
pycached Documentation

Release 0.0.9

Manuel Miranda

May 08, 2019

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
4	Indices and tables	17

CHAPTER 1

Installing

- pip install pycached
- pip install pycached[redis]

CHAPTER 2

Usage

Using a cache is as simple as

```
>>> from pycached import Cache
>>> cache = Cache()
>>> cache.set('key', 'value')
True
>>> 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 from pycached import caches, SimpleMemoryCache, RedisCache
2 from pycached.serializers import StringSerializer, PickleSerializer
3
4 caches.set_config({
```

(continues on next page)

(continued from previous page)

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

(continues on next page)

(continued from previous page)

```
61  
62  
63 if __name__ == "__main__":  
64     test_alias()
```

In examples folder you can check different use cases:

- 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.

CHAPTER 3

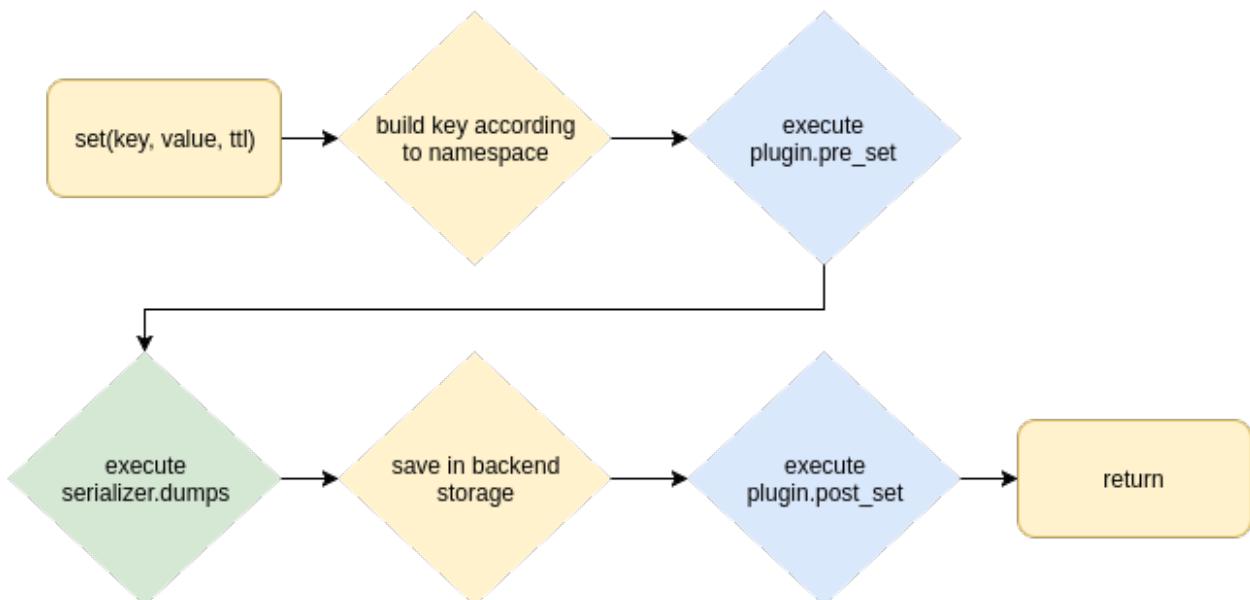
Contents

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 `timeout_decorator.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

3.2 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 pycached import SimpleMemoryCache
>>> from pycached.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 MsgPackSerializer

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 zlib
2
3 from pycached import Cache
4 from pycached.serializers import BaseSerializer
5
6
7 class CompressionSerializer(BaseSerializer):
8
9     # This is needed because zlib works with bytes.
10    # this way the underlying backend knows how to
11    # store/retrieve values
12    DEFAULT_ENCODING = None
13
14    def dumps(self, value):
15        print("I've received:\n{}".format(value))
16        compressed = zlib.compress(value.encode())
17        print("But I'm storing:\n{}".format(compressed))
18        return compressed
19
20    def loads(self, value):
21        print("I've retrieved:\n{}".format(value))
22        decompressed = zlib.decompress(value).decode()
23        print("But I'm returning:\n{}".format(decompressed))
24        return decompressed
25
26
27 cache = Cache(Cache.REDIS, serializer=CompressionSerializer(), namespace="main")
28
29
30 def serializer():
31     text = (
32         "Lorem ipsum dolor sit amet, consectetur adipiscing elit, sed do eiusmod\u2022"
33         "tempor incididunt\u2022"
34         "ut labore et dolore magna aliqua. Ut enim ad minim veniam, quis nostrud\u2022"
35         "exercitation\u2022"
36         "ullamco laboris nisi ut aliquip ex ea commodo consequat. Duis aute irure\u2022"
37         "dolor in\u2022"
38         "reprehenderit in voluptate velit esse cillum dolore eu fugiat nulla pariatur."
39         "Excepteur\u2022"
40         "sint occaecat cupidatat non proident, sunt in culpa qui officia deserunt\u2022"
41         "mollit\u2022"
42         "anim id est laborum.")
43     cache.set("key", text)

```

(continues on next page)

(continued from previous page)

```

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

```

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

```

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

```

(continues on next page)

(continued from previous page)

```

40
41     @post_load
42     def build_my_type(self, data):
43         return RandomModel(**data)
44
45     class Meta:
46         strict = True
47
48
49 cache = Cache(serializer=MarshmallowSerializer(), namespace="main")
50
51
52 def serializer():
53     model = RandomModel()
54     cache.set("key", model)
55
56     result = cache.get("key")
57
58     assert result.int_type == model.int_type
59     assert result.str_type == model.str_type
60     assert result.dict_type == model.dict_type
61     assert result.list_type == model.list_type
62
63
64 def test_serializer():
65     serializer()
66     cache.delete("key")
67
68
69 if __name__ == "__main__":
70     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 pycached import Cache
>>> from pycached.plugins import TimingPlugin
cache = Cache(plugins=[HitMissRatioPlugin()])
cache.plugins += [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 logging
 2 import random
 3
 4 from pycached import Cache
 5 from pycached.plugins import HitMissRatioPlugin, TimingPlugin, BasePlugin
 6
 7 logger = logging.getLogger(__name__)
 8
 9
10 class MyCustomPlugin(BasePlugin):
11
12     def pre_set(self, *args, **kwargs):
13         logger.info("I'm the pre_set hook being called with %s %s" % (args, kwargs))
14
15     def post_set(self, *args, **kwargs):
16         logger.info("I'm the post_set hook being called with %s %s" % (args, kwargs))
17
18
19 cache = Cache(
20     plugins=[HitMissRatioPlugin(), TimingPlugin(), MyCustomPlugin()],
21     namespace="main")
22
23
24 def run():
25     cache.set("a", "1")
26     cache.set("b", "2")
27     cache.set("c", "3")
28     cache.set("d", "4")
29
30     possible_keys = ["a", "b", "c", "d", "e", "f"]
31
32     for t in range(1000):
33         cache.get(random.choice(possible_keys))
34
35     assert cache.hit_miss_ratio["hit_ratio"] > 0.5
36     assert cache.hit_miss_ratio["total"] == 1000
37
38     assert cache.profiling["get_min"] > 0
39     assert cache.profiling["set_min"] > 0
40     assert cache.profiling["get_max"] > 0
41     assert cache.profiling["set_max"] > 0
42
43     print(cache.hit_miss_ratio)
44     print(cache.profiling)
45
46
47 def test_run():
48     run()
49     cache.delete("a")
50     cache.delete("b")
51     cache.delete("c")
52     cache.delete("d")
53
54 if __name__ == "__main__":
55     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  from pycached import caches, SimpleMemoryCache, RedisCache
2  from pycached.serializers import StringSerializer, PickleSerializer
3
4  caches.set_config({
5      'default': {
6          'cache': "pycached.SimpleMemoryCache",
7          'serializer': {
8              'class': "pycached.serializers.StringSerializer"
9          }
10     },
11     'redis_alt': {
12         'cache': "pycached.RedisCache",
13         'endpoint': "127.0.0.1",
14         'port': 6379,
15         'timeout': 1,
16         'serializer': {
17             'class': "pycached.serializers.PickleSerializer"
18         },
19         'plugins': [
20             {'class': "pycached.plugins.HitMissRatioPlugin"},
21             {'class': "pycached.plugins.TimingPlugin"}
22         ]
23     }
24 })
25
26
27 def default_cache():
28     cache = caches.get('default') # This always returns the same instance
29     cache.set("key", "value")
30
31     assert cache.get("key") == "value"
32     assert isinstance(cache, SimpleMemoryCache)
33     assert isinstance(cache.serializer, StringSerializer)
34
35
36 def alt_cache():
37     # This generates a new instance every time! You can also use `caches.create('alt
38     ↵')`
```

(continues on next page)

(continued from previous page)

```
38     # or even `caches.create('alt', namespace="test", etc...)` to override extra args
39     cache = caches.create(**caches.get_alias_config('redis_alt'))
40     cache.set("key", "value")
41
42     assert cache.get("key") == "value"
43     assert isinstance(cache, RedisCache)
44     assert isinstance(cache.serializer, PickleSerializer)
45     assert len(cache.plugins) == 2
46     assert cache.endpoint == "127.0.0.1"
47     assert cache.timeout == 1
48     assert cache.port == 6379
49     cache.close()
50
51
52 def test_alias():
53     default_cache()
54     alt_cache()
55
56     cache = RedisCache()
57     cache.delete("key")
58     cache.close()
59
60     caches.get('default').close()
61
62
63 if __name__ == "__main__":
64     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

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

3.5.1 cached

```
1
2 from collections import namedtuple
3
4 from pycached import cached, RedisCache
5 from pycached.serializers import PickleSerializer
6
7 Result = namedtuple('Result', 'content, status')
8
9
10 @cached(
11     ttl=10, cache=RedisCache, key="key", serializer=PickleSerializer(), port=6379, ↴
12     namespace="main")
```

(continues on next page)

(continued from previous page)

```

12 def cached_call():
13     return Result("content", 200)
14
15
16 def test_cached():
17     cache = RedisCache(endpoint="127.0.0.1", port=6379, namespace="main")
18     cached_call()
19     assert cache.exists("key") is True
20     cache.delete("key")
21     cache.close()
22
23
24 if __name__ == "__main__":
25     test_cached()

```

3.5.2 multi_cached

```

1 from pycached import multi_cached, RedisCache
2
3 DICT = {
4     'a': "Z",
5     'b': "Y",
6     'c': "X",
7     'd': "W"
8 }
9
10
11 @multi_cached("ids", cache=RedisCache, namespace="main")
12 def multi_cached_ids(ids=None):
13     return {id_: DICT[id_] for id_ in ids}
14
15
16 @multi_cached("keys", cache=RedisCache, namespace="main")
17 def multi_cached_keys(keys=None):
18     return {id_: DICT[id_] for id_ in keys}
19
20
21 cache = RedisCache(endpoint="127.0.0.1", port=6379, namespace="main")
22
23
24 def test_multi_cached():
25     multi_cached_ids(ids=['a', 'b'])
26     multi_cached_ids(ids=['a', 'c'])
27     multi_cached_keys(keys=['d'])
28
29     assert cache.exists('a')
30     assert cache.exists('b')
31     assert cache.exists('c')
32     assert cache.exists('d')
33
34     cache.delete("a")
35     cache.delete("b")
36     cache.delete("c")
37     cache.delete("d")
38     cache.close()

```

(continues on next page)

(continued from previous page)

```
39
40
41 if __name__ == "__main__":
42     test_multi_cached()
```

Warning: null

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

CHAPTER 4

Indices and tables

- genindex
- modindex
- search