
redis-py-cluster Documentation

Release 2.1.3

Johan Andersson

Jun 25, 2021

1	Installation	3
2	Basic usage example	5
3	Library Dependencies	7
4	Supported python versions	9
4.1	Python 2 Compatibility Note	9
5	Regarding duplicate package name on pypi	11
6	The Usage Guide	13
6.1	RedisCluster client configuration options	13
6.2	Implemented redis commands in RedisCluster	15
6.3	Limitations and differences	29
6.4	Pipelines	30
6.5	Pubsub	36
6.6	Readonly mode	37
6.7	Setup client logging	38
6.8	Redis cluster setup	38
6.9	Benchmarks	39
7	The Community Guide	41
7.1	Project status	41
7.2	Testing	41
7.3	Development	42
7.4	Upgrading redis-py-cluster	42
7.5	Release process	45
7.6	Release Notes	46
7.7	Project Authors	52
7.8	Licensing	52
7.9	Disclaimer	53

This project is a port of *redis-rb-cluster* by antirez, with a lot of added functionality.

The original source can be found at <https://github.com/antirez/redis-rb-cluster>.

The source code for this project is [available on github](#).

CHAPTER 1

Installation

Latest stable release from pypi

```
$ pip install redis-py-cluster
```

or from source code

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

Basic usage example

Small sample script that shows how to get started with RedisCluster. It can also be found in the file *examples/basic.py*. Additional code examples of more advance functionality can be found in the *examples/* folder in the source code git repo.

```
>>> from rediscluster import RedisCluster

>>> # Requires at least one node for cluster discovery. Multiple nodes is recommended.
>>> startup_nodes = [{"host": "127.0.0.1", "port": "7000"}]

>>> # Note: See note on Python 3 for decode_responses behaviour
>>> rc = RedisCluster(startup_nodes=startup_nodes, decode_responses=True)

>>> rc.set("foo", "bar")
True
>>> print(rc.get("foo"))
'bar'
```

Note: Python 3

Since Python 3 changed to Unicode strings from Python 2's ASCII, the return type of *most* commands will be binary strings, unless the class is instantiated with the option `decode_responses=True`.

In this case, the responses will be Python 3 strings (Unicode).

For the init argument `decode_responses`, when set to False, redis-py-cluster will not attempt to decode the responses it receives.

In Python 3, this means the responses will be of type *bytes*. In Python 2, they will be native strings (*str*).

If `decode_responses` is set to True, for Python 3 responses will be *str*, for Python 2 they will be *unicode*.

Library Dependencies

Even if the goal is to support all major versions of redis-py in the 3.x.x track, this is not a guarantee that all versions will work.

It is always recommended to use the latest version of the dependencies of this project.

- Redis-py: 'redis>=3.0.0,<4.0.0' is required in this major version of this cluster lib.
- Optional Python: hiredis >= 0.2.0. Older versions might work but is not tested.
- A working Redis cluster based on version >=3.0.0 is required.

Supported python versions

Python versions should follow the same supported python versions as specified by the upstream package *redis-py*, based on what major version(s) that is specified.

If this library supports more than one major version line of *redis-py*, then the supported python versions must include the set of supported python versions by all major version lines.

- 2.7
- 3.5
- 3.6
- 3.7
- 3.8

4.1 Python 2 Compatibility Note

This library follows the announced change from our upstream package *redis-py*. Due to this, we will follow the same python 2.7 deprecation timeline as stated in there.

redis-py-cluster 2.1.x will be the last major version release that supports Python 2.7. The 2.1.x line will continue to get bug fixes and security patches that support Python 2 until August 1, 2020. *redis-py-cluster* 3.0.x will be the next major version and will require Python 3.5+.

Regarding duplicate package name on pypi

It has been found that the python module name that is used in this library (rediscluster) is already shared with a similar but older project.

This lib will *NOT* change the naming of the module to something else to prevent collisions between the libs.

My reasoning for this is the following

- Changing the namespace is a major task and probably should only be done in a complete rewrite of the lib, or if the lib had plans for a version 2.0.0 where this kind of backwards incompatibility could be introduced.
- This project is more up to date, the last merged PR in the other project was 3 years ago.
- This project is aimed for implement support for the cluster support in 3.0+, the other lib do not have that right now, but they implement almost the same cluster solution as the 3.0+ but in much more in the client side.
- The 2 libs is not compatible to be run at the same time even if the name would not collide. It is not recommended to run both in the same python interpreter.

An issue has been raised in each repository to have tracking of the problem.

redis-py-cluster: <https://github.com/Grokzen/redis-py-cluster/issues/150>

rediscluster: <https://github.com/salimane/rediscluster-py/issues/11>

6.1 RedisCluster client configuration options

This chapter will describe all the configuration options and flags that can be sent into the RedisCluster class instance.

Each option will be described in a separate topic to describe how it works and what it does. This will only describe any options that does anything else when compared to the options that redis-py already provides, or new options that is cluster specific. To find out what options redis-py provides please consult the documentation and/or git repo for that project.

6.1.1 Host port remapping

This option exists to enable the client to fix a problem where the redis-server internally tracks a different ip:port compared to what your clients would like to connect to.

A simple example to describe this problem is if you start a redis cluster through docker on your local machine. If we assume that you start the docker image grokzen/redis-cluster, when the redis cluster is initialized it will track the docker network IP for each node in the cluster.

For example this could be 172.18.0.2. The problem is that a client that runs outside on your local machine will receive from the redis cluster that each node is reachable on the ip 172.18.0.2. But in some cases this IP is not available on your host system. To solve this we need a remapping table where we can tell this client that if you get back from your cluster 172.18.0.2 then you should remap it to localhost instead. When the client does this it can now connect and reach all nodes in your cluster.

Remapping works off a rules list. Each rule is a dictionary of the form shown below

Remapping properties:

- This host_port_remap feature will not work on the startup_nodes so you still need to put in a valid and reachable set of startup nodes.
- The remapping logic treats host_port_remap list as a “rules list” and only the first matching remapping entry will be applied

- A remapping rule may contain just host or just port mapping, but both sides of the mapping (i.e. from_host and to_host or from_port and to_port) are required for either
- If both from_host and from_port are specified, then both will be used to decide if a remapping rule applies

Examples of valid rules:

```
{'from_host': "1.2.3.4", 'from_port': 1000, 'to_host': "2.2.2.2", 'to_port': 2000}

{'from_host': "1.1.1.1", 'to_host': "127.0.0.1"}

{'from_port': 1000, 'to_port': 2000}
```

Example scripts:

```
from rediscluster import RedisCluster

startup_nodes = [{"host": "127.0.0.1", "port": "7000"}]

rc = RedisCluster(
    startup_nodes=startup_nodes,
    decode_responses=True,
    host_port_remap=[
        {
            'from_host': '172.18.0.2',
            'from_port': 7000,
            'to_host': 'localhost',
            'to_port': 7000,
        },
        {
            'from_host': '172.22.0.1',
            'from_port': 7000,
            'to_host': 'localhost',
            'to_port': 7000,
        },
    ]
)

## Debug output to show the client config/setup after client has been initialized.
## It should point to localhost:7000 for those nodes.
print(rc.connection_pool.nodes.nodes)

## Test the client that it can still send and receive data from the nodes after the
↪ remap has been done
print(rc.set('foo', 'bar'))
```

This feature is also useful in cases such as when one is trying to access AWS ElastiCache cluster secured by Stunnel (<https://www.stunnel.org/>)

```
from rediscluster import RedisCluster

startup_nodes = [
    {"host": "127.0.0.1", "port": "17000"},
    {"host": "127.0.0.1", "port": "17001"},
    {"host": "127.0.0.1", "port": "17002"},
    {"host": "127.0.0.1", "port": "17003"},
    {"host": "127.0.0.1", "port": "17004"},
    {"host": "127.0.0.1", "port": "17005"}
]
```

(continues on next page)

(continued from previous page)

```

host_port_remap=[
    {'from_host': '41.1.3.1', 'from_port': 6379, 'to_host': '127.0.0.1', 'to_port': ↵
↵17000},
    {'from_host': '41.1.3.5', 'from_port': 6379, 'to_host': '127.0.0.1', 'to_port': ↵
↵17001},
    {'from_host': '41.1.4.2', 'from_port': 6379, 'to_host': '127.0.0.1', 'to_port': ↵
↵17002},
    {'from_host': '50.0.1.7', 'from_port': 6379, 'to_host': '127.0.0.1', 'to_port': ↵
↵17003},
    {'from_host': '50.0.7.3', 'from_port': 6379, 'to_host': '127.0.0.1', 'to_port': ↵
↵17004},
    {'from_host': '32.0.1.1', 'from_port': 6379, 'to_host': '127.0.0.1', 'to_port': ↵
↵17005}
]

# Note: decode_responses must be set to True when used with python3
rc = RedisCluster(
    startup_nodes=startup_nodes,
    host_port_remap=host_port_remap,
    decode_responses=True,
    ssl=True,
    ssl_cert_reqs=None,
    # Needed for Elasticache Clusters
    skip_full_coverage_check=True)

print(rc.connection_pool.nodes.nodes)
print(rc.ping())
print(rc.set('foo', 'bar'))
print(rc.get('foo'))

```

6.2 Implemented redis commands in RedisCluster

This document will enumerate and describe all implemented redis commands and if there is any cluster specific customization/changes done to the command to make them work for a cluster workload.

If a command is specified here but there is no comments on it, then you can assume it will work and behave the same way as when using it from *redis-py*.

If a new command has been added to redis-server and it is not documented here then please open a issue on github telling that it is missing and needs to be added to this documentation.

Danger: If a command below begins with *[NYV]* / *Not Yet Verified* it means that the command is documented here but it is not yet verified that it works or is properly implemented or decided what implementation to use in a clustered environment.

6.2.1 Cluster

<https://redis.io/commands#cluster>

- CLUSTER ADDSLOTS slot [slot ...]

Note: Client has custom implementation where the user has to route the command to the correct node manually.

- [NYV] - CLUSTER BUMPEPOCH
- CLUSTER COUNT_FAILURE-REPORTS node-id

Note: Client has custom implementation where the user has to route the command to the correct node manually.

- CLUSTER COUNTKEYSINSLOT slot

Note: Client will route command to node that owns the slot

- CLUSTER DELSLOTS slot [slot ...]

Note: Client has custom implementation where the user has to route the command to the correct node manually.

- CLUSTER FAILOVER [FORCE|TAKEOVER]

Note: Client has custom implementation where the user has to route the command to the correct node manually.

- [NYV] - CLUSTER FLUSHSLOTS
- [NYV] - CLUSTER FORGET node-id
- CLUSTER GETKEYSINSLOT slot count

Note: Client will route command to node that owns the slot

- CLUSTER INFO

Note: Command is sent to all nodes in the cluster.

Result is merged into a single dict with node as key.

- CLUSTER KEYSLOT key

Note: Client has custom implementation where the user has to route the command to the correct node manually.

- CLUSTER MEET ip port

Note: Client has custom implementation where the user has to route the command to the correct node manually.

- [NYV] - CLUSTER MYID
- CLUSTER NODES

Note: Command will be sent to random node in the cluster as the data should be the same on all nodes in a stable/working cluster

- CLUSTER REPLICATE node-id

Note: Client has custom implementation where the user has to route the command to the correct node manually.

- CLUSTER RESET [HARD|SOFT]

Note: Client has custom implementation where the user has to route the command to the correct node manually.

- CLUSTER SAVECONFIG

Note: Client has custom implementation where the user has to route the command to the correct node manually.

- CLUSTER SET-CONFIG-EPOCH config-epoch

Note: Client has custom implementation where the user has to route the command to the correct node manually.

- CLUSTER SETSLOT slot IMPORTING|MIGRATING|STABLE|NODE [node-id]

Note: Client has custom implementation where the user has to route the command to the correct node manually.

- CLUSTER SLAVES node-id

Note: Client has custom implementation where the user has to route the command to the correct node manually.

- [NYV] - CLUSTER REPLICAS node-id
- CLUSTER SLOTS

Note: Command will be sent to random node in the cluster as the data should be the same on all nodes in a stable/working cluster

- [NYV] - READONLY
- [NYV] - READWRITE

6.2.2 Connection

<https://redis.io/commands#connection>

- [NYV] - AUTH [username] password
- [NYV] - CLIENT CACHING YES/NO
- CLIENT ID

Warning: Command is sent to all nodes in the cluster.

Result from each node will be aggregated into a dict where the key will be the internal node name.

- CLIENT KILL [ip:port] [ID client-id] [TYPE normal|master|slave|pubsub] [USER username] [ADDR ip:port] [SKIPME yes/no]

Warning: Command is sent to all nodes in the cluster.

Result from each node will be aggregated into a dict where the key will be the internal node name.

- CLIENT LIST [TYPE normal|master|replica|pubsub]

Warning: Command is sent to all nodes in the cluster.

Result from each node will be aggregated into a dict where the key will be the internal node name.

- CLIENT GETNAME

Warning: Command is sent to all nodes in the cluster.

Result from each node will be aggregated into a dict where the key will be the internal node name.

- [NYV] - CLIENT GETREDIR
- [NYV] - CLIENT PAUSE timeout
- [NYV] - CLIENT REPLY ON|OFF|SKIP
- [NYV] - CLIENT SETNAME connection-name
- [NYV] - CLIENT TRACKING ON|OFF [REDIRECT client-id] [PREFIX prefix [PREFIX prefix ...]] [BCAST] [OPTIN] [OPTOUT] [NOLOOP]
- [NYV] - CLIENT UNBLOCK client-id [TIMEOUT|ERROR]
- ECHO message

Warning: Command is sent to all nodes in the cluster.

Result from each node will be aggregated into a dict where the key will be the internal node name.

- [NYV] - HELLO protover [AUTH username password] [SETNAME clientname]
- PING [message]

Warning: Command is sent to all nodes in the cluster.

Result from each node will be aggregated into a dict where the key will be the internal node name.

- [NYV] - QUIT
- [NYV] - SELECT index

6.2.3 Geo

<https://redis.io/commands#geo>

- [NYV] - GEOADD key longitude latitude member [longitude latitude member ...]
- [NYV] - GEOHASH key member [member ...]
- [NYV] - GEOPOS key member [member ...]
- [NYV] - GEODIST key member1 member2 [mkmlftlmi]
- [NYV] - GEORADIUS key longitude latitude radius mkmlftlmi [WITHCOORD] [WITHDIST] [WITHHASH] [COUNT count] [ASCDESC] [STORE key] [STOREDIST key]
- [NYV] - GEORADIUSBYMEMBER key member radius mkmlftlmi [WITHCOORD] [WITHDIST] [WITHHASH] [COUNT count] [ASCDESC] [STORE key] [STOREDIST key]

6.2.4 Hashes

<https://redis.io/commands#hash>

- HDEL key field [field ...]
- HEXISTS key field
- HGET key field
- HGETALL key
- HINCRBY key field increment
- HINCRBYFLOAT key field increment
- HKEYS key
- HLEN key
- HMGET key field [field ...]
- HMSET key field value [field value ...]
- HSET key field value [field value ...]
- HSETNX key field value

- HSTRLEN key field
- HVALS key
- HSCAN key cursor [MATCH pattern] [COUNT count]

Note: HSCAN command has currently a buggy client side implementation.

It is not recommended to use any *SCAN methods.

6.2.5 Hyperloglog

<https://redis.io/commands#hyperloglog>

- [NYV] - PFADD key element [element ...]
- [NYV] - PFCOUNT key [key ...]
- [NYV] - PFMERGE destkey sourcekey [sourcekey ...]

6.2.6 Keys/Generic

<https://redis.io/commands#generic>

- DEL key [key ...]

Note: Method has a custom client side implementation.

Command is no longer atomic.

DEL command is sent for each individual key to redis-server.

- DUMP key
- [NYV] - EXISTS key [key ...]
- EXPIRE key seconds
- EXPIREAT key timestamp
- [NYV] - KEYS pattern
- [NYV] - MIGRATE host port key''' destination-db timeout [COPY] [REPLACE] [AUTH password] [AUTH2 username password] [KEYS key [key ...]]
- MOVE key db

Note: Concept of databases do not exists in a cluter

- OBJECT subcommand [arguments [arguments ...]]

Note: Command is blocked from executing in the client.

- PERSIST key
- PEXPIRE key milliseconds

- PEXPIREAT key milliseconds-timestamp
- PTTL key
- RANDOMKEY
- RENAME key newkey

Note: Method has a custom client side implementation.

Command is no longer atomic.

If the slots is the same RENAME will be sent to that shard. If the source and destination keys have different slots then a dump (old key/slot) -> restore (new key/slot) -> delete (old key) will be performed.

- RENAMENX key newkey

Note: Method has a custom client side implementation.

Command is no longer atomic.

Method will check if key exists and if it does it uses the custom RENAME implementation mentioned above.

- [NYV] - RESTORE key ttl serialized-value [REPLACE] [ABSTTL] [IDLETIME seconds] [FREQ frequency]
- SORT key [BY pattern] [LIMIT offset count] [GET pattern [GET pattern ...]] [ASCIDESC] [ALPHA] [STORE destination]

Note: SORT command will only work on the most basic sorting of lists.

Any additional arguments or more complex sorts can't get guaranteed to work if working with cross slots.

Command works if all used keys is in same slot.

- [NYV] - TOUCH key [key ...]
- TTL key
- TYPE key
- [NYV] - UNLINK key [key ...]
- [NYV] - WAIT numreplicas timeout
- [NYV] - SCAN cursor [MATCH pattern] [COUNT count] [TYPE type]

Note: SCAN command has currently a buggy client side implementation.

It is not recommended to use any *SCAN methods.

6.2.7 Lists

<https://redis.io/commands#list>

- [NYV] - BLPOP key [key ...] timeout
- [NYV] - BRPOP key [key ...] timeout

- [NYV] - BRPOLPUSH source destination timeout
- [NYV] - LINDEX key index
- [NYV] - LINSERT key BEFORE|AFTER pivot element
- [NYV] - LLEN key
- [NYV] - LPOP key
- [NYV] - LPOS key element [RANK rank] [COUNT num-matches] [MAXLEN len]
- [NYV] - LPUSH key element [element ...]
- [NYV] - LPUSHX key element [element ...]
- [NYV] - LRANGE key start stop
- [NYV] - LREM key count element
- [NYV] - LSET key index element
- [NYV] - LTRIM key start stop
- [NYV] - RPOP key
- [NYV] - RPOPLPUSH source destination
- [NYV] - RPUSH key element [element ...]
- [NYV] - RPUSHX key element [element ...]

6.2.8 PubSub

<https://redis.io/commands#pubsub>

Warning: All pubsub commands is possible to execute and be routed to correct node when used. But in general pubsub solution should NOT be used inside a clustered environment unless you really know what you are doing. Please read the documentation section about pubsub to get more information about why.

- PSUBSCRIBE pattern [pattern ...]
- PUBSUB subcommand [argument [argument ...]]
- PUBLISH channel message
- PUNSUBSCRIBE [pattern [pattern ...]]
- SUBSCRIBE channel [channel ...]
- UNSUBSCRIBE [channel [channel ...]]

6.2.9 Scripting

<https://redis.io/commands#scripting>

- EVAL script numkeys key [key ...] arg [arg ...]

Warning: Method has a custom client side implementation.

Command will only work if all keys point to the same slot. Otherwise a CROSSSLOT error will be raised.

- SCRIPT DEBUG YES|SYNC|NO

Warning: Command will only be sent to all master nodes in the cluster and result will be aggregated into a dict where the key will be the internal node name.

- SCRIPT EXISTS sha1 [sha1 ...]

Warning: Command will only be sent to all master nodes in the cluster and result will be aggregated into a dict where the key will be the internal node name.

- SCRIPT FLUSH

Warning: Command will only be sent to all master nodes in the cluster and result will be aggregated into a dict where the key will be the internal node name.

- SCRIPT KILL

Warning: Command has been blocked from executing in a cluster environment

- SCRIPT LOAD script

Warning: Command will only be sent to all master nodes in the cluster and result will be aggregated into a dict where the key will be the internal node name.

6.2.10 Server

<https://redis.io/commands#server>

- ACL LOAD

Warning: Command has been blocked from executing in a cluster environment

- ACL SAVE

Warning: Command has been blocked from executing in a cluster environment

- ACL LIST

Warning: Command has been blocked from executing in a cluster environment

- ACL USERS

Warning: Command has been blocked from executing in a cluster environment

- ACL GETUSER username

Warning: Command has been blocked from executing in a cluster environment

- ACL SETUSER username [rule [rule ...]]

Warning: Command has been blocked from executing in a cluster environment

- ACL DELUSER username [username ...]

Warning: Command has been blocked from executing in a cluster environment

- ACL CAT [categoryname]

Warning: Command has been blocked from executing in a cluster environment

- ACL GENPASS [bits]

Warning: Command has been blocked from executing in a cluster environment

- ACL WHOAMI

Warning: Command has been blocked from executing in a cluster environment

- ACL LOG [count or RESET]

Warning: Command has been blocked from executing in a cluster environment

- ACL HELP

Warning: Command has been blocked from executing in a cluster environment

- BGREWRITEAOF

Warning: Command is sent to all nodes in the cluster.
Result from each node will be aggregated into a dict where the key will be the internal node name.

- BGSAVE [SCHEDULE]

Warning: Command is sent to all nodes in the cluster.

Result from each node will be aggregated into a dict where the key will be the internal node name.

- [NYV] - COMMAND
- [NYV] - COMMAND COUNT
- [NYV] - COMMAND GETKEYS
- [NYV] - COMMAND INFO command-name [command-name ...]
- [NYV] - CONFIG GET parameter
- [NYV] - CONFIG REWRITE
- [NYV] - CONFIG SET parameter value
- [NYV] - CONFIG RESETSTAT
- [NYV] - DBSIZE
- [NYV] - DEBUG OBJECT key
- [NYV] - DEBUG SEGFAULT
- [NYV] - FLUSHALL [ASYNC]
- [NYV] - FLUSHDB [ASYNC]
- [NYV] - INFO [section]
- [NYV] - LOLWUT [VERSION version]
- [NYV] - LASTSAVE
- [NYV] - MEMORY DOCTOR
- [NYV] - MEMORY HELP
- [NYV] - MEMORY MALLOC-STATS
- [NYV] - MEMORY PURGE
- [NYV] - MEMORY STATS
- [NYV] - MEMORY USAGE key [SAMPLES count]
- [NYV] - MODULE LIST
- [NYV] - MODULE LOAD path [arg [arg ...]]
- [NYV] - MODULE UNLOAD name
- [NYV] - MONITOR
- [NYV] - ROLE
- [NYV] - SAVE
- [NYV] - SHUTDOWN [NOSAVE|SAVE]
- [NYV] - SLAVEOF host port
- [NYV] - REPLICAOF host port
- [NYV] - SLOWLOG subcommand [argument]
- [NYV] - SWAPDB index1 index2

- [NYV] - SYNC
- [NYV] - PSYNC replicationid offset
- [NYV] - TIME

Note: Command is sent to all nodes in the cluster.

Result is merged into a single dict with node as key.

- [NYV] - LATENCY DOCTOR
- [NYV] - LATENCY GRAPH event
- [NYV] - LATENCY HISTORY event
- [NYV] - LATENCY LATEST
- [NYV] - LATENCY RESET [event [event ...]]
- [NYV] - LATENCY HELP

6.2.11 Sets

<https://redis.io/commands#set>

- [NYV] - SADD key member [member ...]
- [NYV] - SCARD key
- [NYV] - SDIFF key [key ...]
- [NYV] - SDIFFSTORE destination key [key ...]
- [NYV] - SINTER key [key ...]
- [NYV] - SINTERSTORE destination key [key ...]
- [NYV] - SISMEMBER key member
- [NYV] - SMEMBERS key
- [NYV] - SMOVE source destination member
- [NYV] - SPOP key [count]
- [NYV] - SRANDMEMBER key [count]
- [NYV] - SREM key member [member ...]
- [NYV] - SUNION key [key ...]
- [NYV] - SUNIONSTORE destination key [key ...]
- [NYV] - SSCAN key cursor [MATCH pattern] [COUNT count]

6.2.12 Sorted Sets

https://redis.io/commands#sorted_set

- [NYV] - BZPOPMIN key [key ...] timeout
- [NYV] - BZPOPMAX key [key ...] timeout

- [NYV] - ZADD key [NX|XX] [CH] [INCR] score member [score member ...]
- [NYV] - ZCARD key
- [NYV] - ZCOUNT key min max
- [NYV] - ZINCRBY key increment member
- [NYV] - ZINTERSTORE destination numkeys key [key ...] [WEIGHTS weight [weight ...]] [AGGREGATE SUM|MIN|MAX]
- [NYV] - ZLEXCOUNT key min max
- [NYV] - ZPOPMAX key [count]
- [NYV] - ZPOPMIN key [count]
- [NYV] - ZRANGE key start stop [WITHSCORES]
- [NYV] - ZRANGEBYLEX key min max [LIMIT offset count]
- [NYV] - ZREVRANGEBYLEX key max min [LIMIT offset count]
- [NYV] - ZRANGEBYSCORE key min max [WITHSCORES] [LIMIT offset count]
- [NYV] - ZRANK key member
- [NYV] - ZREM key member [member ...]
- [NYV] - ZREMRANGEBYLEX key min max
- [NYV] - ZREMRANGEBYRANK key start stop
- [NYV] - ZREMRANGEBYSCORE key min max
- [NYV] - ZREVRANGE key start stop [WITHSCORES]
- [NYV] - ZREVRANGEBYSCORE key max min [WITHSCORES] [LIMIT offset count]
- [NYV] - ZREVRANK key member
- [NYV] - ZSCORE key member
- [NYV] - ZUNIONSTORE destination numkeys key [key ...] [WEIGHTS weight [weight ...]] [AGGREGATE SUM|MIN|MAX]
- [NYV] - ZSCAN key cursor [MATCH pattern] [COUNT count]

6.2.13 Streams

<https://redis.io/commands#stream>

- [NYV] - XINFO [CONSUMERS key groupname] [GROUPS key] [STREAM key] [HELP]
- [NYV] - XADD key ID field value [field value ...]
- [NYV] - XTRIM key MAXLEN [~] count
- [NYV] - XDEL key ID [ID ...]
- [NYV] - XRANGE key start end [COUNT count]
- [NYV] - XREVRANGE key end start [COUNT count]
- [NYV] - XLEN key
- [NYV] - XREAD [COUNT count] [BLOCK milliseconds] STREAMS key [key ...] id [id ...]

- [NYV] - XGROUP [CREATE key groupname id-or- $\$$] [SETID key groupname id-or- $\$$] [DESTROY key groupname] [DELCONSUMER key groupname consumername]
- [NYV] - XREADGROUP GROUP group consumer [COUNT count] [BLOCK milliseconds] [NOACK] STREAMS key [key ...] ID [ID ...]
- [NYV] - XACK key group ID [ID ...]
- [NYV] - XCLAIM key group consumer min-idle-time ID [ID ...] [IDLE ms] [TIME ms-unix-time] [RETRYCOUNT count] [FORCE] [JUSTID]
- [NYV] - XPENDING key group [start end count] [consumer]

6.2.14 Strings

<https://redis.io/commands#string>

- [NYV] - APPEND key value
- [NYV] - BITCOUNT key [start end]
- [NYV] - BITFIELD key [GET type offset] [SET type offset value] [INCRBY type offset increment] [OVERFLOW WRAPISATIFAIL]
- BITOP operation destkey key [key ...]

Note: Command only works if keys is in same slot. No custom client implementation exists.

- BITPOS key bit [start] [end]
- DECR key
- DECRBY key decrement
- GET key
- GETBIT key offset
- GETRANGE key start end
- GETSET key value
- INCR key
- INCRBY key increment
- INCRBYFLOAT key increment
- [NYV] - MGET key [key ...]
- [NYV] - MSET key value [key value ...]
- [NYV] - MSETNX key value [key value ...]
- [NYV] - PSETEX key milliseconds value
- SET key value [EX seconds|PX milliseconds|KEEPTTL] [NX|XX]
- SETBIT key offset value
- SETEX key seconds value
- SETNX key value
- [NYV] - SETRANGE key offset value

- [NYV] - STRALGO LCS algo-specific-argument [algo-specific-argument ...]
- [NYV] - STRLEN key

6.2.15 Transactions

<https://redis.io/commands#transactions>

- [NYV] - DISCARD
- [NYV] - EXEC
- [NYV] - MULTI
- [NYV] - UNWATCH
- [NYV] - WATCH key [key ...]

6.2.16 Sentinel

<https://redis.io/topics/sentinel>

Sentinel commands is no longer needed or really supported by redis now when cluster solution is in place. All *SENTINEL* commands have been blocked by this client to be executed on any node in the cluster.

- SENTINEL GET-MASTER-ADDR-BY-NAME
- SENTINEL MASTER
- SENTINEL MASTERS
- SENTINEL MONITOR
- SENTINEL REMOVE
- SENTINEL SENTINELS
- SENTINEL SET
- SENTINEL SLAVES

6.3 Limitations and differences

This will compare against *redis-py*

There is a lot of differences that have to be taken into consideration when using redis cluster.

Any method that can operate on multiple keys have to be reimplemented in the client and in some cases that is not possible to do. In general any method that is overridden in RedisCluster have lost the ability of being atomic.

Pipelines do not work the same way in a cluster. In *Redis* it batches all commands so that they can be executed at the same time when requested. But with RedisCluster pipelines will send the command directly to the server when it is called, but it will still store the result internally and return the same data from `.execute()`. This is done so that the code still behaves like a pipeline and no code will break. A better solution will be implemented in the future.

A lot of methods will behave very different when using RedisCluster. Some methods send the same request to all servers and return the result in another format than *Redis* does. Some methods are blocked because they do not work / are not implemented / are dangerous to use in redis cluster.

Some of the commands are only partially supported when using RedisCluster. The commands `zinterstore` and `zunionstore` are only supported if all the keys map to the same key slot in the cluster. This can be achieved by namespacing related keys with a prefix followed by a bracketed common key. Example:

```
r.zunionstore('d{foo}', ['a{foo}', 'b{foo}', 'c{foo}'])
```

This corresponds to how redis behaves in cluster mode. Eventually these commands will likely be more fully supported by implementing the logic in the client library at the expense of atomicity and performance.

6.4 Pipelines

6.4.1 How pipelining works

In `redis-py-cluster`, pipelining is all about trying to achieve greater network efficiency. Transaction support is disabled in `redis-py-cluster`. Use pipelines to avoid extra network round-trips, not to ensure atomicity.

Just like in `redis-py`, `redis-py-cluster` queues up all the commands inside the client until `execute` is called. But, once `execute` is called, `redis-py-cluster` internals work slightly differently. It still packs the commands to efficiently transmit multiple commands across the network. But since different keys may be mapped to different nodes, `redis-py-cluster` must first map each key to the expected node. It then packs all the commands destined for each node in the cluster into its own packed sequence of commands. It uses the `redis-py` library to communicate with each node in the cluster.

Ideally all the commands should be sent to each node in the cluster in parallel so that all the commands can be processed as fast as possible. We do this by first writing all of the commands to the sockets sequentially before reading any of the responses. This allows us to parallelize the network i/o without the overhead of managing python threads.

In previous versions of the library there were some bugs associated with pipelining operations. In an effort to simplify the logic and lessen the likelihood of bugs, if we get back connection errors, `MOVED` errors, `ASK` errors or any other error that can safely be retried, we fall back to sending these remaining commands sequentially to each individual node just as we would in a normal redis call. We still buffer the results inside the pipeline response so there will be no change in client behavior. During normal cluster operations, pipelined commands should work nearly efficiently as pipelined commands to a single instance redis. When there is a disruption to the cluster topography, like when keys are being resharded, or when a slave takes over for a master, there will be a slight loss of network efficiency. Commands that are rejected by the server are tried one at a time as we rebuild the slot mappings. Once the slots table is rebuilt correctly (usually in a second or so), the client resumes efficient networking behavior. We felt it was more important to prioritize correctness of behavior and reliable error handling over networking efficiency for the rare cases where the cluster topography is in flux.

6.4.2 Connection Error handling

The other way pipelines differ in `redis-py-cluster` from `redis-py` is in error handling and retries. With the normal `redis-py` client, if you hit a connection error during a pipeline command it raises the error right there. But we expect `redis-cluster` to be more resilient to failures.

If you hit a connection problem with one of the nodes in the cluster, most likely a stand-by slave will take over for the down master pretty quickly. In this case, we try the commands bound for that particular node to another random node. The other random node will not just blindly accept these commands. It only accepts them if the keys referenced in those commands actually map to that node in the cluster configuration.

Most likely it will respond with a `MOVED` error telling the client the new master for those commands. Our code handles these `MOVED` commands according to the redis cluster specification and re-issues the commands to the correct server transparently inside of `pipeline.execute()` method. You can disable this behavior if you'd like as well.

```
# ASKED and MOVED errors
```

The other tricky part of the redis-cluster specification is that if any command response comes back with an *ASK* or *MOVED* error, the command is to be retried against the specified node.

In previous versions of *redis-py-cluster* treated *ASKED* and *MOVED* errors the same, but they really need to be handled differently. *MOVED* error means that the client can safely update its own representation of the slots table to point to a new node for all future commands bound for that slot.

An *ASK* error means the slot is only partially migrated and that the client can only successfully issue that command to the new server if it prefixes the request with an ‘ASKING’ ‘command first. This lets the new node taking over that slot know that the original server said it was okay to run that command for the given key against the new node even though the slot is not yet completely migrated. Our current implementation now handles this case correctly.

6.4.3 The philosophy on pipelines

After playing around with pipelines and thinking about possible solutions that could be used in a cluster setting this document will describe how pipelines work, strengths and weaknesses of the implementation that was chosen.

Why can't we reuse the pipeline code in *redis-py*? In short it is almost the same reason why code from the normal redis client can't be reused in a cluster environment and that is because of the slots system. Redis cluster consist of a number of slots that is distributed across a number of servers and each key belongs in one of these slots.

In the normal pipeline implementation in *redis-py* we can batch send all the commands and send them to the server at once, thus speeding up the code by not issuing many requests one after another. We can say that we have defined and guaranteed execution order because of this.

One problem that appears when you want to do pipelines in a cluster environment is that you can't have guaranteed execution order in the same way as a single server pipeline. The problem is that because you can queue a command to any key, we will end up in most of the cases having to talk to 2 or more nodes in the cluster to execute the pipeline. The problem with that is that there is no single place/node/way to send the pipeline and redis will sort everything out by itself via some internal mechanisms. Because of that when we build a pipeline for a cluster we have to build several smaller pipelines that we each send to the designated node in the cluster.

When the pipeline is executed in the client each key is checked to what slot it should be sent to and the pipeline is built up based on that information. One thing to note here is that there will be partial correct execution order if you look over the entire cluster because for each pipeline the ordering will be correct. It can also be argued that the correct execution order is applied/valid for each slot in the cluster.

The next thing to take into consideration is what commands should be available and which should be blocked/locked.

In most cases and in almost all solutions multi key commands have to be blocked hard from being executed inside a pipeline. This would only be possible in the case you have a pipeline implementation that always executes immediately each command is queued up. That solution would only give the interface of working like a pipeline to ensure old code will still work, but it would not give any benefits or advantages other than all commands would work and old code would work.

In the solution for this lib multikey commands are blocked hard and will probably not be enabled in pipelines. If you really need to use them you need to execute them through the normal cluster client if they are implemented and work in there. Why can't multi key commands work? In short again it is because the keys can live in different slots on different nodes in the cluster. It is possible in theory to have any command work in a cluster, but only if the keys operated on belong to the same cluster slot. This lib have decided that currently no serious support for that will be attempted.

Examples on commands that do not work is *MGET*, *MSET*, *MOVE*.

One good thing that comes out of blocking multi key commands is that correct execution order is less of a problem and as long as it applies to each slot in the cluster we should be fine.

Consider the following example. Create a pipeline and issue 6 commands *A*, *B*, *C*, *D*, *E*, *F* and then execute it. The pipeline is calculated and 2 sub pipelines is created with *A*, *C*, *D*, *F* in the first and *B*, *E* in the second. Both pipelines

are then sent to each node in the cluster and a response is sent back. For the first node `[True, MovedException(12345), MovedException(12345), True]` and from the second node `[True, True]`. After this response is parsed we see that 2 commands in the first pipeline did not work and must be sent to another node. This case happens if the client slots cache is wrong because a slot was migrated to another node in the cluster. After parsing the response we then build a third pipeline object with commands `[C, D]` to the second node. The third object is executed and passes and from the client perspective the entire pipeline was executed.

If we look back at the order we executed the commands we get `[A, F]` for the first node and `[B, E, C, D]` for the second node. At first glance this looks like it is out of order because command `E` is executed before `C & D`. Why is this not matter? Because no multi key operations can be done in a pipeline, we only have to care the execution order is correct for each slot and in this case it was because `B & E` belongs to the same slot and `C & D` belongs to the same slot. There should be no possible way to corrupt any data between slots if multi key commands are blocked by the code.

What is good with this pipeline solution? First we can actually have a pipeline solution that will work in most cases with few commands blocked (only multi key commands). Secondly we can run it in parallel to increase the performance of the pipeline even further, making the benefits even greater.

6.4.4 Packing Commands

When issuing only a single command, there is only one network round trip to be made. But what if you issue 100 pipelined commands? In a single-instance redis configuration, you still only need to make one network hop. The commands are packed into a single request and the server responds with all the data for those requests in a single response. But with redis cluster, those keys could be spread out over many different nodes.

The client is responsible for figuring out which commands map to which nodes. Let's say for example that your 100 pipelined commands need to route to 3 different nodes? The first thing the client does is break out the commands that go to each node, so it only has 3 network requests to make instead of 100.

6.4.5 Parallel execution of pipeline

In older version of *redis-py-cluster*, there was a thread implementation that helped to increase the performance of running pipelines by running the connections and execution of all commands to all nodes in the pipeline in parallel. This implementation was later removed in favor of a much simpler and faster implementation.

In this new implementation we execute everything in the same thread, but we do all the writing to all sockets in order to each different server and then start to wait for them in sequence until all of them is complete. There is no real need to run them in parallel since we still have to wait for a thread join of all parallel executions before the code can continue, so we can wait in sequence for all of them to complete. This is not the absolute fastest implementation, but it much simpler to implement and maintain and cause less issues because there is no threads or other parallel implementation that will use some overhead and add complexity to the method.

This feature is implemented by default and will be used in all pipeline requests.

6.4.6 Transactions and WATCH

Support for transactions and `WATCH`:es in pipelines. If we look on the entire pipeline across all nodes in the cluster there is no possible way to have a complete transaction across all nodes because if we need to issue commands to 3 servers, each server is handled by its own and there is no way to tell other nodes to abort a transaction if only one of the nodes fail but not the others. A possible solution for that could be to implement a 2 step commit process. The 2 steps would consist of building 2 batches of commands for each node where the first batch would consist of validating the state of each slot that the pipeline wants to operate on. If any of the slots is migrating or moved then the client can correct its slots cache and issue a more correct pipeline batch. The second step would be to issue the actual commands and the data would be committed to redis. The big problem with this is that 99% of the time this would work really well if you have a very stable cluster with no migrations/resharding/servers down. But there can be times where a

slot has begun migration in between the 2 steps of the pipeline and that would cause a race condition where the client thinks it has corrected the pipeline and wants to commit the data but when it does it will still fail.

Why *MULTI/EXEC* support won't work in a cluster environment. There is some test code in the second *MULTI/EXEC cluster test code* of this document that tests if *MULTI/EXEC* is possible to use in a cluster pipeline. The test shows a huge problem when errors occur. If we wrap *MULTI/EXEC* in a packed set of commands then if a slot is migrating we will not get a good error we can parse and use. Currently it will only report *True* or *False* so we can narrow down what command failed but not why it failed. This might work really well if used on a non clustered node because it does not have to take care of *ASK* or *MOVED* errors. But for a cluster we need to know what cluster error occurred so the correct action to fix the problem can be taken. Since there is more than 1 error to take care of it is not possible to take action based on just *True* or *False*.

Because of this problem with error handling *MULTI/EXEC* is blocked hard in the code from being used in a pipeline because the current implementation can't handle the errors.

In theory it could be possible to design a pipeline implementation that can handle this case by trying to determine by itself what it should do with the error by either asking the cluster after a *False* value was found in the response about the current state of the slot or just default to *MOVED* error handling and hope for the best. The problem is that this is not 100% guaranteed to work and can easily cause problems when wrong action was taken on the response.

Currently *WATCH* requires more studying is it possible to use or not, but since it is tied into *MULTI/EXEC* pattern it probably will not be supported for now.

6.4.7 MULTI/EXEC cluster test code

This code does NOT wrap *MULTI/EXEC* around the commands when packed

```
>>> from rediscluster import RedisCluster as s
>>> r = s(startup_nodes=[{"host": "127.0.0.1", "port": "7002"}])
>>> # Simulate that a slot is migrating to another node
>>> r.connection_pool.nodes.slots[14226] = [{
>>>     'host': '127.0.0.1',
>>>     'server_type': 'master',
>>>     'port': 7001,
>>>     'name': '127.0.0.1:7001',
>>> }]
>>> p = r.pipeline()
>>> p.set('ert', 'tre')
>>> p.set('wer', 'rew')
>>> print(p.execute())

ClusterConnection<host=127.0.0.1,port=7001>
[True, ResponseError('MOVED 14226 127.0.0.1:7002',)]
ClusterConnection<host=127.0.0.1,port=7002>
[True]
```

This code DO wrap *MULTI/EXEC* around the commands when packed

```
>>> from rediscluster import RedisCluster as s
>>> r = s(startup_nodes=[{"host": "127.0.0.1", "port": "7002"}])
>>> # Simulate that a slot is migrating to another node
>>> r.connection_pool.nodes.slots[14226] = [{
```

(continues on next page)

(continued from previous page)

```
>>>     'host': '127.0.0.1',
>>>     'server_type': 'master',
>>>     'port': 7001,
>>>     'name': '127.0.0.1:7001',
>>> }]

>>> p = r.pipeline()
>>> p.set('ert', 'tre')
>>> p.set('wer', 'rew')
>>> print(p.execute())

ClusterConnection<host=127.0.0.1,port=7001>
[True, False]
```

6.4.8 Different pipeline solutions

This section will describe different types of pipeline solutions. It will list their main benefits and weaknesses.

Note: This section is mostly random notes and thoughts and not that well written and cleaned up right now. It will be done at some point in the future.

Suggestion one

Simple but yet sequential pipeline. This solution acts more like an interface for the already existing pipeline implementation and only provides a simple backwards compatible interface to ensure that code that exists still will work without any major modifications. This is good because, with this implementation, all commands are run in sequence and it will handle *MOVED* or *ASK* redirections very well and without any problems. The major downside to this solution is that no command is ever batched and run in parallel and thus you do not get any major performance boost from this approach. Another plus is that execution order is preserved across the entire cluster but a major downside is that the commands are no longer atomic on the cluster scale because they are sent in multiple commands to different nodes.

Good

- Sequential execution of the entire pipeline
- Easy *ASK* or *MOVED* handling

Bad

- No batching of commands aka. no execution speedup

Suggestion two

Current pipeline implementation. This implementation is rather good and works well because it combines the existing pipeline interface and functionality and it also provides a basic handling of *ASK* or *MOVED* errors inside the client. One major downside to this is that execution order is not preserved across the cluster. Although the execution order is somewhat broken if you look at the entire cluster level because commands can be split so that cmd1, cmd3, cmd5 get sent to one server and cmd2, cmd4 gets sent to another server. The order is then broken globally but locally for each server it is preserved and maintained correctly. On the other hand I guess that there can't be any commands that can affect different hashslots within the same command so maybe it really doesn't matter if the execution order is not correct because for each slot/key the order is valid. There might be some issues with rebuilding the correct response

ordering from the scattered data because each command might be in different sub pipelines. But I think that our current code still handles this correctly. I think I have to figure out some weird case where the execution order actually matters. There might be some issues with the nonsupported `mget/mset` commands that actually performs different sub commands then it currently supports.

Good

- Sequential execution per node

Bad

- Non sequential execution on the entire pipeline
- Medium difficult *ASK* or *MOVED* handling

Suggestion three

There is a even simpler form of pipelines that can be made where all commands is supported as long as they conform to the same hashslot because REDIS supports that mode of operation. The good thing with this is that since all keys must belong to the same slot there can't be very few *ASK* or *MOVED* errors that happens and if they happen they will be very easy to handle because the entire pipeline is kinda atomic because you talk to the same server and only 1 server. There can't be any multiple server communication happening.

Good

- Super simple *ASK* or *MOVED* handling
- Sequential execution per slot and through the entire pipeline

Bad

- Single slot per pipeline

Suggestion four

One other solution is the 2 step commit solution where you send for each server 2 batches of commands. The first command should somehow establish that each keyslot is in the correct state and able to handle the data. After the client have received OK from all nodes that all data slots is good to use then it will actually send the real pipeline with all data and commands. The big problem with this approach is that there is a gap between the checking of the slots and the actual sending of the data where things can happen to the already established slots setup. But at the same time there is no possibility of merging these 2 steps because if step 2 is automatically run if step 1 is Ok then the pipeline for the first node that will fail will fail but for the other nodes it will succeed but when it should not because if one command gets *ASK* or *MOVED* redirection then all pipeline objects must be rebuilt to match the new specs/setup and then reissued by the client. The major advantage of this solution is that if you have total control of the redis server and do controlled upgrades when no clients is talking to the server then it can actually work really well because there is no possibility that *ASK* or *MOVED* will triggered by migrations in between the 2 batches.

Good

- Still rather safe because of the 2 step commit solution
- Handles *ASK* or *MOVED* before committing the data

Bad

- Big possibility of race conditions that can cause problems

6.5 Pubsub

After testing pubsub in cluster mode one big problem was discovered with the *PUBLISH* command.

According to the current official redis documentation on *PUBLISH*:

```
Integer reply: the number of clients that received the message.
```

It was initially assumed that if we had clients connected to different nodes in the cluster it would still report back the correct number of clients that received the message.

However after some testing of this command it was discovered that it would only report the number of clients that have subscribed on the same server the *PUBLISH* command was executed on.

Because of this, if there is some functionality that relies on an exact and correct number of clients that listen/subscribed to a specific channel it will be broken or behave wrong.

Currently the only known workarounds is to:

- Ignore the returned value
- All clients talk to the same server
- Use a non clustered redis server for pubsub operations

Discussion on this topic can be found here: <https://groups.google.com/forum/?hl=sv#!topic/redis-db/BlwSOYNBUI8>

6.5.1 Scalability issues

The following part is from this discussion https://groups.google.com/forum/?hl=sv#!topic/redis-db/B0_fvfDWLGM and it describes the scalability issue that pubsub has and the performance that goes with it when used in a cluster environment.

according to [1] and [2] PubSub works by broadcasting every publish to every other Redis Cluster node. This limits the PubSub throughput to the bisection bandwidth of the underlying network infrastructure divided by the number of nodes times message size. So if a typical message has 1KB, the cluster has 10 nodes and bandwidth is 1 GBit/s, throughput is already limited to 12.5K RPS. If we increase the message size to 5 KB and the number of nodes to 50, we only get 500 RPS much less than a single Redis instance could service (>100K RPS), while putting maximum pressure on the network. PubSub thus scales linearly wrt. to the cluster size, but in the the negative direction!

6.5.2 How pubsub works in RedisCluster

In release *1.2.0* the pubsub was code was reworked to now work like this.

For *PUBLISH* and *SUBSCRIBE* commands:

- The channel name is hashed and the keyslot is determined.
- Determine the node that handles the keyslot.
- Send the command to the node.

The old solution was that all pubsub connections would talk to the same node all the time. This would ensure that the commands would work.

This new solution is probably future safe and it will probably be a similar solution when *redis* fixes the scalability issues.

6.5.3 Known limitations with pubsub

Pattern subscribe and publish do not work properly because if we hash a pattern like *fo** we will get a keyslot for that string but there is an endless possibility of channel names based on that pattern that we can't know in advance. This feature is not limited but the commands are not recommended to use right now.

The implemented solution will only work if other clients use/adopt the same behaviour. If some other client behaves differently, there might be problems with *PUBLISH* and *SUBSCRIBE* commands behaving wrong.

6.5.4 Other solutions

The simplest solution is to have a separate non clustered redis instance that you have a regular *Redis* instance that works with your pubsub code. It is not recommended to use pubsub until *redis* fixes the implementation in the server itself.

6.6 Readonly mode

By default, Redis Cluster always returns *MOVE* redirection response on accessing slave node. You can overcome this limitation [for scaling read with *READONLY* mode](<http://redis.io/topics/cluster-spec#scaling-reads-using-slave-nodes>).

redis-py-cluster also implements this mode. You can access slave by passing *readonly_mode=True* to *RedisCluster* (or *RedisCluster*) constructor.

```
>>> from rediscluster import RedisCluster
>>> startup_nodes = [{"host": "127.0.0.1", "port": "7000"}]
>>> rc = RedisCluster(startup_nodes=startup_nodes, decode_responses=True)
>>> rc.set("foo16706", "bar")
>>> rc.set("foo81", "foo")
True
>>> rc_readonly = RedisCluster(startup_nodes=startup_nodes, decode_responses=True,
↳readonly_mode=True)
>>> rc_readonly.get("foo16706")
u'bar'
>>> rc_readonly.get("foo81")
u'foo'
```

We can use pipeline via *readonly_mode=True* object.

```
>>> with rc_readonly.pipeline() as readonly_pipe:
...     readonly_pipe.get('foo81')
...     readonly_pipe.get('foo16706')
...     readonly_pipe.execute()
...
[u'foo', u'bar']
```

But this mode has some downside or limitations.

- It is possible that you cannot get the latest data from *READONLY* mode enabled object because Redis implements asynchronous replication.
- **You MUST NOT use SET related operation with *READONLY* mode enabled object**, otherwise you can possibly get 'Too many Cluster redirections' error because we choose master and its slave nodes randomly.
- You should use get related stuff only.

- Ditto with pipeline, otherwise you can get ‘Command # X (XXXX) of pipeline: MOVED’ error.

```
>>> rc_readonly = RedisCluster(startup_nodes=startup_nodes, decode_responses=True,
↳readonly_mode=True)
>>> # NO: This works in almost case, but possibly emits Too many Cluster redirections
↳error...
>>> rc_readonly.set('foo', 'bar')
>>> # OK: You should always use get related stuff...
>>> rc_readonly.get('foo')
```

6.7 Setup client logging

To setup logging for debugging inside the client during development you can add this as an example to your own code to enable *DEBUG* logging when using the library.

```
import logging

from rediscluster import RedisCluster

logging.basicConfig()
logger = logging.getLogger('rediscluster')
logger.setLevel(logging.DEBUG)
logger.propagate = True
```

Note that this logging is not recommended to be used inside production as it can cause a performance drain and a slowdown of your client.

6.8 Redis cluster setup

6.8.1 Manually

- Redis cluster tutorial: <http://redis.io/topics/cluster-tutorial>
- Redis cluster specs: <http://redis.io/topics/cluster-spec>
- This video will describe how to setup and use a redis cluster: <http://vimeo.com/63672368> (This video is outdated but could server as a good tutorial/example)

6.8.2 Docker

A fully functional docker image can be found at <https://github.com/Grokzen/docker-redis-cluster>

See repo *README* for detailed instructions how to setup and run.

6.8.3 Vagrant

A fully functional vagrant box can be found at <https://github.com/72squared/vagrant-redis-cluster>

See repo *README* for detailed instructions how to setup and run.

6.8.4 Simple makefile

A simple makefile solution can be found at <https://github.com/Grokzen/travis-redis-cluster>

See repo *README* for detailed instructions how to setup.

6.9 Benchmarks

These are a few benchmarks that are designed to test specific parts of the code to demonstrate the performance difference between using this lib and the normal Redis client.

6.9.1 Setup benchmarks

Before running any benchmark you should install this lib in editable mode inside a virtualenv so it can import *Redis-Cluster* lib.

Install with

```
pip install -e .
```

You also need a few redis servers to test against. You must have one cluster with at least one node on port 7001 and you must also have a non-clustered server on port 7007.

6.9.2 Implemented benchmarks

- *simple.py*, This benchmark can be used to measure a simple *set* and *get* operation chain. It also supports running pipelines by adding the flag *-pipeline*.

6.9.3 Run predefined benchmarks

These are a set of predefined benchmarks that can be run to measure the performance drop from using this library.

To run the benchmarks run

```
make benchmark
```

Example output and comparison of different runmodes

7.1 Project status

If you have a problem with the code or general questions about this lib, you can ping me inside the gitter channel that you can find here <https://gitter.im/Grokzen/redis-py-cluster> and i will help you out with problems or usage of this lib.

As of release *1.0.0* this project will be considered stable and usable in production. If you are going to use redis cluster in your project, you should read up on all documentation that you can find in the bottom of this Readme file. It will contain usage examples and descriptions of what is and what is not implemented. It will also describe how and why things work the way they do in this client.

On the topic about porting/moving this code into *redis-py* there is currently work over here <https://github.com/andymccurdy/redis-py/pull/604> that will bring cluster support based on this code. But my suggestion is that until that work is completed that you should use this lib.

7.2 Testing

All tests are currently built around a 6 redis server cluster setup (3 masters + 3 slaves). One server must be using port 7000 for redis cluster discovery.

The easiest way to setup a cluster is to use either a Docker or Vagrant. They are both described in [Setup a redis cluster. Manually, Docker & Vagrant](docs/Cluster_Setup.md).

7.2.1 Tox - Multi environment testing

To run all tests in all supported environments with *tox* read this [Tox multienv testing](docs/Tox.md)

Tox is the easiest way to run all tests because it will manage all dependencies and run the correct test command for you.

TravisCI will use tox to run tests on all supported python & hiredis versions.

Install tox with *pip install tox*

To run all environments you need all supported python versions installed on your machine. (See supported python versions list) and you also need the python-dev package for all python versions to build hiredis.

To run a specific python version use either `tox -e py27` or `tox -e py34`

7.3 Development

7.3.1 Documentation

To build and test/view documentation you need to install sphinx and addons to be able to run the local dev server to render the documentation.

Install sphinx plus addons

To start the local development server run from the root folder of this git repo

Open up `localhost:8000` in your web-browser to view the online documentation

7.4 Upgrading redis-py-cluster

This document describes what must be done when upgrading between different versions to ensure that code still works.

7.4.1 2.0.0 → 2.1.0

Python3 version must now be one of 3.5, 3.6, 3.7, 3.8

The following exception example has now a new more specific exception class that will be attempted to be caught and the client to resolve the cluster layout. If enough attempts has been made then `SlotNotCoveredError` will be raised with the same message as before. If you have catch for `RedisClusterException` you either remove it and let the client try to resolve the cluster layout itself, or start to catch `SlotNotCoveredError`. This error usually happens during failover if you run `skip_full_coverage_check=True` when running on AWS ElasticCache for example.

```
## Example exception rediscluster.exceptions.RedisClusterException: Slot "6986" not covered by the
cluster. "skip_full_coverage_check=True"
```

7.4.2 1.3.x → 2.0.0

Redis-py upstream package dependency has now been updated to be any of the releases in the major version line 3.0.x. This means that you must upgrade your dependency from 2.10.6 to the latest version. Several internal components have been updated to reflect the code from 3.0.x.

Class `StrictRedisCluster` was renamed to `RedisCluster`. All usages of this class must be updated.

Class `StrictRedis` has been removed to mirror upstream class structure.

Class `StrictClusterPipeline` was renamed to `ClusterPipeline`.

Method `SORT` has been changed back to only allow execution if keys are in the same slot. No more client side parsing and handling of the keys and values.

7.4.3 1.3.2 → Next Release

If you created the *StrictRedisCluster* (or *RedisCluster*) instance via the *from_url* method and were passing *readonly_mode* to it, the connection pool created will now properly allow selecting read-only slaves from the pool. Previously it always used master nodes only, even in the case of *readonly_mode=True*. Make sure your code don't attempt any write commands over connections with *readonly_mode=True*.

7.4.4 1.3.1 → 1.3.2

If your redis instance is configured to not have the *CONFIG ...* commands enabled due to security reasons you need to pass this into the client object *skip_full_coverage_check=True*. Benefits are that the client class no longer requires the *CONFIG ...* commands to be enabled on the server. A downside is that you can't use the option in your redis server and still use the same feature in this client.

7.4.5 1.3.0 → 1.3.1

Method *scan_iter* was rebuilt because it was broken and did not perform as expected. If you are using this method you should be careful with this new implementation and test it through before using it. The expanded testing for that method indicates it should work without problems. If you find any issues with the new method please open a issue on github.

A major refactoring was performed in the pipeline system that improved error handling and reliability of execution. It also simplified the code, making it easier to understand and to continue development in the future. Because of this major refactoring you should thoroughly test your pipeline code to ensure that none of your code is broken.

7.4.6 1.2.0 → Next release

Class *RedisClusterMgt* has been removed. You should use the *CLUSTER ...* methods that exist in the *StrictRedisCluster* client class.

Method *cluster_delslots* changed argument specification from *self, node_id, *slots* to *self, *slots* and changed the behaviour of the method to now automatically determine the *slot_id* based on the current cluster structure and where each slot that you want to delete is loaded.

Method *pfcount* no longer has custom logic and exceptions to prevent CROSSSLOT errors. If method is used with different slots then a regular CROSSSLOT error (*rediscluster.exceptions.ClusterCrossSlotError*) will be returned.

7.4.7 1.1.0 → 1.2.0

Discontinue passing *pipeline_use_threads* flag to *rediscluster.StrictRedisCluster* or *rediscluster.RedisCluster*.

Also discontinue passing *use_threads* flag to the *pipeline()* method.

In 1.1.0 and prior, you could use *pipeline_use_threads* flag to tell the client to perform queries to the different nodes in parallel via threads. We exposed this as a flag because using threads might have been risky and we wanted people to be able to disable it if needed.

With this release we figured out how parallelize commands without the need for threads. We write to all the nodes before reading from them, essentially multiplexing the connections (but without the need for complicated socket multiplexing). We found this approach to be faster and more scalable as more nodes are added to the cluster.

That means we don't need the *pipeline_use_threads* flag anymore, or the *use_threads* flag that could be passed into the instantiation of the pipeline object itself.

The logic is greatly simplified and the default behavior will now come with a performance boost and no need to use threads.

Publish and subscribe no longer connects to a single instance. It now hashes the channel name and uses that to determine what node to connect to. More work will be done in the future when *redis-server* improves the pubsub implementation. Please read up on the documentation about pubsub in the *docs/pubsub.md* file about the problems and limitations on using a pubsub in a cluster.

Commands Publish and Subscribe now uses the same connections as any other commands. If you are using any pubsub commands you need to test it through thoroughly to ensure that your implementation still works.

To use less strict cluster slots discovery you can add the following config to your redis-server config file “cluster-require-full-coverage=no” and this client will honour that setting and not fail if not all slots is covered.

A bug was fixed in ‘sdiffstore’, if you are using this, verify that your code still works as expected.

Class RedisClusterMgt is now deprecated and will be removed in next release in favor of all cluster commands implemented in the client in this release.

7.4.8 1.0.0 → 1.1.0

The following exceptions have been changed/added and code that use this client might have to be updated to handle the new classes.

raise RedisClusterException(“Too many Cluster redirections”) have been changed to raise ClusterError(‘TTL exhausted.’)

ClusterDownException have been replaced with *ClusterDownError*

Added new *AskError* exception class.

Added new *TryAgainError* exception class.

Added new *MovedError* exception class.

Added new *ClusterCrossSlotError* exception class.

Added optional *max_connections_per_node* parameter to *ClusterConnectionPool* which changes behavior of *max_connections* so that it applies per-node rather than across the whole cluster. The new feature is opt-in, and the existing default behavior is unchanged. Users are recommended to opt-in as the feature fixes two important problems. First is that some nodes could be starved for connections after *max_connections* is used up by connecting to other nodes. Second is that the asymmetric number of connections across nodes makes it challenging to configure file descriptor and redis max client settings.

Reinitialize on *MOVED* errors will not run on every error but instead on every 25 error to avoid excessive cluster reinitialize when used in multiple threads and resharding at the same time. If you want to go back to the old behaviour with reinitialize on every error you should pass in *reinitialize_steps=1* to the client constructor. If you want to increase or decrease the interval of this new behaviour you should set *reinitialize_steps* in the client constructor to a value that you want.

Pipelines in general have received a lot of attention so if you are using pipelines in your code, ensure that you test the new code out a lot before using it to make sure it still works as you expect.

The entire client code should now be safer to use in a threaded environment. Some race conditions was found and have now been fixed and it should prevent the code from behaving weird during reshard operations.

7.4.9 0.2.0 → 0.3.0

In *0.3.0* release the name of the client class was changed from *RedisCluster* to *StrictRedisCluster* and a new implementation of *RedisCluster* was added that is based on *redis.Redis* class. This was done to enable implementation a

cluster enabled version of *redis.Redis* class.

Because of this all imports and usage of *RedisCluster* must be changed to *StrictRedisCluster* so that existing code will remain working. If this is not done some issues could arise in existing code.

7.4.10 0.1.0 → 0.2.0

No major changes was done.

7.5 Release process

This section describes the process and how a release is made of this package.

All steps for twine tool can be found here <https://twine.readthedocs.io/en/latest/>

7.5.1 Install helper tools

We use the standard sdist package build solution to package the source dist and wheel package into the format that pip and pypi understands.

We then use *twine* as the helper tool to upload and interact with pypi to submit the package to both pypi & testpypi.

First create a new venv that uses at least python3.7 but it is recommended to use the latest python version always. Published releases will be built with python 3.9.0+

Install twine with

7.5.2 Build python package

First ensure that your *dist/* folder is empty so that you will not attempt to upload a dev version or other packages to the public index.

Create the source dist and wheel dist by running

The built python packages can be found in 'dist'

7.5.3 Submit to testpypi

It is always good to test out the build first locally so there are no obvious code problems but also to submit the build to testpypi to verify that the upload works and that you get the version number and *README* section working correct.

To upload to *testpypi* run

It will upload everything to <https://test.pypi.org/project/redis-py-cluster/>

7.5.4 Submit build to public pypi

To submit the final package to public official pypi run

7.6 Release Notes

7.6.1 2.1.3 (May 30 2021)

- Add example script `pipelin-readonly-replica.py` to show how to use replica nodes to offload read commands from primary node
- `max_connection` now defaults to 50 in `ClusterBlockingConnectionPool` to avoid issue with infinite loop in queue mechanism
- Using read replica for read commands inside pipeline is now better supported. Feature might be unstable to use as own risk.
- Fixed that in some cases where `ConnectionError` is raised, a non existing connection was attempted to be disconnected and caused a sub exception to be raised.

7.6.2 2.1.2 (Apr 18 2021)

- Fixed bug where `“from rediscluster import *”` would not work correct

7.6.3 2.1.1 (Apr 18 2021)

- `ClusterPipeline` is now exposed when doing `“from rediscluster import *”`
- Fix issue where connection would be `None` in some cases when connection pool fails to initialize
- Ported in a fix from `redis-py` where it now checks if a connection is ready or not before returning the connection for usage
- `ClusterFailover` command option is no longer mandatory but optional as it is intended
- Fixed `“SLOWLOG GET”` kwarg command where it failed on `decode_responses`
- `BaseException` is now caught when executing commands and it will disconnect and the connection before raising the exception.
- Logging exception on `ReseponseError` when doing the initial connection to the `startup_nodes` instances

7.6.4 2.1.0 (Sept 26, 2020)

- Add new config option for `Client` and `Pipeline` classes to controll how many attempts will be made before bailing out from a `ClusterDownError`. Use `“cluster_down_retry_attempts=<int>”` when creating the client class to controll this behaviour.
- Updated `redis-py` compatbile version to support any version in the major version 3.0.x, 3.1.x, 3.2.x, 3.3.x., 3.4.x, 3.5.x (#326) It is always recommended to use the latest version of `redis-py` to avoid issues and compatibility problems.
- Fixed bug preventing reinitialization after getting `MOVED` errors
- Add testing of `redis-esrver 6.0` versions to travis and unit tests
- Add python 2.7 compatiblity note about deprecation and upcomming changes in python 2.7 support for this lib
- Updated tests and cluster tests versions of the same methods to latest tests from upstream `redis-py` package
- Reorganized tests and how cluster specific tests is written and run over the upstream version of the same test to make it easier and much faster to update and keep them in sync over time going into the future (#368)

- Python 3.5.x or higher is now required if running on a python 3 version
- Removed the monkeypatching of RedisCluster, ClusterPubSub & ClusterPipeline class names into the “redis” python package namespace during runtime. They are now exposed in the “rediscluster” namespace to mimic the same feature from redis-py
- cluster_down_retry_attempts can now be configured to any value when creating RedisCluster instance
- Creating RedisCluster from unix socket url:s has been disabled
- Patch the from_url method to use the corret cluster version of the same Connection class
- ConnectionError and TimeoutError is now handled seperately in the main execute loop to better handle each case (#363)
- Update scan_iter custom cluster implementation
- Improve description_format handling for connection classes to simplify how they work
- Implement new connection pool ClusterBlockingConnectionPool (#347)
- Nodemanager initialiize should now handle usernames properly (#365)
- PubSub tests has been all been disabled
- New feature, host_port_remap. Send in a remapping configuration to RedisCluster instance where the nodes configuration recieved from the redis cluster can be altered to allow for connection in certain circumstances. See new section in client.rst in docs/ for usage example.
- When a slot is not covered by the cluster, it will not raise SlotNotCoveredError instead of the old generic RedisClusterException. The client will not attempt to rebuild the cluster layout a few times before giving up and raising that exception to the user. (#350)
- CLIENT SETNAME is now possible to use from the client instance. For setting the name for all connections from the client by default, see issue #802 in redis-py repo for the change that was implemented in redis-py 3.4.0.
- Rewrote implemented commands documentation to mimic the redis.io commands documentation and describe each command and any additional implementation that has been made.
- Added RTD theme to the rendered output when running the documentation in local dev mode.
- Added some basic logging to the client that should make it easier to debug and track down minor issues around the main execution loop. See docs/logging.rst for implementation example into your own code.
- Seperated some of the exception handling inside the main execution loop to get more fine grained controll what to do at certain errors.

7.6.5 2.0.0 (Aug 12, 2019)

Specific changes to redis-py-cluster is mentioned below here.

- Update entire code base to now support all redis-py version in the 3.0.x version line. Any future redis-py version will be supported at a later time.
- Major update to all tests to mirror the code of the same tests from redis-py
- Dropped support for the 2.10.6 redis-py release.
- Add pythoncodestyle lint validation check to travis-ci runs to check for proper linting before accepting PR:s
- Class StrictRedisCluster was renamed to RedisCluster
- Class StrictRedis has been removed to mirror upstream class structure
- Class StrictClusterPipeline was renamed to ClusterPipeline

- Fixed travis-ci tests not running properly on python 3.7
- Fixed documentation regarding threads in pipelines
- Update list of command callbacks and parsers. Added in “CLIENT ID”
- Removed custom implementation of SORT and revert back to use same-slot mechanism for that command.
- Added better exception message to `get_master_node_by_slot` command to help the user understand the error.
- Improved the exception object message parsing when running on python3

7.6.6 1.3.6 (Nov 16, 2018)

- Pin upstream redis-py package to release 2.10.6 to avoid issues with incompatible version 3.0.0

7.6.7 1.3.5 (July 22, 2018)

- Add Redis 4 compatibility fix to CLUSTER NODES command (See issue #217)
- Fixed bug with command “CLUSTER GETKEYSINSLOT” that was throwing exceptions
- Added new methods `cluster_get_keys_in_slot()` to client
- Fixed bug with `StrictRedisCluster.from_url` that was ignoring the `readonly_mode` parameter
- NodeManager will now ignore nodes showing cluster errors when initializing the cluster
- Fix bug where RedisCluster wouldn’t refresh the cluster table when executing commands on specific nodes
- Add redis 5.0 to travis-ci tests
- Change default redis version from 3.0.7 to 4.0.10
- Increase accepted ranges of dependencies specified in dev-requirements.txt
- Several major and minor documentation updates and tweaks
- Add example script “`from_url_password_protected.py`”
- command “CLUSTER GETKEYSINSLOT” is now returned as a list and not int
- Improve support for ssl connections
- Retry on Timeout errors when doing cluster discovery
- Added new error class “MasterDownError”
- Updated requirements for dependency of redis-py to latest version

7.6.8 1.3.4 (Mar 5, 2017)

- Package is now built as a wheel and source package when releases is built.
- Fixed issues with some key types in `NodeManager.keyslot()`.
- Add support for `PUBSUB` subcommands `CHANNELS`, `NUMSUB [arg] [args...]` and `NUMPAT`.
- Add method `set_result_callback(command, callback)` allowing the default reply callbacks to be changed, in the same way `set_response_callback(command, callback)` inherited from Redis-Py does for responses.
- Node manager now honors defined `max_connections` variable so connections that is emitted from that class uses the same variable.

- Fixed a bug in cluster detection when running on python 3.x and `decode_responses=False` was used. Data back from redis for cluster structure is now converted no matter what the data you want to set/get later is using.
- Add `SSLClusterConnection` for connecting over TLS/SSL to Redis Cluster
- Add new option to make the nodemanager to follow the cluster when nodes move around by avoiding to query the original list of startup nodes that was provided when the client object was first created. This could make the client handle drifting clusters on for example AWS easier but there is a higher risk of the client talking to the wrong group of nodes during split-brain event if the cluster is not consistent. This feature is `EXPERIMENTAL` and use it with care.

7.6.9 1.3.3 (Dec 15, 2016)

- Remove print statement that was faulty committed into release 1.3.2 that case logs to fill up with unwanted data.

7.6.10 1.3.2 (Nov 27, 2016)

- Fix a bug where `from_url` was not possible to use without passing in additional variables. Now it works as the same method from `redis-py`. Note that the same rules that is currently in place for passing ip addresses/dns names into `startup_nodes` variable apply the same way through the `from_url` method.
- Added options to skip full coverage check. This flag is useful when the `CONFIG redis` command is disabled by the server.
- Fixed a bug where method `CLUSTER SLOTS` would break in newer redis versions where node id is included in the reponse. Method is not compatible with both old and new redis versions.

7.6.11 1.3.1 (Oct 13, 2016)

- Rebuilt broken method `scan_iter`. Previous tests was to small to detect the problem but is not corrected to work on a bigger dataset during the test of that method. (korvus81, Grokzen, RedWhiteMiko)
- Errors in pipeline that should be retried, like connection errors, moved, errors and ask errors now fall back to single operation logic in `StrictRedisCluster.execute_command`. (72squared).
- Moved `reinitialize_steps` and counter into nodemanager so it can be correctly counted across pipeline operations (72squared).

7.6.12 1.3.0 (Sep 11, 2016)

- Removed `RedisClusterMgt` class and file
- Fixed a bug when using pipelines with `RedisCluster` class (Ozahata)
- Bump `redis-server` during travis tests to 3.0.7
- Added docs about same module name in another python redis cluster project.
- Fix a bug when a connection was to be tracked for a node but the node either do not yet exists or was removed because of resharding was done in another thread. (ashishbaghudana)
- Fixed a bug with “`CLUSTER ...`” commands when a `node_id` argument was needed and the return type was supposed to be converted to bool with `bool_ok` in `redis_compat`.
- Add back gitter chat room link
- Add new client commands - `cluster_reset_all_nodes`

- Command `cluster_delslots` now determines what cluster shard each slot is on and sends each slot deletion command to the correct node. Command have changed argument spec (Read `Upgrading.rst` for details)
- Fixed a bug when hashing the key if it was a python 3 byte string and it would cause it to route to wrong slot in the cluster (fossilet, Grokzen)
- Fixed a bug when reinitialize the nodemanager it would use the old `nodes_cache` instead of the new one that was just parsed (monklof)

7.6.13 1.2.0 (Apr 09, 2016)

- Drop maintained support for python 3.2.
- Remove Vagrant file in favor for repo maintained by 72squared
- Add Support for password protected cluster (etng)
- Removed assertion from code (gmolight)
- Fixed a bug where a regular connection pool was allocated with each `StrictRedisCluster` instance.
- Rework `pfcount` to now work as expected when all arguments points to same hashslot
- New code and important changes from `redis-py 2.10.5` have been added to the codebase.
- Removed the need for threads inside of pipeline. We write the packed commands all nodes before reading the responses which gives us even better performance than threads, especially as we add more nodes to the cluster.
- Allow passing in a custom connection pool
- Provide default `max_connections` value for `ClusterConnectionPool` (`2**31`)
- Travis now tests both `redis 3.0.x` and `3.2.x`
- Add simple `ptpdb` debug script to make it easier to test the client
- Fix a bug in `sdiffstore` (mt3925)
- Fix a bug with `scan_iter` where duplicate keys would be returned during iteration
- Implement all “`CLUSTER ...`” commands as methods in the client class
- Client now follows the service side setting ‘`cluster-require-full-coverage=yes/no`’ (baranbartu)
- Change the pubsub implementation (`PUBLISH/SUBSCRIBE` commands) from using one single node to now determine the hashslot for the channel name and use that to connect to a node in the cluster. Other clients that do not use this pattern will not be fully compatible with this client. Known limitations is pattern subscription that do not work properly because a pattern can’t know all the possible channel names in advance.
- Convert all docs to `ReadTheDocs`
- Rework connection pool logic to be more similar to `redis-py`. This also fixes an issue with pubsub and that connections was never release back to the pool of available connections.

7.6.14 1.1.0 (Oct 27, 2015)

- Refactored exception handling and exception classes.
- Added `READONLY` mode support, scales reads using slave nodes.
- Fix `__repr__` for `ClusterConnectionPool` and `ClusterReadOnlyConnectionPool`
- Add `max_connections_per_node` parameter to `ClusterConnectionPool` so that `max_connections` parameter is calculated per-node rather than across the whole cluster.

- Improve thread safety of `get_connection_by_slot` and `get_connection_by_node` methods (iandyh)
- Improved error handling when sending commands to all nodes, e.g. `info`. Now the connection takes `retry_on_timeout` as an option and retry once when there is a timeout. (iandyh)
- Added support for `SCRIPT LOAD`, `SCRIPT FLUSH`, `SCRIPT EXISTS` and `EVALSHA` commands. (alisaifee)
- Improve thread safety to avoid exceptions when running one client object inside multiple threads and doing resharding of the cluster at the same time.
- Fix `ASKING` error handling so now it really sends `ASKING` to next node during a reshard operation. This improvement was also made to pipelined commands.
- Improved thread safety in pipelined commands, along better explanation of the logic inside pipelining with code comments.

7.6.15 1.0.0 (Jun 10, 2015)

- No change to anything just a bump to 1.0.0 because the lib is now considered stable/production ready.

7.6.16 0.3.0 (Jun 9, 2015)

- simple benchmark now uses `docopt` for cli parsing
- New make target to run some benchmarks ‘make benchmark’
- simple benchmark now support pipelines tests
- Renamed `RedisCluster` → `StrictRedisCluster`
- Implement backwards compatible `redis.Redis` class in cluster mode. It was named `RedisCluster` and everyone updating from 0.2.0 to 0.3.0 should consult `docs/Upgrading.md` for instructions how to change your code.
- Added comprehensive documentation regarding pipelines
- Meta retrieval commands(`slots`, `nodes`, `info`) for Redis Cluster. (iandyh)

7.6.17 0.2.0 (Dec 26, 2014)

- Moved pipeline code into new file.
- Code now uses a proper cluster connection pool class that handles all nodes and connections similar to how `redis-py` do.
- Better support for `pubsub`. All clients will now talk to the same server because `pubsub` commands do not work reliably if it talks to a random server in the cluster.
- Better result callbacks and node routing support. No more ugly decorators.
- Fix `keyslot` command when using non `ascii` characters.
- Add `bitpos` support, `redis-py 2.10.2` or higher required.
- Fixed a bug where `vagrant` users could not build the package via shared folder.
- Better support for `CLUSTERDOWN` error. (Neuront)
- Parallel pipeline execution using threads. (72squared)
- Added `vagrant` support for testing and development. (72squared)
- Improve stability of client during resharding operations (72squared)

7.6.18 0.1.0 (Sep 29, 2014)

- Initial release
- First release uploaded to pypi

7.7 Project Authors

Added in the order they contributed.

If you are mentioned in this document and want your row changed for any reason, open a new PR with changes.

Lead author and maintainer: Grokzen - <https://github.com/Grokzen>

Authors who contributed code or testing:

- Dobrite - <https://github.com/dobrite>
- 72squared - <https://github.com/72squared>
- Neuron Teckid - <https://github.com/neuront>
- iandyh - <https://github.com/iandyh>
- mumumu - <https://github.com/mumumu>
- awestendorf - <https://github.com/awestendorf>
- Ali-Akber Saifee - <https://github.com/alisaifee>
- etng - <https://github.com/etng>
- gmolight - <https://github.com/gmolight>
- baranbartu - <https://github.com/baranbartu>
- monklof - <https://github.com/monklof>
- dutradda - <https://github.com/dutradda>
- AngusP - <https://github.com/AngusP>
- Doug Kent - <https://github.com/dkent>
- VascoVisser - <https://github.com/VascoVisser>
- astrohsy - <https://github.com/astrohsy>
- Artur Stawiarski - <https://github.com/astawiarski>
- Matthew Anderson - <https://github.com/mc3ander>
- Appurv Jain - <https://github.com/appurvj>

7.8 Licensing

Copyright (c) 2013-2021 Johan Andersson

MIT (See docs/License.txt file)

The license should be the same as redis-py (<https://github.com/andymccurdy/redis-py>)

7.9 Disclaimer

Both Redis cluster and redis-py-cluster is considered stable and production ready.

But this depends on what you are going to use clustering for. In the simple use cases with SET/GET and other single key functions there is not issues. If you require multi key functionality or pipelines then you must be very careful when developing because they work slightly different from the normal redis server.

If you require advance features like pubsub or scripting, this lib and redis do not handle that kind of use-cases very well. You either need to develop a custom solution yourself or use a non clustered redis server for that.

Finally, this lib itself is very stable and I know of at least 2 companies that use this in production with high loads and big cluster sizes.