PyRMQ

Alexandre Gerona & Jasper Sibayan

Dec 02, 2021

CONTENTS

Python with RabbitMQ—simplified so you won't have to.

FEATURES

Stop worrying about boilerplating and implementing retry logic on your queues. PyRMQ already does it for you.

- Use out-of-the-box *Consumer* and *Publisher* classes created from pika for your projects and tests.

- Custom DLX-DLK-based retry logic for message consumption.

- Message priorities

- Works with Python 3.

- Production ready

QUICKSTART

PyRMQ is available at PyPI.

```
$ pip install pyrmq
```

Just instantiate the feature you want with their respective settings. PyRMQ already works out of the box with RabbitMQ's default initialization settings.

```python
from pyrmq import Publisher
publisher = Publisher(
    exchange_name="exchange_name",
    queue_name="queue_name",
    routing_key="routing_key",
)
publisher.publish({"pyrmq": "My first message"})
```

PUBLISH MESSAGE WITH PRIORITIES

To enable prioritization of messages, instantiate your queue with the queue argument *x-max-priority*. It takes an integer that sets the number of possible priority values with a higher number commanding more priority. Then, simply publish your message with the priority argument specified. Any number higher than the set max priority is floored or considered the same. Read more about message priorities here

```python
from pyrmq import Publisher
publisher = Publisher(
    exchange_name="exchange_name",
    queue_name="queue_name",
    routing_key="routing_key",
    queue_args={"x-max-priority": 3}
)
publisher.publish({"pyrmq": "My first message"}, priority=1)
```

Warning: Adding arguments on an existing queue is not possible. If you wish to add queue arguments, you will need to either delete the existing queue then recreate the queue with arguments or simply make a new queue with the arguments.

FOUR

CONSUMING

Instantiating a *Consumer* automatically starts it in its own thread making it non-blocking by default. When run after the code from before, you should be able to receive the published data.

```python
from pyrmq import Consumer


def callback(data):
    print(f"Received {data}!")

consumer = Consumer(
    exchange_name="exchange_name",
    queue_name="queue_name",
    routing_key="routing_key",
)

consumer.start()
```

FIVE

DLX-DLK RETRY LOGIC

What if you wanted to retry a failure on a consumed message? PyRMQ offers a custom solution that keeps your message in queues while retrying periodically for a set amount of times.

This approach uses `dead letter exchanges and queues`_ to republish a message to your original queue once it has expired. PyRMQ creates this "retry" queue for you with the default naming convention of appending your original queue with *.retry*.

```python
from pyrmq import Consumer

def callback(data):
    print(f"Received {data}!")
    raise Exception

consumer = Consumer(
    exchange_name="exchange_name",
    queue_name="queue_name",
    routing_key="routing_key",
    callback=callback,
    is_dlk_retry_enabled=True,
)
consumer.start()
```

This will start a loop of passing your message between the original queue and the retry queue until it reaches the default number of `max_retries`.

DLX-DLK RETRY BACKOFF VS PERIODIC RETRIES

Since RabbitMQ does not remove expired messages that aren't at the head of the queue, this leads to a congestion of the retry queue that is bottlenecked with an unexpired message at the head. As such, as of 3.3.0, PyRMQ will be using a simple periodic retry.

USING OTHER EXCHANGE TYPES

You can use another exchange type just by simply specifying it in the Publisher class. The default is `direct`.

This is an example of how to publish to a headers exchange that will get routed based on its headers.

7.1 Binding an exchange to another exchange

By default, the `exchange_name` you pass when initializing a `Consumer` is declared and bound to the passed `queue_name`. What if you want to bind and declare this exchange to another exchange as well?

This is done by using `bound_exchange`. This parameter accepts an object with two keys: `name` of your exchange and its `type`. Let's take a look at an example to see this in action.

In the example above, we want to consume from an exchange called `direct_exchange` that is directly bound to queue `direct_queue`. We want *direct_exchange* to get its messages from another exchange called `headers_exchange_name` of type `headers`. By using `bound_exchange`, PyRMQ declares `direct_exchange` and `direct_queue` along with any queue or exchange arguments you may have first then declares the bound exchange next and binds them together. This is done to alleviate the need to declare your bound exchange manually.

> **Warning:** Since this method uses e2e bindings, if you're using a headers exchange to bind your consumer to, they _and_ your publisher must all have the same routing key to route the messages properly. This is not needed for exchange to queue bindings as the routing key is optional for those.

EIGHT

USER GUIDE

8.1 PyRMQ Installation

There are multiple ways to install PyRMQ as long as multiple versions to choose from.

8.1.1 Stable Version

PyRMQ is available at PyPI.

```
$ pip install pyrmq
```

8.1.2 Development Version

Since PyRMQ is continuously used in a growing number of internal microservices all working with RabbitMQ, you can see or participate in its active development in its GitHub repository.

There are two ways to work or collaborate with its development version.

Git Checkout

Clone the code from GitHub and run it in a *virtualenv*.

```
$ git clone git@github.com:first-digital-finance/pyrmq.git
$ virtualenv venv --distribute
$ . venv/bin/activate
$ python setup.py install
```

This will setup PyRMQ and its dependencies on your local machine. Just fetch/pull code from the master branch to keep your copy up to date.

```
$ mkdir pyrmq
$ cd pyrmq
$ virtualenv venv --distribute
$ . venv/bin/activate
$ pip install git+git://github.com/first-digital-finance/pyrmq.git
```

8.2 How to use PyRMQ

8.2.1 Publishing

Instantiate the *Publisher* class and plug in your application specific settings. PyRMQ already works out of the box with RabbitMQ's default initialization settings.

```python
from pyrmq import Publisher
publisher = Publisher(
    exchange_name="exchange_name",
    queue_name="queue_name",
    routing_key="routing_key",
)
publisher.publish({"pyrmq": "My first message"})
```

This publishes a message that uses a BlockingConnection on its own thread with default settings and and provides a handler for its retries.

Retries

PyRMQ's *Publisher* retries happen on two levels: connecting and publishing.

Connecting

PyRMQ instantiates a BlockingConnection when connecting. If this fails, it will retry for 2 more times by default with a delay of 5 seconds, a backoff base of 2 seconds, and a backoff constant of 5 seconds. All these settings are configurable via the *Publisher* class.

Publishing

PyRMQ calls pika's basic_publish when publishing. If this fails, it will retry for 2 more times by default with a delay of 5 seconds, a backoff base of 2 seconds, and a backoff constant of 5 seconds. All these settings are configurable via the *Publisher* class.

Max retries reached

When PyRMQ has tried one too many times, it will call your specified callback.

8.2.2 Publish message with priorities

To enable prioritization of messages, instantiate your queue with the queue argument *x-max-priority*. It takes an integer that sets the number of possible priority values with a higher number commanding more priority. Then, simply publish your message with the priority argument specified. Any number higher than the set max priority is floored or considered the same. Read more about message priorities here

```python
from pyrmq import Publisher
publisher = Publisher(
    exchange_name="exchange_name",
    queue_name="queue_name",
    routing_key="routing_key",
    queue_args={"x-max-priority": 3}
)
publisher.publish({"pyrmq": "My first message"}, priority=1)
```

> **Warning:** Adding arguments on an existing queue is not possible. If you wish to add queue arguments, you will need to either delete the existing queue then recreate the queue with arguments or simply make a new queue with the arguments.

8.2.3 Consuming

Instantiate the *Consumer* class and plug in your application specific settings. PyRMQ already works out of the box with RabbitMQ's default initialization settings.

```python
from pyrmq import Consumer


def callback(data):
    print(f"Received {data}!")

consumer = Consumer(
    exchange_name="exchange_name",
    queue_name="queue_name",
    routing_key="routing_key",
)

consumer.start()
```

Once the *Consumer* class is instantiated, just run `start()` to start its own thread that targets pika's start_consuming method on its own thread with default settings and and provides a handler for its retries. Consumption calls basic_ack with `delivery_tag` set to what the message's `method`'s was.

Retries

PyRMQ's *Consumer* retries happen on two levels: connecting and consuming.

Connecting

PyRMQ instantiates a BlockingConnection when connecting. If this fails, it will retry for 2 more times by default with a delay of 5 seconds, a backoff base of 2 seconds, and a backoff constant of 5 seconds. All these settings are configurable via the *Consumer* class.

DLX-DLK Consumption Retry Logic

PyRMQ calls pika's start_consuming when *Consumer* is instantiated. If your consumption callback throws an exception, PyRMQ uses dead letter exchanges and queues to republish your messages to your original queue once it has expired. PyRMQ already creates this "retry" queue for you with the default naming convention of appending your original queue with *.retry*. This is simply enabled by setting the `is_dlk_retry_enabled` flag on the *Consumer* class to `True`.

```python
from pyrmq import Consumer

def callback(data):
    print(f"Received {data}!")
    raise Exception

consumer = Consumer(
    exchange_name="exchange_name",
    queue_name="queue_name",
    routing_key="routing_key",
    callback=callback,
    is_dlk_retry_enabled=True,
)
consumer.start()
```

This will start a loop of passing your message between the original queue and the retry queue until it reaches the default number of `max_retries`.

Max retries reached

When PyRMQ has tried one too many times, it will call your specified callback.

8.3 API Documentation

8.3.1 Publisher Class

class pyrmq.**Publisher**(*exchange_name: str*, *queue_name: Optional[str] = ''*, *routing_key: Optional[str] = ''*, *exchange_type: Optional[str] = 'direct'*, ***kwargs*)

   This class uses a `BlockingConnection` from pika that automatically handles queue declarations and bindings plus retry logic built for its connection and publishing.

   __create_connection() → pika.adapters.blocking_connection.BlockingConnection
      Create pika's `BlockingConnection` from the given connection parameters.

__init__(*exchange_name: str, queue_name: Optional[str] = '', routing_key: Optional[str] = '',*
*exchange_type: Optional[str] = 'direct', **kwargs*)

> **Parameters**
>
> - **exchange_name** – Your exchange name.
>
> - **queue_name** – Your queue name.
>
> - **routing_key** – Your queue name.
>
> - **exchange_type** – Exchange type to declare. Default: `"direct"`
>
> - **host** – Your RabbitMQ host. Checks env var `RABBITMQ_HOST`. Default: `"localhost"`
>
> - **port** – Your RabbitMQ port. Checks env var `RABBITMQ_PORT`. Default: `5672`
>
> - **username** – Your RabbitMQ username. Default: `"guest"`
>
> - **password** – Your RabbitMQ password. Default: `"guest"`
>
> - **connection_attempts** – How many times should PyRMQ try?. Default: `3`
>
> - **retry_delay** – Seconds between connection retries. Default: `5`
>
> - **error_callback** – Callback function to be called when connection_attempts is reached.
>
> - **infinite_retry** – Tells PyRMQ to keep on retrying to publish while firing error_callback, if any. Default: `False`
>
> - **exchange_args** – Your exchange arguments. Default: `None`
>
> - **queue_args** – Your queue arguments. Default: `None`

__send_reconnection_error_message(*error*, *retry_count*) → None
    Send error message to your preferred location. :param error: Error that prevented the Publisher from sending the message. :param retry_count: Amount retries the Publisher tried before sending an error message.

__weakref__
    list of weak references to the object (if defined)

connect(*retry_count=1*) → pika.adapters.blocking_connection.BlockingChannel
    Create pika's `BlockingConnection` and initialize queue bindings. :param retry_count: Amount retries the Publisher tried before sending an error message.

declare_queue(*channel*) → None
    Declare and bind a channel to a queue. :param channel: pika Channel

publish(*data: dict, priority: Optional[int] = None, message_properties: Optional[dict] = None, attempt: int = 0, retry_count: int = 1*) → None
    Publish data to RabbitMQ. :param data: Data to be published. :param priority: Message priority. Only works if `x-max-priority` is defined as queue argument. :param message_properties: Message properties. Default: {`"delivery_mode":  2`} :param attempt: Number of attempts made. :param retry_count: Amount retries the Publisher tried before sending an error message.

8.3.2 Consumer Class

class pyrmq.**Consumer**(*exchange_name: str*, *queue_name: str*, *routing_key: str*, *callback: Callable*,
                                       exchange_type: Optional[str] = 'direct', ***kwargs*)

This class uses a `BlockingConnection` from pika that automatically handles queue declarations and bindings plus retry logic built for its connection and consumption. It starts its own thread upon initialization and runs pika's `start_consuming()`.

__create_connection() → pika.adapters.blocking_connection.BlockingConnection
    Create pika's `BlockingConnection` from the given connection parameters.

__init__(*exchange_name: str*, *queue_name: str*, *routing_key: str*, *callback: Callable*, *exchange_type: Optional[str] = 'direct'*, ***kwargs*)

> **Parameters**
>
> - **exchange_name** – Your exchange name.
>
> - **queue_name** – Your queue name.
>
> - **routing_key** – Your queue name.
>
> - **callback** – Your callback that should handle a consumed message
>
> - **host** – Your RabbitMQ host. Default: `"localhost"`
>
> - **port** – Your RabbitMQ port. Default: `5672`
>
> - **username** – Your RabbitMQ username. Default: `"guest"`
>
> - **password** – Your RabbitMQ password. Default: `"guest"`
>
> - **connection_attempts** – How many times should PyRMQ try? Default: `3`
>
> - **is_dlk_retry_enabled** – Flag to enable DLK-based retry logic of consumed messages. Default: `False`
>
> - **retry_delay** – Seconds between connection retries. Default: `5`
>
> - **retry_interval** – Seconds between consumption retries. Default: `900`
>
> - **retry_queue_suffix** – The suffix that will be appended to the `queue_name` to act as the name of the retry_queue. Default: `retry`
>
> - **max_retries** – Number of maximum retries for DLK retry logic. Default: `20`
>
> - **exchange_args** – Your exchange arguments. Default: `None`
>
> - **queue_args** – Your queue arguments. Default: `None`
>
> - **bound_exchange** – The exchange this consumer needs to bind to. This is an object that has two keys, `name` and `type`. Default: `None`
>
> - **auto_ack** – Flag whether to ack or nack the consumed message regardless of its outcome. Default: `True`
>
> - **prefetch_count** – How many messages should the consumer retrieve at a time for consumption. Default: `1`

__run_error_callback(*message: str*, *error: Exception*, *error_type: str*) → None
    Log error message :param message: Message to be logged in error_callback :param error: Error encountered in consuming the message :param error_type: Type of error (CONNECT_ERROR or CONSUME_ERROR)

__send_consume_error_message(*error: Exception*, *retry_count: int = 1*) → None
> Send error message to your preferred location. :param error: Error that prevented the Consumer from processing the message. :param retry_count: Amount retries the Consumer tried before sending an error message.

__send_reconnection_error_message(*error: Union[pika.exceptions.AMQPConnectionError, ConnectionResetError, pika.exceptions.ChannelClosedByBroker], retry_count: int*) → None
> Send error message to your preferred location. :param error: Error that prevented the Consumer from processing the message. :param retry_count: Amount retries the Consumer tried before sending an error message.

__weakref__
> list of weak references to the object (if defined)

_consume_message(*channel*, *method*, *properties*, *data: dict*) → None
> Wrap the user-provided callback, gracefully handle its errors, and call pika's `basic_ack` once successful. :param channel: pika's Channel this message was received. :param method: pika's basic Return :param properties: pika's BasicProperties :param data: Data received in bytes.

_publish_to_retry_queue(*data: dict*, *properties*, *retry_reason: Exception*) → None
> Publish message to retry queue with the appropriate metadata in the headers.

close() → None
> Manually close a connection to RabbitMQ. This is useful for debugging and tests.

connect(*retry_count=1*) → None
> Create pika's `BlockingConnection` and initialize queue bindings. :param retry_count: Amount retries the Consumer tried before sending an error message.

consume(*retry_count=1*) → None
> Wrap pika's `basic_consume()` and `start_consuming()` with retry logic.

declare_queue() → None
> Declare and bind a channel to a queue.

8.4 Testing PyRMQ

We're not gonna lie. Testing RabbitMQ, mocks or not, is infuriating. Much harder than a traditional integration testing with a database. That said, we hope that you could help us expand on what we have started should you feel our current tests aren't enough.

8.4.1 RabbitMQ

Since PyRMQ strives to be as complete with testing as it can be, it has several integration tests that need a running RabbitMQ to pass. Currently, PyRMQ is tested against `rabbitmq:3.8`.

Run Docker image (recommended)

```
$ docker run -d --hostname my-rabbit --name rabbitmq -p 5672:5672 rabbitmq:alpine
```

This allows you to connect to RabbitMQ via localhost through port 5672. Default credentials are `guest/guest`.

Install and run RabbitMQ locally

```
$ # Depending on your OS
$ # Ubuntu
$ sudo apt install rabbitmq
$ # Arch Linux
$ sudo pacman -S rabbitmq
```

8.4.2 Using tox

Install `pip install tox` and run:

```
$ tox
$ tox -e py38  # If this is what you have installed or don't want to bother testing for
→other versions
```

Symbols

C

D

P