only.
- *Always* add tests and docs for your code. This is a hard rule; patches with missing tests or documentation can't be accepted.
- Make sure your changes pass our [CI](#). You won't get any feedback until it's green unless you ask for it.
- Once you've addressed review feedback, make sure to bump the pull request with a short note. Maintainers don't receive notifications when you push new commits.
- Don't break [backward compatibility](#).

## Code

- Obey [PEP 8](#) and [PEP 257](#). We use the `"""`-on-separate-lines style for docstrings:

```
def func(x):  
    """  
    Do something.  
  
    :param str x: A very important parameter.  
  
    :rtype: str  
    """
```

- If you add or change public APIs, tag the docstring using `.. versionadded:: 16.0.0 WHAT` or `.. versionchanged:: 17.1.0 WHAT`.
- Prefer double quotes (`"`) over single quotes (`'`) unless the string contains double quotes itself.

## Tests

- Write your asserts as `expected == actual` to line them up nicely:

```
x = f()  
  
assert 42 == x.some_attribute  
assert "foo" == x._a_private_attribute
```

- To run the test suite, all you need is a recent [tox](#). It will ensure the test suite runs with all dependencies against all Python versions just as it will on Travis CI. If you lack some Python versions, you can always limit the environments like `tox -e py27,py35` (in that case you may want to look into [pyenv](#), which makes it very easy to install many different Python versions in parallel).
- Write [good test docstrings](#).

## Documentation

- Use [semantic newlines](#) in [reStructuredText](#) files (files ending in `.rst`):

```
This is a sentence.  
This is another sentence.
```

- If you start a new section, add two blank lines before and one blank line after the header except if two headers follow immediately after each other:

```
Last line of previous section.

Header of New Top Section
-----

Header of New Section
^^^^^^^^^^^^^^^^^^^^

First line of new section.
```

- If your change is noteworthy, add an entry to the [changelog](#). Use [semantic newlines](#), and add a link to your pull request:

```
- Added ``structlog.func()`` that does foo.
  It's pretty cool.
  [#1 <https://github.com/hynek/structlog/pull/1>`_]
- ``structlog.func()`` now doesn't crash the Large Hadron Collider anymore.
  That was a nasty bug!
  [#2 <https://github.com/hynek/structlog/pull/2>`_]`_
```

## Local Development Environment

You can (and should) run our test suite using `tox` however you'll probably want a more traditional environment too. We highly recommend to develop using the latest Python 3 release because you're more likely to catch certain bugs earlier.

First create a [virtual environment](#). It's out of scope for this document to list all the ways to manage virtual environments in Python but if you don't have already a pet way, take some time to look at tools like [pew](#), [virtualfish](#), and [virtualenvwrapper](#).

Next get an up to date checkout of the `structlog` repository:

```
git checkout git@github.com:hynek/structlog.git
```

Change into the newly created directory and **after activating your virtual environment** install an editable version of `structlog`:

```
cd structlog
pip install -e .
```

If you run the virtual environment's Python and try to import `structlog` it should work!

To run the test suite, you'll need our development dependencies which can be installed using

```
pip install -r dev-requirements.txt
```

At this point

```
python -m pytest
```

should work and pass!

Again, this list is mainly to help you to get started by codifying tribal knowledge and expectations. If something is unclear, feel free to ask for help!

Please note that this project is released with a Contributor [Code of Conduct](#). By participating in this project you agree to abide by its terms. Please report any harm to [Hynek Schlawack](#) in any way you find appropriate.

Thank you for considering contributing to `structlog`!

## License and Hall of Fame

`structlog` is licensed both under the [Apache License, Version 2](#) and the [MIT license](#).

The reason for that is to be both protected against patent claims by own contributors and still allow the usage within GPLv2 software. For more legal details, see [this issue](#) on the bug tracker of PyCA's cryptography.

The full license texts can be also found in the source code repository:

- [Apache License 2.0](#)
- [MIT](#)

## Authors

`structlog` is written and maintained by [Hynek Schlawack](#). It's inspired by previous work done by [Jean-Paul Calderone](#) and [David Reid](#).

The development is kindly supported by [Variomedia AG](#).

A full list of contributors can be found on GitHub's [overview](#). Some of them disapprove of the addition of thread local context data. :)

The `structlog` logo has been contributed by [Russell Keith-Magee](#).

## Changelog

Versions are year-based with a strict backward compatibility policy. The third digit is only for regressions.

### 17.2.0 (2017-05-15)

#### Backward-incompatible changes:

*none*

#### Deprecations:

*none*

**Changes:**

- `structlog.stdlib.ProcessorFormatter` now accepts `keep_exc_info` and `keep_stack_info` arguments to control what to do with this information on log records. Most likely you want them both to be `False` therefore it's the default. [#109](#)
  - `structlog.stdlib.add_logger_name()` now works in `structlog.stdlib.ProcessorFormatter`'s `foreign_pre_chain`. [#112](#)
  - Clear log record args in `structlog.stdlib.ProcessorFormatter` after rendering. This fix is for you if you tried to use it and got `TypeError: not all arguments converted during string formatting` exceptions. [#116](#) [#117](#)
- 

**17.1.0 (2017-04-24)**

The main features of this release are massive improvements in standard library's logging integration. Have a look at the updated [standard library chapter](#) on how to use them! Special thanks go to [Fabian Büchler](#), [Gilbert Gilb's](#), [Iva Kaneva](#), [insolite](#), and [sky-code](#), that made them possible.

**Backward-incompatible changes:**

- The default renderer now is `structlog.dev.ConsoleRenderer` if you don't configure `structlog`. Colors are used if available and human-friendly timestamps are prepended. This is in line with our [backward compatibility policy](#) that explicitly excludes default settings.

**Changes:**

- Added `structlog.stdlib.render_to_log_kwargs()`. This allows you to use logging-based formatters to take care of rendering your entries. [#98](#)
  - Added `structlog.stdlib.ProcessorFormatter` which does the opposite: This allows you to run `structlog` processors on arbitrary logging `LogRecords`. [#79](#) [#105](#)
  - UNIX epoch timestamps from `structlog.processors.TimeStamper` are more precise now.
  - Added `repr_native_str` to `structlog.processors.KeyValueRenderer` and `structlog.dev.ConsoleRenderer`. This allows for human-readable non-ASCII output on Python 2 (`repr()` on Python 2 behaves like `ascii()` on Python 3 in that regard). As per compatibility policy, it's on (original behavior) in `KeyValueRenderer` and off (humand-friendly behavior) in `ConsoleRenderer`. [#94](#)
  - Added `colors` argument to `structlog.dev.ConsoleRenderer` and made it the default renderer. [#78](#)
  - Fixed bug with Python 3 and `structlog.stdlib BoundLogger.log()`. Error log level was not reproducible and was logged as exception one time out of two. [#92](#)
  - Positional arguments are now removed even if they are empty. [#82](#)
-

## 16.1.0 (2016-05-24)

### Backward-incompatible changes:

- Python 3.3 and 2.6 aren't supported anymore. They may work by chance but any effort to keep them working has ceased.

The last Python 2.6 release was on October 29, 2013 and isn't supported by the CPython core team anymore. Major Python packages like Django and Twisted dropped Python 2.6 a while ago already.

Python 3.3 never had a significant user base and wasn't part of any distribution's LTS release.

### Changes:

- Add a `drop_missing` argument to `KeyValueRenderer`. If `key_order` is used and a key is missing a value, it's not rendered at all instead of being rendered as `None`. [#67](#)
  - Exceptions without a `__traceback__` are now also rendered on Python 3.
  - Don't cache loggers in lazy proxies returned from `get_logger()`. This lead to in-place mutation of them if used before configuration which in turn lead to the problem that configuration was applied only partially to them later. [#72](#)
- 

## 16.0.0 (2016-01-28)

### Changes:

- `structlog.processors.ExceptionPrettyPrinter` and `structlog.processors.format_exc_info` now support passing of Exceptions on Python 3.
  - Clean up the context when exiting `structlog.threadlocal.tmp_bind` in case of exceptions. [#64](#)
  - Be more more lenient about missing `__name__`s. [#62](#)
  - Add `structlog.dev.ConsoleRenderer` that renders the event dictionary aligned and with colors.
  - Use `six` for compatibility.
  - Add `structlog.processors.UnicodeDecoder` that will decode all byte string values in an event dictionary to Unicode.
  - Add `serializer` parameter to `structlog.processors.JSONRenderer` which allows for using different (possibly faster) JSON encoders than the standard library.
- 

## 15.3.0 (2015-09-25)

### Changes:

- Tolerate frames without a `__name__`, better. [#58](#)
- Officially support Python 3.5.
- Add `structlog.ReturnLogger.failure` and `structlog.PrintLogger.failure` as preparation for the new Twisted logging system.

## 15.2.0 (2015-06-10)

### Changes:

- Allow empty lists of processors. This is a valid use case since [#26](#) has been merged. Before, supplying an empty list resulted in the defaults being used.
  - Prevent Twisted's `log.err` from quoting strings rendered by `structlog.twisted.JSONRenderer`.
  - Better support of logging.Logger.exception within structlog. [#52](#)
  - Add option to specify target key in `structlog.processors.TimeStamper` processor. [#51](#)
- 

## 15.1.0 (2015-02-24)

### Changes:

- Tolerate frames without a `__name__`.
- 

## 15.0.0 (2015-01-23)

### Changes:

- Add `structlog.stdlib.add_log_level` and `structlog.stdlib.add_logger_name` processors. [#44](#)
  - Add `structlog.stdlib.BindLogger.log`. [#42](#)
  - Pass positional arguments to stdlib wrapped loggers that use string formatting. [#19](#)
  - structlog is now dually licensed under the [Apache License, Version 2](#) and the [MIT](#) license. Therefore it is now legal to use structlog with GPLv2-licensed projects. [#28](#)
  - Add `structlog.stdlib.BindLogger.exception`. [#22](#)
- 

## 0.4.2 (2014-07-26)

### Changes:

- Fixed a memory leak in greenlet code that emulates thread locals. It shouldn't matter in practice unless you use multiple wrapped dicts within one program that is rather unlikely. [#8](#)
- `structlog.PrintLogger` now is thread-safe.
- Test Twisted-related code on Python 3 (with some caveats).

- Drop support for Python 3.2. There is no justification to add complexity for a Python version that nobody uses. If you are one of the [0.350%](#) that use Python 3.2, please stick to the 0.4 branch; critical bugs will still be fixed.
  - Officially support Python 3.4.
  - Allow final processor to return a dictionary. See the adapting chapter. [#26](#)
  - `from structlog import *` works now (but you still shouldn't use it).
- 

### 0.4.1 (2013-12-19)

#### Changes:

- Don't cache proxied methods in `structlog.threadlocal._ThreadLocalDictWrapper`. This doesn't affect regular users.
  - Various doc fixes.
- 

### 0.4.0 (2013-11-10)

#### Backward-incompatible changes:

#### Changes:

- Add `structlog.processors.StackInfoRenderer` for adding stack information to log entries without involving exceptions. Also added it to default processor chain. [#6](#)
  - Allow optional positional arguments for `structlog.get_logger` that are passed to logger factories. The standard library factory uses this for explicit logger naming. [#12](#)
  - Add `structlog.processors.ExceptionPrettyPrinter` for development and testing when multi-line log entries aren't just acceptable but even helpful.
  - Allow the standard library name guesser to ignore certain frame names. This is useful together with frameworks.
  - Add meta data (e.g. function names, line numbers) extraction for wrapped stdlib loggers. [#5](#)
- 

### 0.3.2 (2013-09-27)

#### Changes:

- Fix stdlib's name guessing.
-



### 0.3.1 (2013-09-26)

#### Changes:

- Add forgotten `structlog.processors.TimeStamper` to API documentation.
- 

### 0.3.0 (2013-09-23)

#### Changes:

- Greatly enhanced and polished the documentation and added a new theme based on Write The Docs, requests, and Flask.
  - Add Python Standard Library-specific `BoundLogger` that has an explicit API instead of intercepting unknown method calls. See `structlog.stdlib.BoundLogger`.
  - `structlog.ReturnLogger` now allows arbitrary positional and keyword arguments.
  - Add Twisted-specific `BoundLogger` that has an explicit API instead of intercepting unknown method calls. See `structlog.twisted.BoundLogger`.
  - Allow logger proxies that are returned by `structlog.get_logger` and `structlog.wrap_logger` to cache the `BoundLogger` they assemble according to configuration on first use. See the chapter on performance and the `cache_logger_on_first_use` argument of `structlog.configure` and `structlog.wrap_logger`.
  - Extract a common base class for loggers that does nothing except keeping the context state. This makes writing custom loggers much easier and more straight-forward. See `structlog.BoundLoggerBase`.
- 

### 0.2.0 (2013-09-17)

#### Changes:

- Promote to stable, thus henceforth a strict backward compatibility policy is put into effect.
  - Add `key_order` option to `structlog.processors.KeyValueRenderer` for more predictable log entries with any `dict` class.
  - `structlog.PrintLogger` now uses proper I/O routines and is thus viable not only for examples but also for production.
  - Enhance Twisted support by offering JSONification of non-structlog log entries.
  - Allow for custom serialization in `structlog.twisted.JSONRenderer` without abusing `__repr__`.
- 

### 0.1.0 (2013-09-16)

Initial release.



## CHAPTER 9

---

### Indices and tables

---

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



### S

- `structlog`, [39](#)
- `structlog.dev`, [44](#)
- `structlog.processors`, [45](#)
- `structlog.stdlib`, [48](#)
- `structlog.threadlocal`, [45](#)
- `structlog.twisted`, [51](#)



## Symbols

`__call__()` (structlog.stdlib.LoggerFactory method), 49  
`__call__()` (structlog.twisted.LoggerFactory method), 51  
`_logger` (structlog.BoundLoggerBase attribute), 43  
`_process_event()` (structlog.BoundLoggerBase method), 43  
`_proxy_to_logger()` (structlog.BoundLoggerBase method), 43

## A

`add_log_level()` (in module structlog.stdlib), 50  
`add_logger_name()` (in module structlog.stdlib), 50  
`as_immutable()` (in module structlog.threadlocal), 45

## B

`bind()` (structlog.BoundLogger method), 41  
`bind()` (structlog.BoundLoggerBase method), 44  
`bind()` (structlog.stdlib.BoundLogger method), 48  
`bind()` (structlog.twisted.BoundLogger method), 51  
`BoundLogger` (class in structlog), 41  
`BoundLogger` (class in structlog.stdlib), 48  
`BoundLogger` (class in structlog.twisted), 51  
`BoundLoggerBase` (class in structlog), 43

## C

`configure()` (in module structlog), 40  
`configure_once()` (in module structlog), 40  
`ConsoleRenderer` (class in structlog.dev), 44  
`critical()` (structlog.PrintLogger method), 41  
`critical()` (structlog.ReturnLogger method), 42  
`critical()` (structlog.stdlib.BoundLogger method), 48

## D

`debug()` (structlog.PrintLogger method), 41  
`debug()` (structlog.ReturnLogger method), 42  
`debug()` (structlog.stdlib.BoundLogger method), 48  
`DropEvent`, 43

## E

`err()` (structlog.PrintLogger method), 41  
`err()` (structlog.ReturnLogger method), 42  
`err()` (structlog.twisted.BoundLogger method), 51  
`error()` (structlog.PrintLogger method), 41  
`error()` (structlog.ReturnLogger method), 42  
`error()` (structlog.stdlib.BoundLogger method), 48  
`EventAdapter` (class in structlog.twisted), 51  
`exception()` (structlog.stdlib.BoundLogger method), 48  
`ExceptionPrettyPrinter` (class in structlog.processors), 47

## F

`failure()` (structlog.PrintLogger method), 41  
`failure()` (structlog.ReturnLogger method), 42  
`filter_by_level()` (in module structlog.stdlib), 49  
`format_exc_info()` (in module structlog.processors), 47

## G

`get_logger()` (in module structlog), 39  
`getLogger()` (in module structlog), 39

## I

`info()` (structlog.PrintLogger method), 41  
`info()` (structlog.ReturnLogger method), 42  
`info()` (structlog.stdlib.BoundLogger method), 48

## J

`JSONLogObserverWrapper()` (in module structlog.twisted), 52  
`JSONRenderer` (class in structlog.processors), 45  
`JSONRenderer` (class in structlog.twisted), 52

## K

`KeyValueRenderer` (class in structlog.processors), 46

## L

`log()` (structlog.PrintLogger method), 42  
`log()` (structlog.ReturnLogger method), 42  
`log()` (structlog.stdlib.BoundLogger method), 48

LoggerFactory (class in structlog.stdlib), [49](#)  
LoggerFactory (class in structlog.twisted), [51](#)

## M

msg() (structlog.PrintLogger method), [42](#)  
msg() (structlog.ReturnLogger method), [42](#)  
msg() (structlog.twisted.BoundLogger method), [51](#)

## N

new() (structlog.BoundLogger method), [41](#)  
new() (structlog.BoundLoggerBase method), [44](#)  
new() (structlog.stdlib.BoundLogger method), [48](#)  
new() (structlog.twisted.BoundLogger method), [51](#)

## P

PlainFileLogObserver (class in structlog.twisted), [52](#)  
plainJSONStdOutLogger() (in module structlog.twisted),  
[52](#)  
PositionalArgumentsFormatter (class in structlog.stdlib),  
[50](#)  
PrintLogger (class in structlog), [41](#)  
PrintLoggerFactory (class in structlog), [42](#)  
ProcessorFormatter (class in structlog.stdlib), [50](#)

## R

render\_to\_log\_kwargs() (in module structlog.stdlib), [49](#)  
reset\_defaults() (in module structlog), [40](#)  
ReturnLogger (class in structlog), [42](#)  
ReturnLoggerFactory (class in structlog), [42](#)

## S

StackInfoRenderer (class in structlog.processors), [47](#)  
structlog (module), [39](#)  
structlog.dev (module), [44](#)  
structlog.processors (module), [45](#)  
structlog.stdlib (module), [48](#)  
structlog.threadlocal (module), [45](#)  
structlog.twisted (module), [51](#)

## T

TimeStamper (class in structlog.processors), [47](#)  
tmp\_bind() (in module structlog.threadlocal), [45](#)

## U

unbind() (structlog.BoundLogger method), [41](#)  
unbind() (structlog.BoundLoggerBase method), [44](#)  
unbind() (structlog.stdlib.BoundLogger method), [48](#)  
unbind() (structlog.twisted.BoundLogger method), [51](#)  
UnicodeDecoder (class in structlog.processors), [46](#)  
UnicodeEncoder (class in structlog.processors), [46](#)

## W

warn() (structlog.stdlib.BoundLogger method), [49](#)

warning() (structlog.PrintLogger method), [42](#)  
warning() (structlog.ReturnLogger method), [42](#)  
warning() (structlog.stdlib.BoundLogger method), [49](#)  
wrap\_dict() (in module structlog.threadlocal), [45](#)  
wrap\_for\_formatter() (struct-  
log.stdlib.ProcessorFormatter static method),  
[50](#)  
wrap\_logger() (in module structlog), [40](#)