
aiocache Documentation

Release 0.8.0

Manuel Miranda

Jan 23, 2018

Contents

1	Installing	1
2	Usage	3
3	Contents	7
3.1	Caches	7
3.2	Serializers	8
3.3	Plugins	11
3.4	Configuration	13
3.5	Decorators	14
3.6	Locking	16
3.7	Testing	16
4	Indices and tables	19

CHAPTER 1

Installing

- `pip install aiocache`
- `pip install aiocache[redis]`
- `pip install aiocache[memcached]`
- `pip install aiocache[redis,memcached]`

Using a cache is as simple as

```
>>> import asyncio
>>> loop = asyncio.get_event_loop()
>>> from aiocache import SimpleMemoryCache
>>> cache = SimpleMemoryCache()
>>> loop.run_until_complete(cache.set('key', 'value'))
True
>>> loop.run_until_complete(cache.get('key'))
'value'
```

Here we are using the *SimpleMemoryCache* but you can use any other listed in *Caches*. All caches contain the same minimum interface which consists on the following functions:

- `add`: Only adds key/value if key does not exist. Otherwise raises `ValueError`.
- `get`: Retrieve value identified by key.
- `set`: Sets key/value.
- `multi_get`: Retrieves multiple key/values.
- `multi_set`: Sets multiple key/values.
- `exists`: Returns `True` if key exists `False` otherwise.
- `increment`: Increment the value stored in the given key.
- `delete`: Deletes key and returns number of deleted items.
- `clear`: Clears the items stored.
- `raw`: Executes the specified command using the underlying client.

You can also setup cache aliases like in Django settings:

```
1 import asyncio
2
3 from aiocache import caches, SimpleMemoryCache, RedisCache
```

```
4 from aiocache.serializers import StringSerializer, PickleSerializer
5
6 caches.set_config({
7     'default': {
8         'cache': "aiocache.SimpleMemoryCache",
9         'serializer': {
10            'class': "aiocache.serializers.StringSerializer"
11        }
12    },
13    'redis_alt': {
14        'cache': "aiocache.RedisCache",
15        'endpoint': "127.0.0.1",
16        'port': 6379,
17        'timeout': 1,
18        'serializer': {
19            'class': "aiocache.serializers.PickleSerializer"
20        },
21        'plugins': [
22            {'class': "aiocache.plugins.HitMissRatioPlugin"},
23            {'class': "aiocache.plugins.TimingPlugin"}
24        ]
25    }
26 })
27
28
29 async def default_cache():
30     cache = caches.get('default') # This always returns the same instance
31     await cache.set("key", "value")
32
33     assert await cache.get("key") == "value"
34     assert isinstance(cache, SimpleMemoryCache)
35     assert isinstance(cache.serializer, StringSerializer)
36
37
38 async def alt_cache():
39     # This generates a new instance every time! You can also use `caches.create('alt
40     ↪')`
41     # or even `caches.create('alt', namespace="test", etc...)` to override extra args
42     cache = caches.create(**caches.get_alias_config('redis_alt'))
43     await cache.set("key", "value")
44
45     assert await cache.get("key") == "value"
46     assert isinstance(cache, RedisCache)
47     assert isinstance(cache.serializer, PickleSerializer)
48     assert len(cache.plugins) == 2
49     assert cache.endpoint == "127.0.0.1"
50     assert cache.timeout == 1
51     assert cache.port == 6379
52     await cache.close()
53
54 def test_alias():
55     loop = asyncio.get_event_loop()
56     loop.run_until_complete(default_cache())
57     loop.run_until_complete(alt_cache())
58
59     cache = RedisCache()
60     loop.run_until_complete(cache.delete("key"))
```



```
61     loop.run_until_complete(cache.close())
62
63     loop.run_until_complete(caches.get('default').close())
64
65
66 if __name__ == "__main__":
67     test_alias()
```

In examples folder you can check different use cases:

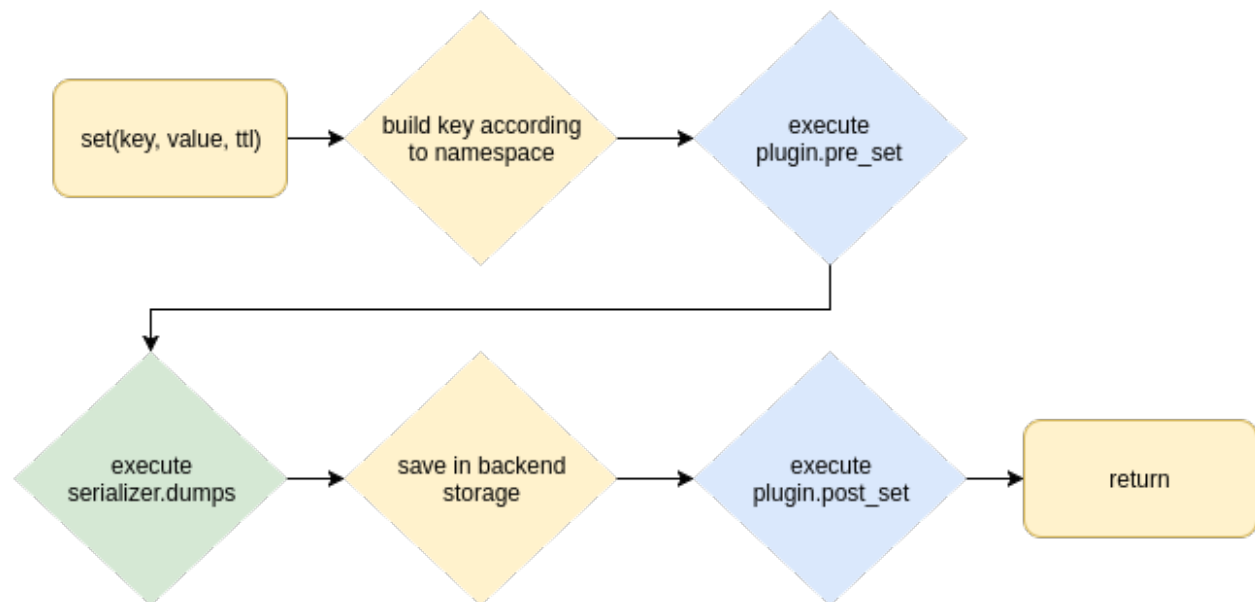
- Sanic, Aiohttp and Tornado
- Python object in Redis
- Custom serializer for compressing data
- TimingPlugin and HitMissRatioPlugin demos
- Using marshmallow as a serializer
- Using cached decorator.
- Using multi_cached decorator.

3.1 Caches

You can use different caches according to your needs. All the caches implement the same interface.

Caches are always working together with a serializer which transforms data when storing and retrieving from the backend. It may also contain plugins that are able to enrich the behavior of your cache (like adding metrics, logs, etc).

This is the flow of the `set` command:



Let's go with a more specific case. Let's pick Redis as the cache with namespace "test" and PickleSerializer as the serializer:

1. We receive `set("key", "value")`.

2. Hook `pre_set` of all attached plugins (none by default) is called.
3. “key” will become “test:key” when calling `build_key`.
4. “value” will become an array of bytes when calling `serializer.dumps` because of `PickleSerializer`.
5. the byte array is stored together with the key using `set` cmd in Redis.
6. Hook `post_set` of all attached plugins is called.

By default, all commands are covered by a timeout that will trigger an `asyncio.TimeoutError` in case of timeout. Timeout can be set at instance level or when calling the command.

The supported commands are:

- `add`
- `get`
- `set`
- `multi_get`
- `multi_set`
- `delete`
- `exists`
- `increment`
- `expire`
- `clear`
- `raw`

If you feel a command is missing here do not hesitate to [open an issue](#)

3.1.1 BaseCache

3.1.2 RedisCache

3.1.3 SimpleMemoryCache

3.1.4 MemcachedCache

3.2 Serializers

Serializers can be attached to backends in order to serialize/deserialize data sent and retrieved from the backend. This allows to apply transformations to data in case you want it to be saved in a specific format in your cache backend. For example, imagine you have your `Model` and want to serialize it to something that Redis can understand (Redis can't store python objects). This is the task of a serializer.

To use a specific serializer:

```
>>> from aiocache import SimpleMemoryCache
>>> from aiocache.serializers import PickleSerializer
cache = SimpleMemoryCache(serializer=PickleSerializer())
```

Currently the following are built in:

3.2.1 NullSerializer

3.2.2 StringSerializer

3.2.3 PickleSerializer

3.2.4 JsonSerializer

3.2.5 MessagePackSerializer

In case the current serializers are not covering your needs, you can always define your custom serializer as shown in `examples/serializer_class.py`:

```

1 import asyncio
2 import zlib
3
4 from aiocache import RedisCache
5 from aiocache.serializers import BaseSerializer
6
7
8 class CompressionSerializer(BaseSerializer):
9
10     # This is needed because zlib works with bytes.
11     # this way the underlying backend knows how to
12     # store/retrieve values
13     DEFAULT_ENCODING = None
14
15     def dumps(self, value):
16         print("I've received:\n{}".format(value))
17         compressed = zlib.compress(value.encode())
18         print("But I'm storing:\n{}".format(compressed))
19         return compressed
20
21     def loads(self, value):
22         print("I've retrieved:\n{}".format(value))
23         decompressed = zlib.decompress(value).decode()
24         print("But I'm returning:\n{}".format(decompressed))
25         return decompressed
26
27
28 cache = RedisCache(serializer=CompressionSerializer(), namespace="main")
29
30
31 async def serializer():
32     text = (
33         "Lorem ipsum dolor sit amet, consectetur adipiscing elit, sed do eiusmod_
↵tempor incididunt"
34         "ut labore et dolore magna aliqua. Ut enim ad minim veniam, quis nostrud_
↵exercitation"
35         "ullamco laboris nisi ut aliquip ex ea commodo consequat. Duis aute irure_
↵dolor in"
36         "reprehenderit in voluptate velit esse cillum dolore eu fugiat nulla pariatur.
↵ Excepteur"
37         "sint occaecat cupidatat non proident, sunt in culpa qui officia deserunt_
↵mollit"
38         "anim id est laborum.")
39     await cache.set("key", text)

```

```

40     print("-----")
41     real_value = await cache.get("key")
42     compressed_value = await cache.raw("get", "main:key")
43     assert len(compressed_value) < len(real_value.encode())
44
45
46 def test_serializer():
47     loop = asyncio.get_event_loop()
48     loop.run_until_complete(serializer())
49     loop.run_until_complete(cache.delete("key"))
50     loop.run_until_complete(cache.close())
51
52
53 if __name__ == "__main__":
54     test_serializer()

```

You can also use marshmallow as your serializer (examples/marshmallow_serializer_class.py):

```

1  import random
2  import string
3  import asyncio
4
5  from marshmallow import fields, Schema, post_load
6
7  from aiocache import SimpleMemoryCache
8  from aiocache.serializers import BaseSerializer
9
10
11 class RandomModel:
12     MY_CONSTANT = "CONSTANT"
13
14     def __init__(self, int_type=None, str_type=None, dict_type=None, list_type=None):
15         self.int_type = int_type or random.randint(1, 10)
16         self.str_type = str_type or random.choice(string.ascii_lowercase)
17         self.dict_type = dict_type or {}
18         self.list_type = list_type or []
19
20     def __eq__(self, obj):
21         return self.__dict__ == obj.__dict__
22
23
24 class MarshmallowSerializer(Schema, BaseSerializer):
25     int_type = fields.Integer()
26     str_type = fields.String()
27     dict_type = fields.Dict()
28     list_type = fields.List(fields.Integer())
29
30     # marshmallow Schema class doesn't play nicely with multiple inheritance and won
31     ↪ 't call
32     # BaseSerializer.__init__
33     encoding = 'utf-8'
34
35     def dumps(self, *args, **kwargs):
36         # dumps returns (data, errors), we just want to save data
37         return super().dumps(*args, **kwargs).data
38
39     def loads(self, *args, **kwargs):
40         # dumps returns (data, errors), we just want to return data

```

```

40     return super().loads(*args, **kwargs).data
41
42     @post_load
43     def build_my_type(self, data):
44         return RandomModel(**data)
45
46     class Meta:
47         strict = True
48
49
50 cache = SimpleMemoryCache(serializer=MarshmallowSerializer(), namespace="main")
51
52
53 async def serializer():
54     model = RandomModel()
55     await cache.set("key", model)
56
57     result = await cache.get("key")
58
59     assert result.int_type == model.int_type
60     assert result.str_type == model.str_type
61     assert result.dict_type == model.dict_type
62     assert result.list_type == model.list_type
63
64
65 def test_serializer():
66     loop = asyncio.get_event_loop()
67     loop.run_until_complete(serializer())
68     loop.run_until_complete(cache.delete("key"))
69
70
71 if __name__ == "__main__":
72     test_serializer()

```

By default cache backends assume they are working with `str` types. If your custom implementation transform data to bytes, you will need to set the class attribute `encoding` to `None`.

3.3 Plugins

Plugins can be used to enrich the behavior of the cache. By default all caches are configured without any plugin but can add new ones in the constructor or after initializing the cache class:

```

>>> from aiocache import SimpleMemoryCache
>>> from aiocache.plugins import TimingPlugin
cache = SimpleMemoryCache(plugings=[HitMissRatioPlugin()])
cache.plugings += [TimingPlugin()]

```

You can define your custom plugin by inheriting from `BasePlugin` and overriding the needed methods (the overrides NEED to be `async`). All commands have `pre_<command_name>` and `post_<command_name>` hooks.

Warning: Both pre and post hooks are executed awaiting the coroutine. If you perform expensive operations with the hooks, you will add more latency to the command being executed and thus, there are more probabilities of raising a timeout error.

A complete example of using plugins:

```
1 import asyncio
2 import random
3 import logging
4
5 from aiocache import SimpleMemoryCache
6 from aiocache.plugins import HitMissRatioPlugin, TimingPlugin, BasePlugin
7
8
9 logger = logging.getLogger(__name__)
10
11
12 class MyCustomPlugin(BasePlugin):
13
14     async def pre_set(self, *args, **kwargs):
15         logger.info("I'm the pre_set hook being called with %s %s" % (args, kwargs))
16
17     async def post_set(self, *args, **kwargs):
18         logger.info("I'm the post_set hook being called with %s %s" % (args, kwargs))
19
20
21 cache = SimpleMemoryCache(
22     plugins=[HitMissRatioPlugin(), TimingPlugin(), MyCustomPlugin()],
23     namespace="main")
24
25
26 async def run():
27     await cache.set("a", "1")
28     await cache.set("b", "2")
29     await cache.set("c", "3")
30     await cache.set("d", "4")
31
32     possible_keys = ["a", "b", "c", "d", "e", "f"]
33
34     for t in range(1000):
35         await cache.get(random.choice(possible_keys))
36
37     assert cache.hit_miss_ratio["hit_ratio"] > 0.5
38     assert cache.hit_miss_ratio["total"] == 1000
39
40     assert cache.profiling["get_min"] > 0
41     assert cache.profiling["set_min"] > 0
42     assert cache.profiling["get_max"] > 0
43     assert cache.profiling["set_max"] > 0
44
45     print(cache.hit_miss_ratio)
46     print(cache.profiling)
47
48
49 def test_run():
50     loop = asyncio.get_event_loop()
51     loop.run_until_complete(run())
52     loop.run_until_complete(cache.delete("a"))
53     loop.run_until_complete(cache.delete("b"))
54     loop.run_until_complete(cache.delete("c"))
55     loop.run_until_complete(cache.delete("d"))
56
57
```



```

58 if __name__ == "__main__":
59     test_run()

```

3.3.1 BasePlugin

3.3.2 TimingPlugin

3.3.3 HitMissRatioPlugin

3.4 Configuration

3.4.1 Cache aliases

The caches module allows to setup cache configurations and then use them either using an alias or retrieving the config explicitly. To set the config, call `caches.set_config`:

To retrieve a copy of the current config, you can use `caches.get_config` or `caches.get_alias_config` for an alias config.

Next snippet shows an example usage:

```

1  import asyncio
2
3  from aiocache import caches, SimpleMemoryCache, RedisCache
4  from aiocache.serializers import StringSerializer, PickleSerializer
5
6  caches.set_config({
7      'default': {
8          'cache': "aiocache.SimpleMemoryCache",
9          'serializer': {
10             'class': "aiocache.serializers.StringSerializer"
11         }
12     },
13     'redis_alt': {
14         'cache': "aiocache.RedisCache",
15         'endpoint': "127.0.0.1",
16         'port': 6379,
17         'timeout': 1,
18         'serializer': {
19             'class': "aiocache.serializers.PickleSerializer"
20         },
21         'plugins': [
22             {'class': "aiocache.plugins.HitMissRatioPlugin"},
23             {'class': "aiocache.plugins.TimingPlugin"}
24         ]
25     }
26 })
27
28
29 async def default_cache():
30     cache = caches.get('default') # This always returns the same instance
31     await cache.set("key", "value")
32
33     assert await cache.get("key") == "value"

```

```

34     assert isinstance(cache, SimpleMemoryCache)
35     assert isinstance(cache.serializer, StringSerializer)
36
37
38 async def alt_cache():
39     # This generates a new instance every time! You can also use `caches.create('alt
40     ↪')`
41     # or even `caches.create('alt', namespace="test", etc...)` to override extra args
42     cache = caches.create(**caches.get_alias_config('redis_alt'))
43     await cache.set("key", "value")
44
45     assert await cache.get("key") == "value"
46     assert isinstance(cache, RedisCache)
47     assert isinstance(cache.serializer, PickleSerializer)
48     assert len(cache.plugins) == 2
49     assert cache.endpoint == "127.0.0.1"
50     assert cache.timeout == 1
51     assert cache.port == 6379
52     await cache.close()
53
54 def test_alias():
55     loop = asyncio.get_event_loop()
56     loop.run_until_complete(default_cache())
57     loop.run_until_complete(alt_cache())
58
59     cache = RedisCache()
60     loop.run_until_complete(cache.delete("key"))
61     loop.run_until_complete(cache.close())
62
63     loop.run_until_complete(caches.get('default').close())
64
65
66 if __name__ == "__main__":
67     test_alias()

```

When you do `caches.get('alias_name')`, the cache instance is built lazily the first time. Next accesses will return the **same** instance. If instead of reusing the same instance, you need a new one every time, use `caches.create('alias_name')`. One of the advantages of `caches.create` is that it accepts extra args that then are passed to the cache constructor. This way you can override args like `namespace`, `endpoint`, etc.

3.5 Decorators

aiocache comes with a couple of decorators for caching results from asynchronous functions. Do not use the decorator in synchronous functions, it may lead to unexpected behavior.

3.5.1 cached

```

1 import asyncio
2
3 from collections import namedtuple
4
5 from aiocache import cached, RedisCache
6 from aiocache.serializers import PickleSerializer

```

```

7
8 Result = namedtuple('Result', "content, status")
9
10
11 @cached(
12     ttl=10, cache=RedisCache, key="key", serializer=PickleSerializer(), port=6379,
13     ↪namespace="main")
14 async def cached_call():
15     return Result("content", 200)
16
17 def test_cached():
18     cache = RedisCache(endpoint="127.0.0.1", port=6379, namespace="main")
19     loop = asyncio.get_event_loop()
20     loop.run_until_complete(cached_call())
21     assert loop.run_until_complete(cache.exists("key")) is True
22     loop.run_until_complete(cache.delete("key"))
23     loop.run_until_complete(cache.close())
24
25
26 if __name__ == "__main__":
27     test_cached()

```

3.5.2 multi_cached

```

1 import asyncio
2
3 from aiocache import multi_cached, RedisCache
4
5 DICT = {
6     'a': "Z",
7     'b': "Y",
8     'c': "X",
9     'd': "W"
10 }
11
12
13 @multi_cached("ids", cache=RedisCache, namespace="main")
14 async def multi_cached_ids(ids=None):
15     return {id_: DICT[id_] for id_ in ids}
16
17
18 @multi_cached("keys", cache=RedisCache, namespace="main")
19 async def multi_cached_keys(keys=None):
20     return {id_: DICT[id_] for id_ in keys}
21
22
23 cache = RedisCache(endpoint="127.0.0.1", port=6379, namespace="main")
24
25
26 def test_multi_cached():
27     loop = asyncio.get_event_loop()
28     loop.run_until_complete(multi_cached_ids(ids=['a', 'b']))
29     loop.run_until_complete(multi_cached_ids(ids=['a', 'c']))
30     loop.run_until_complete(multi_cached_keys(keys=['d']))
31

```

```
32     assert loop.run_until_complete(cache.exists('a'))
33     assert loop.run_until_complete(cache.exists('b'))
34     assert loop.run_until_complete(cache.exists('c'))
35     assert loop.run_until_complete(cache.exists('d'))
36
37     loop.run_until_complete(cache.delete("a"))
38     loop.run_until_complete(cache.delete("b"))
39     loop.run_until_complete(cache.delete("c"))
40     loop.run_until_complete(cache.delete("d"))
41     loop.run_until_complete(cache.close())
42
43
44 if __name__ == "__main__":
45     test_multi_cached()
```

Warning: This was added in version 0.7.0 and the API is new. This means its open to breaking changes in future versions until the API is considered stable.

3.6 Locking

Warning: The implementations provided are **NOT** intended for consistency/synchronization purposes. If you need a locking mechanism focused on consistency, consider implementing your mechanism based on more serious tools like <https://zookeeper.apache.org/>.

There are a couple of locking implementations than can help you to protect against different scenarios:

3.6.1 RedLock

3.6.2 OptimisticLock

3.7 Testing

It's really easy to cut the dependency with aiocache functionality:

```
import asyncio

from unittest import MagicMock

from aiocache.base import BaseCache

async def async_main():
    mocked_cache = MagicMock(spec=BaseCache)
    mocked_cache.get.return_value = "world"
    print(await mocked_cache.get("hello"))

if __name__ == "__main__":
```

```
loop = asyncio.get_event_loop()
loop.run_until_complete(async_main())
```

Note that we are passing the *BaseCache* as the spec for the Mock (you need to install `asynctest`).

Also, for debugging purposes you can use `AIOCACHE_DISABLE = 1 python myscript.py` to disable caching.

CHAPTER 4

Indices and tables

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