rejected Documentation

Release 3.18.0

Gavin M. Roy

Mar 09, 2017

Contents

1 Features **3**

2 Installation **5**

3 Getting Started **7**
- 3.1 Consumer Examples . . . 7
- 3.2 Configuration File Syntax . . . 8
- 3.3 Configuration Example . . . 12
- 3.4 Command-Line Options . . . 14

4 API Documentation **15**
- 4.1 Consumer API . . . 15
- 4.2 Testing Support . . . 39

5 Issues **43**

6 Source **45**

7 Version History **47**

8 Indices and tables **49**

Python Module Index **51**

Rejected is a AMQP consumer daemon and message processing framework. It allows for rapid development of message processing consumers by handling all of the core functionality of communicating with RabbitMQ and management of consumer processes.

Rejected runs as a master process with multiple consumer configurations that are each run it an isolated process. It has the ability to collect statistical data from the consumer processes and report on it.

Rejected supports Python 2.7 and 3.4+.

CHAPTER 1

Features

- Automatic exception handling including connection management and consumer restarting
- Smart consumer classes that can automatically decode and deserialize message bodies based upon message headers
- Metrics logging and submission to statsd and InfluxDB
- Built-in profiling of consumer code
- Ability to write asynchronous code in consumers allowing for parallel communication with external resources

CHAPTER 2

Installation

rejected is available from the Python Package Index and can be installed by running **`pip install rejected`**.

For additional dependencies for optional features:

- To install HTML support, run **`pip install rejected[html]`**
- To install InfluxDB support, run **`pip install rejected[influxdb]`**
- To install MessagePack support, run **`pip install rejected[msgpack]`**
- To install Sentry support, run **`pip install rejected[sentry]`**

CHAPTER 3

Getting Started

Consumer Examples

The following example illustrates a very simple consumer that simply logs each message body as it's received.

```
from rejected import consumer
import logging

__version__ = '1.0.0'

LOGGER = logging.getLogger(__name__)


class ExampleConsumer(consumer.Consumer):

    def process(self):
        LOGGER.info(self.body)
```

All interaction with RabbitMQ with regard to connection management and message handling, including acknowledgements and rejections are automatically handled for you.

The `__version__` variable provides context in the rejected log files when consumers are started and can be useful for investigating consumer behaviors in production.

In this next example, a contrived `ExampleConsumer._connect_to_database` method is added that will return `False`. When `ExampleConsumer.process` evaluates if it could connect to the database and finds it can not, it will raise a *`rejected.consumer.ConsumerException`* which will requeue the message in RabbitMQ and increment an error counter. When too many errors occur, rejected will automatically restart the consumer after a brief quiet period. For more information on these exceptions, check out the *consumer API documentation*.

```
from rejected import consumer
import logging

__version__ = '1.0.0'
```

```
LOGGER = logging.getLogger(__name__)


class ExampleConsumer(consumer.Consumer):

    def _connect_to_database(self):
        return False

    def process(self):
        if not self._connect_to_database:
            raise consumer.ConsumerException('Database error')

        LOGGER.info(self.body)
```

Some consumers are also publishers. In this next example, the message body will be republished to a new exchange on the same RabbitMQ connection:

```
from rejected import consumer
import logging

__version__ = '1.0.0'

LOGGER = logging.getLogger(__name__)


class ExampleConsumer(consumer.PublishingConsumer):

    def process(self):
        LOGGER.info(self.body)
        self.publish('new-exchange', 'routing-key', {}, self.body)
```

Note that the previous example extends `rejected.consumer.PublishingConsumer` instead of `rejected.consumer.Consumer`. For more information about what base consumer classes exist, be sure to check out the *consumer API documentation*.

Configuration File Syntax

The rejected configuration uses YAML as the markup language. YAML's format, like Python code is whitespace dependent for control structure in blocks. If you're having problems with your rejected configuration, the first thing you should do is ensure that the YAML syntax is correct. yamllint.com is a good resource for validating that your configuration file can be parsed.

The configuration file is split into three main sections: Application, Daemon, and Logging.

The *example configuration* file provides a good starting point for creating your own configuration file.

Application

The application section of the configuration is broken down into multiple top-level options:

poll_interval	How often rejected should poll consumer processes for status in seconds (int/float)
sentry_dsn	If Sentry support is installed, optionally set a global DSN for all consumers (str)
stats	Enable and configure statsd metric submission (obj)
Connections	A subsection with RabbitMQ connection information for consumers (obj)
Consumers	Where each consumer type is configured (obj)

stats

stats		
	log	Toggle top-level logging of consumer process stats (bool)
	influxdb	Configure the submission of per-message measurements to InfluxDB (obj)
	statsd	Configure the submission of per-message measurements to statsd (obj)

influxdb

stats > influxdb		
	scheme	The scheme to use when submitting metrics to the InfluxDB server. Default: `http` (str)
	host	The hostname or ip address of the InfluxDB server. Default: `localhost` (str)
	port	The port of the influxdb server. Default: `8086` (int)
	user	An optional username to use when submitting measurements. (str)
	pass-word	An optional password to use when submitting measurements. (str)
	database	The InfluxDB database to submit measurements to. Default: `rejected` (str)

statsd

stats > statsd		
	enabled	Toggle statsd reporting off and on (bool)
	prefix	An optional prefix to use when creating the statsd metric path (str)
	host	The hostname or ip address of the statsd server (str)
	port	The port of the statsd server. Default: `8125` (int)

Connections

Each RabbitMQ connection entry should be a nested object with a unique name with connection attributes.

Connection Name		
	host	The hostname or ip address of the RabbitMQ server (str)
	port	The port of the RabbitMQ server (int)
	vhost	The virtual host to connect to (str)
	user	The username to connect as (str)
	pass	The password to use (str)
	heartbeat_interval	Optional: the AMQP heartbeat interval (int) default: 300 sec

Consumers

Each consumer entry should be a nested object with a unique name with consumer attributes.

Consumer Name		
	consumer	The package.module.Class path to the consumer code (str)
	connections	The connections to connect to (list) - See *Consumer Connections*
	qty	The number of consumers per connection to run (int)
	queue	The RabbitMQ queue name to consume from (int)
	ack	Explicitly acknowledge messages (no_ack = not ack) (bool)
	max_errors	Number of errors encountered before restarting a consumer (int)
	sentry_dsn	If Sentry support is installed, set a consumer specific sentry DSN (str)
	drop_exchange	The exchange to publish a message to when it is dropped. If not specified, dropped messages are not republished anywhere.
	drop_invalid_messages	Drop a message if the type property doesn't match the specified message type (str)
	message_type	Used to validate the message type of a message before processing. This attribute can be set to a string that is matched against the AMQP message type or a list of acceptable message types. (str, array)
	error_exchange	The exchange to publish messages that raise `ProcessingException` to (str)
	error_max_retry	The number of `ProcessingException` raised on a message before a message is dropped. If not specified messages will never be dropped (int)
	influxdb_measurement	When using InfluxDB, the measurement name for per-message measurements. Defaults to the consumer name. (str)
	config	Free-form key-value configuration section for the consumer (obj)

Consumer Connections

The consumer connections configuration allows for one or more connections to be made by a single consumer. This configuration section supports two formats. If a list of connection names are specified, the consumer will connect to and consume from the all of the specified connections.

```
Consumer Name:
    connections:
      - connection1
      - connection2
```

If the `connections` list include structured values, additional settings can be set. For example, you may want to consume from one RabbitMQ broker and publish to another, as is illustrated below:

```
Consumer Name:
    connections:
      - name: connection1
        consume: True
        publisher_confirmation: False
      - name: connection2
        consume: False
        publisher_confirmation: True
```

In the above example, the consumer will have two connections, `connection1` and `connection2`. It will only consume from `connection1` but can publish messages `connection2` by specifying the connection name in the `publish_message()` method.

Structured Connections

When specifying a structured consumer connection, the following attributes are available.

Consumer Name > connections		
	name	The connection name, as specified in the Connections section of the application configuration.
	consume	Specify if the connection should consume on the connection. (bool)
	publisher_confirmation	Enable publisher confirmations. (bool)

Daemon

This section contains the settings required to run the application as a daemon. They are as follows:

user	The username to run as when the process is daemonized (bool)
group	Optional The group name to switch to when the process is daemonized (str)
pidfile	The pidfile to write when the process is daemonized (str)

Logging

rejected uses `logging.config.dictConfig` to create a flexible method for configuring the python standard logging module. If rejected is being run in Python 2.6, logutils.dictconfig.dictConfig is used instead.

The following basic example illustrates all of the required sections in the dictConfig format, implemented in YAML:

```
version: 1
formatters: []
verbose:
  format: '%(levelname) -10s %(asctime)s %(process)-6d %(processName) -15s %(name) -
↪10s %(funcName) -20s: %(message)s'
  datefmt: '%Y-%m-%d %H:%M:%S'
handlers:
  console:
    class: logging.StreamHandler
    formatter: verbose
    debug_only: True
loggers:
  rejected:
    handlers: [console]
    level: INFO
    propagate: true
  myconsumer:
    handlers: [console]
    level: DEBUG
    propagate: true
disable_existing_loggers: true
incremental: false
```

Note: The debug_only node of the Logging > handlers > console section is not part of the standard dictConfig format. Please see the *Logging Caveats* section below for more information.

Logging Caveats

In order to allow for customizable console output when running in the foreground and no console output when daemonized, a `debug_only` node has been added to the standard dictConfig format in the handler section. This method is evaluated when logging is configured and if present, it is removed prior to passing the dictionary to dictConfig if present.

If the value is set to true and the application is not running in the foreground, the configuration for the handler and references to it will be removed from the configuration dictionary.

Troubleshooting

If you find that your application is not logging anything or sending output to the terminal, ensure that you have created a logger section in your configuration for your consumer package. For example if your Consumer instance is named `myconsumer.MyConsumer` make sure there is a `myconsumer` logger in the logging configuration.

Configuration Example

The following example will configure rejected to a consumer that connects to two different RabbitMQ servers, running two instances per connection, for a total of four consumer processes. It will consume from a queue named `generated_messages` and provides configuration for the consumer code itself that would consist of a dict with the keys `foo` and `bar`.

```
%YAML 1.2
---
Application:
  poll_interval: 10.0
  stats:
    log: True
    influxdb:
      host: localhost
      port: 8086
      database: rejected
    statsd:
      host: localhost
      port: 8125
  Connections:
    rabbit1:
      host: rabbit1
      port: 5672
      user: rejected
      pass: password
      ssl: False
      vhost: /
      heartbeat_interval: 300
    rabbit2:
      host: rabbit2
      port: 5672
      user: rejected
      pass: password
      ssl: False
      vhost: /
      heartbeat_interval: 300
  Consumers:
    example:
```

```
      consumer: example.Consumer
      connections: [rabbit1, rabbit2]
      drop_exchange: dlxname
      qty: 2
      queue: generated_messages
      dynamic_qos: True
      ack: True
      max_errors: 100
      config:
        foo: True
        bar: baz

Daemon:
  user: rejected
  group: daemon
  pidfile: /var/run/rejected.pid

Logging:
  version: 1
  formatters:
    verbose:
      format: "%(levelname) -10s %(asctime)s %(process)-6d %(processName) -25s
↪%(name) -30s %(funcName) -30s: %(message)s"
      datefmt: "%Y-%m-%d %H:%M:%S"
    syslog:
      format: "%(levelname)s <PID %(process)d:%(processName)s> %(name)s.
↪%(funcName)s(): %(message)s"
  filters: []
  handlers:
    console:
      class: logging.StreamHandler
      formatter: verbose
      debug_only: false
    syslog:
      class: logging.handlers.SysLogHandler
      facility: daemon
      address: /var/run/syslog
      #address: /dev/log
      formatter: syslog
  loggers:
    example:
      level: INFO
      propagate: true
      handlers: [console, syslog]
    helper:
      level: INFO
      propagate: true
      handlers: [console, syslog]
    rejected:
      level: INFO
      propagate: true
      handlers: [console, syslog]
    sprockets_influxdb:
      level: WARNING
      propagate: false
      handlers: [console, syslog]
  root:
    level: INFO
```

```
    propagate: true
    handlers: [console, syslog]
  disable_existing_loggers: true
  incremental: false
```

Command-Line Options

The rejected command line application allows you to spawn the rejected process as a daemon. Additionally it has options for running interactively (`-f`), which along with the `-o` switch for specifying a single consumer to run and `-q` to specify quantity, makes for easier debugging.

If you specify `-P /path/to/write/data/to`, rejected will automatically enable `cProfile`, writing the profiling data to the path specified. This can be used in conjunction with graphviz to diagram code execution and hotspots.

Help

```
usage: rejected [-h] [-c CONFIG] [-f] [-P PROFILE] [-o CONSUMER]
                     [-p PREPEND_PATH] [-q QUANTITY]

RabbitMQ consumer framework

optional arguments:
  -h, --help            show this help message and exit
  -c CONFIG, --config CONFIG
                        Path to the configuration file
  -f, --foreground      Run the application interactively
  -P PROFILE, --profile PROFILE
                        Profile the consumer modules, specifying the output
                        directory.
  -o CONSUMER, --only CONSUMER
                        Only run the consumer specified
  -p PREPEND_PATH, --prepend-path PREPEND_PATH
                        Prepend the python path with the value.
  -q QUANTITY, --qty QUANTITY
                        Run the specified quanty of consumer processes when
                        used in conjunction with -o
```

CHAPTER 4

API Documentation

Consumer API

The *Consumer* and *SmartConsumer* classes to extend for consumer applications.

While the *Consumer* class provides all the structure required for implementing a rejected consumer, the *SmartConsumer* adds functionality designed to make writing consumers even easier. When messages are received by consumers extending *SmartConsumer*, if the message's `content_type` property contains one of the supported mime-types, the message body will automatically be deserialized, making the deserialized message body available via the `body` attribute. Additionally, should one of the supported `content_encoding` types (`gzip` or `bzip2`) be specified in the message's property, it will automatically be decoded.

Message Type Validation

In any of the consumer base classes, if the `MESSAGE_TYPE` attribute is set, the `type` property of incoming messages will be validated against when a message is received, checking for string equality against the `MESSAGE_TYPE` attribute. If they are not matched, the consumer will not process the message and will drop the message without an exception if the `DROP_INVALID_MESSAGES` attribute is set to `True`. If it is `False`, a *ConsumerException* is raised.

Republishing of Dropped Messages

If the consumer is configured by specifying `DROP_EXCHANGE` as an attribute of the consumer class or in the consumer configuration with the `drop_exchange` configuration variable, when a message is dropped, it is published to that exchange prior to the message being rejected in RabbitMQ. When the message is republished, four new values are added to the AMQP `headers` message property: `X-Dropped-By`, `X-Dropped-Reason`, `X-Dropped-Timestamp`, `X-Original-Exchange`.

The `X-Dropped-By` header value contains the configured name of the consumer that dropped the message. `X-Dropped-Reason` contains the reason the message was dropped (eg invalid message type or maximum error count). `X-Dropped-Timestamp` value contains the ISO-8601 formatted timestamp of when the message was

dropped. Finally, the `X-Original-Exchange` value contains the original exchange that the message was published to.

Consumer Classes

Consumer

class `rejected.consumer.`**`Consumer`**(*settings*, *process*, *drop_invalid_messages=None*, *message_type=None*, *error_exchange=None*, *error_max_retry=None*, *drop_exchange=None*)

Base consumer class that defines the contract between rejected and consumer applications.

In any of the consumer base classes, if the `message_type` is specified in the configuration (or set with the `MESSAGE_TYPE` attribute), the `type` property of incoming messages will be validated against when a message is received. If there is no match, the consumer will not process the message and will drop the message without an exception if the `drop_invalid_messages` setting is set to `True` in the configuration (or if the `DROP_INVALID_MESSAGES` attribute is set to `True`). If it is `False`, a *`MessageException`* is raised.

If `DROP_EXCHANGE` is specified either as an attribute of the consumer class or in the consumer configuration, if a message is dropped, it is published to the that exchange prior to rejecting the message in RabbitMQ. When the message is republished, four new values are added to the AMQP `headers` message property: `X-Dropped-By`, `X-Dropped-Reason`, `X-Dropped-Timestamp`, `X-Original-Exchange`.

The `X-Dropped-By` header value contains the configured name of the consumer that dropped the message. `X-Dropped-Reason` contains the reason the message was dropped (eg invalid message type or maximum error count). `X-Dropped-Timestamp` value contains the ISO-8601 formatted timestamp of when the message was dropped. Finally, the `X-Original-Exchange` value contains the original exchange that the message was published to.

If a consumer raises a *`ProcessingException`*, the message that was being processed will be republished to the exchange specified by the `error` exchange configuration value or the `ERROR_EXCHANGE` attribute of the consumer's class. The message will be published using the routing key that was last used for the message. The original message body and properties will be used and two additional header property values may be added:

- **`X-Processing-Exception` contains the string value of the** exception that was raised, if specified.
- **`X-Processing-Exceptions` contains the quantity of processing** exceptions that have been raised for the message.

In combination with a queue that has `x-message-ttl` set and `x-dead-letter-exchange` that points to the original exchange for the queue the consumer is consuming off of, you can implement a delayed retry cycle for messages that are failing to process due to external resource or service issues.

If `error_max_retry` is specified in the configuration or `ERROR_MAX_RETRY` is set on the class, the headers for each method will be inspected and if the value of `X-Processing-Exceptions` is greater than or equal to the specified value, the message will be dropped.

Note: Since 3.17, *`Consumer`* and *`PublishingConsumer`* have been combined into the same class.

`app_id`

Access the current message's `app-id` property as an attribute of the consumer class.

Return type str

`body`

Access the opaque body from the current message.

Return type str

content_encoding
Access the current message's `content-encoding` AMQP message property as an attribute of the consumer class.

Return type str

content_type
Access the current message's `content-type` AMQP message property as an attribute of the consumer class.

Return type str

correlation_id
Access the current message's `correlation-id` AMAP message property as an attribute of the consumer class. If the message does not have a `correlation-id` then, each message is assigned a new UUIDv4 based `correlation-id` value.

Return type str

exchange
Access the AMQP exchange the message was published to as an attribute of the consumer class.

Return type str

expiration
Access the current message's `expiration` AMQP message property as an attribute of the consumer class.

Return type str

finish()
Finishes message processing for the current message. If this is called in *prepare()*, the *process()* method is not invoked for the current message.

headers
Access the current message's `headers` AMQP message property as an attribute of the consumer class.

Return type dict

initialize()
Extend this method for any initialization tasks that occur only when the *Consumer* class is created.

message_id
Access the current message's `message-id` AMQP message property as an attribute of the consumer class.

Return type str

message_type
Access the current message's `type` AMQP message property as an attribute of the consumer class.

Return type str

name
Property returning the name of the consumer class.

Return type str

on_blocked(*name*)
Called when a connection for this consumer is blocked.

Override this method to respond to being blocked.

New in version 3.17.

Parameters **name** (*str*) – The connection name that is blocked

on_finish()

Called after a message has been processed.

Override this method to perform cleanup, logging, etc. This method is a counterpart to *prepare()*. `on_finish` may not produce any output, as it is called after all processing has taken place.

If an exception is raised during the processing of a message, *prepare()* is not invoked.

Note: Asynchronous support: Decorate this method with `tornado.gen.coroutine()` to make it asynchronous.

on_unblocked(*name*)

Called when a connection for this consumer is unblocked.

Override this method to respond to being blocked.

New in version 3.17.

Parameters **name** (*str*) – The connection name that is blocked

prepare()

Called when a message is received before *process()*.

Note: Asynchronous support: Decorate this method with `tornado.gen.coroutine()` to make it asynchronous.

If this method returns a `Future`, execution will not proceed until the Future has completed.

priority

Access the current message's `priority` AMQP message property as an attribute of the consumer class.

Return type int

process()

Extend this method for implementing your Consumer logic.

If the message can not be processed and the Consumer should stop after n failures to process messages, raise the *ConsumerException*.

Note: Asynchronous support: Decorate this method with `tornado.gen.coroutine()` to make it asynchronous.

Raises *rejected.consumer.ConsumerException*

Raises *rejected.consumer.MessageException*

Raises *rejected.consumer.ProcessingException*

properties

Access the current message's AMQP message properties in dict form as an attribute of the consumer class.

Return type dict

publish_message(*exchange*, *routing_key*, *properties*, *body*, *channel=None*)

Publish a message to RabbitMQ on the same channel the original message was received on.

Parameters

- **exchange** (*str*) – The exchange to publish to
- **routing_key** (*str*) – The routing key to publish with
- **properties** (*dict*) – The message properties
- **body** (*str*) – The message body
- **channel** (*str*) – The channel/connection name to use. If it is not specified, the channel that the message was delivered on is used.

redelivered

Indicates if the current message has been redelivered.

Return type bool

reply(*response_body*, *properties*, *auto_id=True*, *exchange=None*, *reply_to=None*)

Reply to the received message.

If `auto_id` is `True`, a new UUIDv4 value will be generated for the `message_id` AMQP message property. The `correlation_id` AMQP message property will be set to the `message_id` of the original message. In addition, the `timestamp` will be assigned the current time of the message. If `auto_id` is `False`, neither the `message_id` and the `correlation_id` AMQP properties will be changed in the properties.

If `exchange` is not set, the exchange the message was received on will be used.

If `reply_to` is set in the original properties, it will be used as the routing key. If the `reply_to` is not set in the properties and it is not passed in, a `ValueError` will be raised. If reply to is set in the properties, it will be cleared out prior to the message being republished.

Parameters

- **response_body** (*any*) – The message body to send
- **properties** (`rejected.data.Properties`) – Message properties to use
- **auto_id** (*bool*) – Automatically shuffle `message_id` & `correlation_id`
- **exchange** (*str*) – Override the exchange to publish to
- **reply_to** (*str*) – Override the `reply_to` AMQP property

Raises `ValueError`

reply_to

Access the current message's `reply-to` AMQP message property as an attribute of the consumer class.

Return type str

returned

Indicates if the message was delivered by consumer previously and returned from RabbitMQ.

New in version 3.17.

Return type bool

routing_key

Access the routing key for the current message.

Return type str

send_exception_to_sentry(*exc_info*)

Send an exception to Sentry if enabled.

Parameters **exc_info** (*tuple*) – exception information as returned from `sys.exc_info()`

sentry_client

Access the Sentry raven `Client` instance or `None`

Use this object to add tags or additional context to Sentry error reports (see `raven.base.Client.tags_context()`) or to report messages (via `raven.base.Client.captureMessage()`) directly to Sentry.

Return type `raven.base.Client`

set_sentry_context(*tag*, *value*)

Set a context tag in Sentry for the given key and value.

Parameters

- **tag** (*str*) – The context tag name
- **value** (*str*) – The context value

settings

Access the consumer settings as specified by the `config` section for the consumer in the rejected configuration.

Return type dict

shutdown()

Override to cleanly shutdown when rejected is stopping the consumer.

This could be used for closing database connections or other such activities.

stats_add_timing(*key*, *duration*)

Add a timing to the per-message measurements

New in version 3.13.0.

Parameters

- **key** (*str*) – The key to add the timing to
- **duration** (*int* | *float*) – The timing value

stats_incr(*key*, *value=1*)

Increment the specified key in the per-message measurements

New in version 3.13.0.

Parameters

- **key** (*str*) – The key to increment
- **value** (*int*) – The value to increment the key by

stats_set_tag(*key*, *value=1*)

Set the specified tag/value in the per-message measurements

New in version 3.13.0.

Parameters

- **key** (*str*) – The key to increment
- **value** (*int*) – The value to increment the key by

stats_set_value(*key*, *value=1*)
Set the specified key/value in the per-message measurements

New in version 3.13.0.

Parameters

- **key** (*str*) – The key to increment
- **value** (*int*) – The value to increment the key by

stats_track_duration(**args*, ***kwds*)
Time around a context and add to the the per-message measurements

New in version 3.13.0.

Parameters **key** (*str*) – The key for the timing to track

statsd_add_timing(*key*, *duration*)
Add a timing to the per-message measurements

Parameters

- **key** (*str*) – The key to add the timing to
- **duration** (*int* | *float*) – The timing value

Deprecated since version 3.13.0.

statsd_incr(*key*, *value=1*)
Increment the specified key in the per-message measurements

Parameters

- **key** (*str*) – The key to increment
- **value** (*int*) – The value to increment the key by

Deprecated since version 3.13.0.

statsd_track_duration(*key*)
Time around a context and add to the the per-message measurements

Parameters **key** (*str*) – The key for the timing to track

Deprecated since version 3.13.0.

timestamp
Access the unix epoch timestamp value from the AMQP message properties of the current message.

Return type int

unset_sentry_context(*tag*)
Remove a context tag from sentry

Parameters **tag** (*str*) – The context tag to remove

user_id
Access the `user-id` AMQP message property from the current message's properties.

Return type str

yield_to_ioloop(**args*, ***kwargs*)
Function that will allow Rejected to process IOLoop events while in a tight-loop inside an asynchronous consumer.

class `rejected.consumer.`**`PublishingConsumer`**(**args*, ***kwargs*)

Deprecated, functionality moved to *`rejected.consumer.Consumer`*

Deprecated since version 3.17.0.

`app_id`

Access the current message's `app-id` property as an attribute of the consumer class.

Return type str

`body`

Access the opaque body from the current message.

Return type str

`content_encoding`

Access the current message's `content-encoding` AMQP message property as an attribute of the consumer class.

Return type str

`content_type`

Access the current message's `content-type` AMQP message property as an attribute of the consumer class.

Return type str

`correlation_id`

Access the current message's `correlation-id` AMAP message property as an attribute of the consumer class. If the message does not have a `correlation-id` then, each message is assigned a new UUIDv4 based `correlation-id` value.

Return type str

`exchange`

Access the AMQP exchange the message was published to as an attribute of the consumer class.

Return type str

`expiration`

Access the current message's `expiration` AMQP message property as an attribute of the consumer class.

Return type str

`finish`()

Finishes message processing for the current message. If this is called in *`prepare()`*, the *`process()`* method is not invoked for the current message.

`headers`

Access the current message's `headers` AMQP message property as an attribute of the consumer class.

Return type dict

`initialize`()

Extend this method for any initialization tasks that occur only when the *`Consumer`* class is created.

`message_id`

Access the current message's `message-id` AMQP message property as an attribute of the consumer class.

Return type str

`message_type`

Access the current message's `type` AMQP message property as an attribute of the consumer class.

Return type str

name

Property returning the name of the consumer class.

Return type str

on_blocked(*name*)

Called when a connection for this consumer is blocked.

Override this method to respond to being blocked.

New in version 3.17.

Parameters **name** (*str*) – The connection name that is blocked

on_finish()

Called after a message has been processed.

Override this method to perform cleanup, logging, etc. This method is a counterpart to *prepare()*. `on_finish` may not produce any output, as it is called after all processing has taken place.

If an exception is raised during the processing of a message, *prepare()* is not invoked.

Note: Asynchronous support: Decorate this method with `tornado.gen.coroutine()` to make it asynchronous.

on_unblocked(*name*)

Called when a connection for this consumer is unblocked.

Override this method to respond to being blocked.

New in version 3.17.

Parameters **name** (*str*) – The connection name that is blocked

prepare()

Called when a message is received before *process()*.

Note: Asynchronous support: Decorate this method with `tornado.gen.coroutine()` to make it asynchronous.

If this method returns a `Future`, execution will not proceed until the Future has completed.

priority

Access the current message's `priority` AMQP message property as an attribute of the consumer class.

Return type int

process()

Extend this method for implementing your Consumer logic.

If the message can not be processed and the Consumer should stop after n failures to process messages, raise the *ConsumerException*.

Note: Asynchronous support: Decorate this method with `tornado.gen.coroutine()` to make it asynchronous.

Raises *rejected.consumer.ConsumerException*

Raises `rejected.consumer.MessageException`

Raises `rejected.consumer.ProcessingException`

properties

Access the current message's AMQP message properties in dict form as an attribute of the consumer class.

Return type dict

publish_message(*exchange*, *routing_key*, *properties*, *body*, *channel=None*)

Publish a message to RabbitMQ on the same channel the original message was received on.

Parameters

- **exchange** (*str*) – The exchange to publish to
- **routing_key** (*str*) – The routing key to publish with
- **properties** (*dict*) – The message properties
- **body** (*str*) – The message body
- **channel** (*str*) – The channel/connection name to use. If it is not specified, the channel that the message was delivered on is used.

redelivered

Indicates if the current message has been redelivered.

Return type bool

reply(*response_body*, *properties*, *auto_id=True*, *exchange=None*, *reply_to=None*)

Reply to the received message.

If `auto_id` is `True`, a new UUIDv4 value will be generated for the `message_id` AMQP message property. The `correlation_id` AMQP message property will be set to the `message_id` of the original message. In addition, the `timestamp` will be assigned the current time of the message. If `auto_id` is `False`, neither the `message_id` and the `correlation_id` AMQP properties will be changed in the properties.

If `exchange` is not set, the exchange the message was received on will be used.

If `reply_to` is set in the original properties, it will be used as the routing key. If the `reply_to` is not set in the properties and it is not passed in, a `ValueError` will be raised. If reply to is set in the properties, it will be cleared out prior to the message being republished.

Parameters

- **response_body** (*any*) – The message body to send
- **properties** (`rejected.data.Properties`) – Message properties to use
- **auto_id** (*bool*) – Automatically shuffle `message_id` & `correlation_id`
- **exchange** (*str*) – Override the exchange to publish to
- **reply_to** (*str*) – Override the `reply_to` AMQP property

Raises `ValueError`

reply_to

Access the current message's `reply-to` AMQP message property as an attribute of the consumer class.

Return type str

returned
Indicates if the message was delivered by consumer previously and returned from RabbitMQ.

New in version 3.17.

Return type bool

routing_key
Access the routing key for the current message.

Return type str

send_exception_to_sentry(*exc_info*)
Send an exception to Sentry if enabled.

Parameters **exc_info** (*tuple*) – exception information as returned from `sys.exc_info()`

sentry_client
Access the Sentry raven `Client` instance or `None`

Use this object to add tags or additional context to Sentry error reports (see `raven.base.Client.tags_context()`) or to report messages (via `raven.base.Client.captureMessage()`) directly to Sentry.

Return type `raven.base.Client`

set_sentry_context(*tag*, *value*)
Set a context tag in Sentry for the given key and value.

Parameters

- **tag** (*str*) – The context tag name
- **value** (*str*) – The context value

settings
Access the consumer settings as specified by the `config` section for the consumer in the rejected configuration.

Return type dict

shutdown()
Override to cleanly shutdown when rejected is stopping the consumer.

This could be used for closing database connections or other such activities.

stats_add_timing(*key*, *duration*)
Add a timing to the per-message measurements

New in version 3.13.0.

Parameters

- **key** (*str*) – The key to add the timing to
- **duration** (*int | float*) – The timing value

stats_incr(*key*, *value=1*)
Increment the specified key in the per-message measurements

New in version 3.13.0.

Parameters

- **key** (*str*) – The key to increment

- **value** (*int*) – The value to increment the key by

stats_set_tag(*key*, *value=1*)

Set the specified tag/value in the per-message measurements

New in version 3.13.0.

Parameters

- **key** (*str*) – The key to increment
- **value** (*int*) – The value to increment the key by

stats_set_value(*key*, *value=1*)

Set the specified key/value in the per-message measurements

New in version 3.13.0.

Parameters

- **key** (*str*) – The key to increment
- **value** (*int*) – The value to increment the key by

stats_track_duration(**args*, ***kwds*)

Time around a context and add to the the per-message measurements

New in version 3.13.0.

Parameters **key** (*str*) – The key for the timing to track

statsd_add_timing(*key*, *duration*)

Add a timing to the per-message measurements

Parameters

- **key** (*str*) – The key to add the timing to
- **duration** (*int* | *float*) – The timing value

Deprecated since version 3.13.0.

statsd_incr(*key*, *value=1*)

Increment the specified key in the per-message measurements

Parameters

- **key** (*str*) – The key to increment
- **value** (*int*) – The value to increment the key by

Deprecated since version 3.13.0.

statsd_track_duration(*key*)

Time around a context and add to the the per-message measurements

Parameters **key** (*str*) – The key for the timing to track

Deprecated since version 3.13.0.

timestamp

Access the unix epoch timestamp value from the AMQP message properties of the current message.

Return type int

unset_sentry_context(*tag*)

Remove a context tag from sentry

Parameters **tag** (*str*) – The context tag to remove

user_id
Access the `user-id` AMQP message property from the current message's properties.

Return type str

yield_to_ioloop(**args*, ***kwargs*)
Function that will allow Rejected to process IOLoop events while in a tight-loop inside an asynchronous consumer.

SmartConsumer

class `rejected.consumer.`**SmartConsumer**(*settings*, *process*, *drop_invalid_messages=None*, *message_type=None*, *error_exchange=None*, *error_max_retry=None*, *drop_exchange=None*)

Base class to ease the implementation of strongly typed message consumers that validate and automatically decode and deserialize the inbound message body based upon the message properties. Additionally, should one of the supported `content_encoding` types (`gzip` or `bzip2`) be specified in the message's property, it will automatically be decoded.

When publishing a message, the message can be automatically serialized and encoded. If the `content_type` property is specified, the consumer will attempt to automatically serialize the message body. If the `content_encoding` property is specified using a supported encoding (`gzip` or `bzip2`), it will automatically be encoded as well.

Supported MIME types for automatic serialization and deserialization are:

- application/json
- application/pickle
- application/x-pickle
- application/x-plist
- application/x-vnd.python.pickle
- application/vnd.python.pickle
- text/csv
- text/html (with beautifulsoup4 installed)
- text/xml (with beautifulsoup4 installed)
- text/yaml
- text/x-yaml

In any of the consumer base classes, if the `MESSAGE_TYPE` attribute is set, the `type` property of incoming messages will be validated against when a message is received, checking for string equality against the `MESSAGE_TYPE` attribute. If they are not matched, the consumer will not process the message and will drop the message without an exception if the `DROP_INVALID_MESSAGES` attribute is set to `True`. If it is `False`, a `ConsumerException` is raised.

Note: Since 3.17, *`SmartConsumer`* and *`SmartPublishingConsumer`* have been combined into the same class.

app_id
Access the current message's `app-id` property as an attribute of the consumer class.

Return type str

body
Return the message body, unencoded if needed, deserialized if possible.

Return type any

content_encoding
Access the current message's `content-encoding` AMQP message property as an attribute of the consumer class.

Return type str

content_type
Access the current message's `content-type` AMQP message property as an attribute of the consumer class.

Return type str

correlation_id
Access the current message's `correlation-id` AMAP message property as an attribute of the consumer class. If the message does not have a `correlation-id` then, each message is assigned a new UUIDv4 based `correlation-id` value.

Return type str

exchange
Access the AMQP exchange the message was published to as an attribute of the consumer class.

Return type str

expiration
Access the current message's `expiration` AMQP message property as an attribute of the consumer class.

Return type str

finish()
Finishes message processing for the current message. If this is called in *prepare()*, the *process()* method is not invoked for the current message.

headers
Access the current message's `headers` AMQP message property as an attribute of the consumer class.

Return type dict

initialize()
Extend this method for any initialization tasks that occur only when the *Consumer* class is created.

message_id
Access the current message's `message-id` AMQP message property as an attribute of the consumer class.

Return type str

message_type
Access the current message's `type` AMQP message property as an attribute of the consumer class.

Return type str

name
Property returning the name of the consumer class.

Return type str

on_blocked(*name*)

Called when a connection for this consumer is blocked.

Override this method to respond to being blocked.

New in version 3.17.

Parameters **name** (*str*) – The connection name that is blocked

on_finish()

Called after a message has been processed.

Override this method to perform cleanup, logging, etc. This method is a counterpart to *prepare()*. `on_finish` may not produce any output, as it is called after all processing has taken place.

If an exception is raised during the processing of a message, *prepare()* is not invoked.

Note: Asynchronous support: Decorate this method with `tornado.gen.coroutine()` to make it asynchronous.

on_unblocked(*name*)

Called when a connection for this consumer is unblocked.

Override this method to respond to being blocked.

New in version 3.17.

Parameters **name** (*str*) – The connection name that is blocked

prepare()

Called when a message is received before *process()*.

Note: Asynchronous support: Decorate this method with `tornado.gen.coroutine()` to make it asynchronous.

If this method returns a `Future`, execution will not proceed until the Future has completed.

priority

Access the current message's `priority` AMQP message property as an attribute of the consumer class.

Return type int

process()

Extend this method for implementing your Consumer logic.

If the message can not be processed and the Consumer should stop after n failures to process messages, raise the *ConsumerException*.

Note: Asynchronous support: Decorate this method with `tornado.gen.coroutine()` to make it asynchronous.

Raises *rejected.consumer.ConsumerException*

Raises *rejected.consumer.MessageException*

Raises *rejected.consumer.ProcessingException*

properties

Access the current message's AMQP message properties in dict form as an attribute of the consumer class.

Return type dict

publish_message(*exchange*, *routing_key*, *properties*, *body*, *no_serialization=False*, *no_encoding=False*, *channel=None*)

Publish a message to RabbitMQ on the same channel the original message was received on.

By default, if you pass a non-string object to the body and the properties have a supported content-type set, the body will be auto-serialized in the specified content-type.

If the properties do not have a timestamp set, it will be set to the current time.

If you specify a content-encoding in the properties and the encoding is supported, the body will be auto-encoded.

Both of these behaviors can be disabled by setting no_serialization or no_encoding to True.

Parameters

- **exchange** (*str*) – The exchange to publish to
- **routing_key** (*str*) – The routing key to publish with
- **properties** (*dict*) – The message properties
- **body** (*mixed*) – The message body to publish
- **no_serialization** (*bool*) – Turn off auto-serialization of the body
- **no_encoding** (*bool*) – Turn off auto-encoding of the body
- **channel** (*str*) – The channel/connection name to use. If it is not specified, the channel that the message was delivered on is used.

redelivered

Indicates if the current message has been redelivered.

Return type bool

reply(*response_body*, *properties*, *auto_id=True*, *exchange=None*, *reply_to=None*)

Reply to the received message.

If `auto_id` is `True`, a new UUIDv4 value will be generated for the `message_id` AMQP message property. The `correlation_id` AMQP message property will be set to the `message_id` of the original message. In addition, the `timestamp` will be assigned the current time of the message. If `auto_id` is `False`, neither the `message_id` and the `correlation_id` AMQP properties will be changed in the properties.

If `exchange` is not set, the exchange the message was received on will be used.

If `reply_to` is set in the original properties, it will be used as the routing key. If the `reply_to` is not set in the properties and it is not passed in, a `ValueError` will be raised. If reply to is set in the properties, it will be cleared out prior to the message being republished.

Parameters

- **response_body** (*any*) – The message body to send
- **properties** (`rejected.data.Properties`) – Message properties to use
- **auto_id** (*bool*) – Automatically shuffle `message_id` & `correlation_id`
- **exchange** (*str*) – Override the exchange to publish to
- **reply_to** (*str*) – Override the `reply_to` AMQP property

Raises `ValueError`

reply_to

Access the current message's `reply-to` AMQP message property as an attribute of the consumer class.

Return type str

returned

Indicates if the message was delivered by consumer previously and returned from RabbitMQ.

New in version 3.17.

Return type bool

routing_key

Access the routing key for the current message.

Return type str

send_exception_to_sentry(*exc_info*)

Send an exception to Sentry if enabled.

Parameters **exc_info** (*tuple*) – exception information as returned from `sys.exc_info()`

sentry_client

Access the Sentry raven `Client` instance or `None`

Use this object to add tags or additional context to Sentry error reports (see `raven.base.Client.tags_context()`) or to report messages (via `raven.base.Client.captureMessage()`) directly to Sentry.

Return type `raven.base.Client`

set_sentry_context(*tag*, *value*)

Set a context tag in Sentry for the given key and value.

Parameters

- **tag** (*str*) – The context tag name
- **value** (*str*) – The context value

settings

Access the consumer settings as specified by the `config` section for the consumer in the rejected configuration.

Return type dict

shutdown()

Override to cleanly shutdown when rejected is stopping the consumer.

This could be used for closing database connections or other such activities.

stats_add_timing(*key*, *duration*)

Add a timing to the per-message measurements

New in version 3.13.0.

Parameters

- **key** (*str*) – The key to add the timing to
- **duration** (*int | float*) – The timing value

stats_incr(*key*, *value=1*)
Increment the specified key in the per-message measurements

New in version 3.13.0.

Parameters

- **key** (*str*) – The key to increment
- **value** (*int*) – The value to increment the key by

stats_set_tag(*key*, *value=1*)
Set the specified tag/value in the per-message measurements

New in version 3.13.0.

Parameters

- **key** (*str*) – The key to increment
- **value** (*int*) – The value to increment the key by

stats_set_value(*key*, *value=1*)
Set the specified key/value in the per-message measurements

New in version 3.13.0.

Parameters

- **key** (*str*) – The key to increment
- **value** (*int*) – The value to increment the key by

stats_track_duration(**args*, ***kwds*)
Time around a context and add to the the per-message measurements

New in version 3.13.0.

Parameters **key** (*str*) – The key for the timing to track

statsd_add_timing(*key*, *duration*)
Add a timing to the per-message measurements

Parameters

- **key** (*str*) – The key to add the timing to
- **duration** (*int* | *float*) – The timing value

Deprecated since version 3.13.0.

statsd_incr(*key*, *value=1*)
Increment the specified key in the per-message measurements

Parameters

- **key** (*str*) – The key to increment
- **value** (*int*) – The value to increment the key by

Deprecated since version 3.13.0.

statsd_track_duration(*key*)
Time around a context and add to the the per-message measurements

Parameters **key** (*str*) – The key for the timing to track

Deprecated since version 3.13.0.

timestamp
Access the unix epoch timestamp value from the AMQP message properties of the current message.

Return type int

unset_sentry_context(*tag*)
Remove a context tag from sentry

Parameters **tag** (*str*) – The context tag to remove

user_id
Access the `user-id` AMQP message property from the current message's properties.

Return type str

yield_to_ioloop(**args*, ***kwargs*)
Function that will allow Rejected to process IOLoop events while in a tight-loop inside an asynchronous consumer.

class rejected.consumer.**SmartPublishingConsumer**(**args*, ***kwargs*)
Deprecated, functionality moved to *rejected.consumer.SmartConsumer*

Deprecated since version 3.17.0.

app_id
Access the current message's `app-id` property as an attribute of the consumer class.

Return type str

body
Return the message body, unencoded if needed, deserialized if possible.

Return type any

content_encoding
Access the current message's `content-encoding` AMQP message property as an attribute of the consumer class.

Return type str

content_type
Access the current message's `content-type` AMQP message property as an attribute of the consumer class.

Return type str

correlation_id
Access the current message's `correlation-id` AMAP message property as an attribute of the consumer class. If the message does not have a `correlation-id` then, each message is assigned a new UUIDv4 based `correlation-id` value.

Return type str

exchange
Access the AMQP exchange the message was published to as an attribute of the consumer class.

Return type str

expiration
Access the current message's `expiration` AMQP message property as an attribute of the consumer class.

Return type str

finish()
Finishes message processing for the current message. If this is called in *prepare()*, the *process()* method is not invoked for the current message.

headers
Access the current message's `headers` AMQP message property as an attribute of the consumer class.

Return type dict

initialize()
Extend this method for any initialization tasks that occur only when the *Consumer* class is created.

message_id
Access the current message's `message-id` AMQP message property as an attribute of the consumer class.

Return type str

message_type
Access the current message's `type` AMQP message property as an attribute of the consumer class.

Return type str

name
Property returning the name of the consumer class.

Return type str

on_blocked(*name*)
Called when a connection for this consumer is blocked.

Override this method to respond to being blocked.

New in version 3.17.

Parameters **name** (*str*) – The connection name that is blocked

on_finish()
Called after a message has been processed.

Override this method to perform cleanup, logging, etc. This method is a counterpart to *prepare()*. `on_finish` may not produce any output, as it is called after all processing has taken place.

If an exception is raised during the processing of a message, *prepare()* is not invoked.

Note: Asynchronous support: Decorate this method with `tornado.gen.coroutine()` to make it asynchronous.

---

on_unblocked(*name*)
Called when a connection for this consumer is unblocked.

Override this method to respond to being blocked.

New in version 3.17.

Parameters **name** (*str*) – The connection name that is blocked

prepare()
Called when a message is received before *process()*.

Note: Asynchronous support: Decorate this method with `tornado.gen.coroutine()` to make it asynchronous.

If this method returns a `Future`, execution will not proceed until the Future has completed.

priority

Access the current message's `priority` AMQP message property as an attribute of the consumer class.

Return type int

process()

Extend this method for implementing your Consumer logic.

If the message can not be processed and the Consumer should stop after n failures to process messages, raise the *ConsumerException*.

Note: Asynchronous support: Decorate this method with `tornado.gen.coroutine()` to make it asynchronous.

Raises *rejected.consumer.ConsumerException*

Raises *rejected.consumer.MessageException*

Raises *rejected.consumer.ProcessingException*

properties

Access the current message's AMQP message properties in dict form as an attribute of the consumer class.

Return type dict

publish_message(*exchange*, *routing_key*, *properties*, *body*, *no_serialization=False*, *no_encoding=False*, *channel=None*)

Publish a message to RabbitMQ on the same channel the original message was received on.

By default, if you pass a non-string object to the body and the properties have a supported content-type set, the body will be auto-serialized in the specified content-type.

If the properties do not have a timestamp set, it will be set to the current time.

If you specify a content-encoding in the properties and the encoding is supported, the body will be auto-encoded.

Both of these behaviors can be disabled by setting no_serialization or no_encoding to True.

Parameters

- **exchange** (*str*) – The exchange to publish to
- **routing_key** (*str*) – The routing key to publish with
- **properties** (*dict*) – The message properties
- **body** (*mixed*) – The message body to publish
- **no_serialization** (*bool*) – Turn off auto-serialization of the body
- **no_encoding** (*bool*) – Turn off auto-encoding of the body
- **channel** (*str*) – The channel/connection name to use. If it is not specified, the channel that the message was delivered on is used.

redelivered
Indicates if the current message has been redelivered.

Return type bool

reply(*response_body*, *properties*, *auto_id=True*, *exchange=None*, *reply_to=None*)
Reply to the received message.

If `auto_id` is `True`, a new UUIDv4 value will be generated for the `message_id` AMQP message property. The `correlation_id` AMQP message property will be set to the `message_id` of the original message. In addition, the `timestamp` will be assigned the current time of the message. If `auto_id` is `False`, neither the `message_id` and the `correlation_id` AMQP properties will be changed in the properties.

If `exchange` is not set, the exchange the message was received on will be used.

If `reply_to` is set in the original properties, it will be used as the routing key. If the `reply_to` is not set in the properties and it is not passed in, a `ValueError` will be raised. If reply to is set in the properties, it will be cleared out prior to the message being republished.

Parameters

- **response_body** (*any*) – The message body to send
- **properties** (`rejected.data.Properties`) – Message properties to use
- **auto_id** (*bool*) – Automatically shuffle `message_id` & `correlation_id`
- **exchange** (*str*) – Override the exchange to publish to
- **reply_to** (*str*) – Override the `reply_to` AMQP property

Raises `ValueError`

reply_to
Access the current message's `reply-to` AMQP message property as an attribute of the consumer class.

Return type str

returned
Indicates if the message was delivered by consumer previously and returned from RabbitMQ.

New in version 3.17.

Return type bool

routing_key
Access the routing key for the current message.

Return type str

send_exception_to_sentry(*exc_info*)
Send an exception to Sentry if enabled.

Parameters **exc_info** (*tuple*) – exception information as returned from `sys.exc_info()`

sentry_client
Access the Sentry raven `Client` instance or `None`

Use this object to add tags or additional context to Sentry error reports (see `raven.base.Client.tags_context()`) or to report messages (via `raven.base.Client.captureMessage()`) directly to Sentry.

Return type `raven.base.Client`

set_sentry_context (*tag*, *value*)
Set a context tag in Sentry for the given key and value.

Parameters

- **tag** (*str*) – The context tag name
- **value** (*str*) – The context value

settings
Access the consumer settings as specified by the `config` section for the consumer in the rejected configuration.

Return type dict

shutdown ()
Override to cleanly shutdown when rejected is stopping the consumer.

This could be used for closing database connections or other such activities.

stats_add_timing (*key*, *duration*)
Add a timing to the per-message measurements

New in version 3.13.0.

Parameters

- **key** (*str*) – The key to add the timing to
- **duration** (*int* | *float*) – The timing value

stats_incr (*key*, *value=1*)
Increment the specified key in the per-message measurements

New in version 3.13.0.

Parameters

- **key** (*str*) – The key to increment
- **value** (*int*) – The value to increment the key by

stats_set_tag (*key*, *value=1*)
Set the specified tag/value in the per-message measurements

New in version 3.13.0.

Parameters

- **key** (*str*) – The key to increment
- **value** (*int*) – The value to increment the key by

stats_set_value (*key*, *value=1*)
Set the specified key/value in the per-message measurements

New in version 3.13.0.

Parameters

- **key** (*str*) – The key to increment
- **value** (*int*) – The value to increment the key by

stats_track_duration (**args*, ***kwds*)
Time around a context and add to the the per-message measurements

New in version 3.13.0.

Parameters **key** (*str*) – The key for the timing to track

statsd_add_timing(*key*, *duration*)
Add a timing to the per-message measurements

Parameters

- **key** (*str*) – The key to add the timing to
- **duration** (*int | float*) – The timing value

Deprecated since version 3.13.0.

statsd_incr(*key*, *value=1*)
Increment the specified key in the per-message measurements

Parameters

- **key** (*str*) – The key to increment
- **value** (*int*) – The value to increment the key by

Deprecated since version 3.13.0.

statsd_track_duration(*key*)
Time around a context and add to the the per-message measurements

Parameters **key** (*str*) – The key for the timing to track

Deprecated since version 3.13.0.

timestamp
Access the unix epoch timestamp value from the AMQP message properties of the current message.

Return type int

unset_sentry_context(*tag*)
Remove a context tag from sentry

Parameters **tag** (*str*) – The context tag to remove

user_id
Access the `user-id` AMQP message property from the current message's properties.

Return type str

yield_to_ioloop(**args*, ***kwargs*)
Function that will allow Rejected to process IOLoop events while in a tight-loop inside an asynchronous consumer.

Exceptions

There are three exception types that consumer applications should raise to handle problems that may arise when processing a message. When these exceptions are raised, rejected will reject the message delivery, letting RabbitMQ know that there was a failure.

The *ConsumerException* should be raised when there is a problem in the consumer itself, such as inability to contact a database server or other resources. When a *ConsumerException* is raised, the message will be rejected *and* requeued, adding it back to the RabbitMQ it was delivered back to. Additionally, rejected keeps track of consumer exceptions and will shutdown the consumer process and start a new one once a consumer has exceeded its configured maximum error count within a `60` second window. The default maximum error count is `5`.

The *MessageException* should be raised when there is a problem with the message. When this exception is raised, the message will be rejected on the RabbitMQ server *without* requeue, discarding the message. This should be

done when there is a problem with the message itself, such as a malformed payload or non-supported properties like `content-type` or `type`.

If a consumer raises a *`ProcessingException`*, the message that was being processed will be republished to the exchange specified by the `error` exchange configuration value or the `ERROR_EXCHANGE` attribute of the consumer's class. The message will be published using the routing key that was last used for the message. The original message body and properties will be used and two additional header property values may be added:

- `X-Processing-Exception` contains the string value of the exception that was

raised, if specified. - `X-Processing-Exceptions` contains the quantity of processing exceptions that have been raised for the message.

In combination with a queue that has `x-message-ttl` set and `x-dead-letter-exchange` that points to the original exchange for the queue the consumer is consuming off of, you can implement a delayed retry cycle for messages that are failing to process due to external resource or service issues.

If `ERROR_MAX_RETRY` is set on the class, the headers for each method will be inspected and if the value of `X-Processing-Exceptions` is greater than or equal to the `ERROR_MAX_RETRY` value, the message will be dropped.

Note: If unhandled exceptions are raised by a consumer, they will be caught by rejected, logged, and turned into a *`ConsumerException`*.

---

class `rejected.consumer.`**`ConsumerException`**

May be called when processing a message to indicate a problem that the Consumer may be experiencing that should cause it to stop.

class `rejected.consumer.`**`MessageException`**

Invoke when a message should be rejected and not requeued, but not due to a processing error that should cause the consumer to stop.

class `rejected.consumer.`**`ProcessingException`**

Invoke when a message should be rejected and not requeued, but not due to a processing error that should cause the consumer to stop. This should be used for when you want to reject a message which will be republished to a retry queue, without anything being stated about the exception.

Testing Support

The *`rejected.testing.AsyncTestCase`* provides a based class for the easy creation of tests for your consumers. The test cases exposes multiple methods to make it easy to setup a consumer and process messages. It is build on top of `tornado.testing.AsyncTestCase` which extends `unittest.TestCase`.

To get started, override the *`rejected.testing.AsyncTestCase.get_consumer()`* method.

Next, the *`rejected.testing.AsyncTestCase.get_settings()`* method can be overridden to define the settings that are passed into the consumer.

Finally, to invoke your Consumer as if it were receiving a message, the *`process_message()`* method should be invoked.

Note: Tests are asynchronous, so each test should be decorated with *`gen_test()`*.

Example

The following example expects that when the message is processed by the consumer, the consumer will raise a `MessageException`.

```
from rejected import consumer, testing

import my_package


class ConsumerTestCase(testing.AsyncTestCase):

    def get_consumer(self):
        return my_package.Consumer

    def get_settings(self):
        return {'remote_url': 'http://foo'}

    @testing.gen_test
    def test_consumer_raises_message_exception(self):
        with self.assertRaises(consumer.MessageException):
            yield self.process_message({'foo': 'bar'})
```

class `rejected.testing.`**`AsyncTestCase`**(*methodName='runTest'*)

`TestCase` subclass for testing `IOLoop`-based asynchronous code.

`create_message`(*message*, *properties=None*, *exchange='rejected'*, *routing_key='test'*)

Create a message instance for use with the consumer in testing.

Parameters

- **`message`** (*`any`*) – the body of the message to create
- **`properties`** (*`dict`*) – AMQP message properties
- **`exchange`** (*`str`*) – The exchange the message should appear to be from
- **`routing_key`** (*`str`*) – The message's routing key

Return type `rejected.data.Message`

`get_consumer`()

Override to return the consumer class for testing.

Return type *`rejected.consumer.Consumer`*

`get_settings`()

Override this method to provide settings to the consumer during construction.

Return type dict

`process_message`(**args*, ***kwargs*)

Process a message as if it were being delivered by RabbitMQ. When invoked, an AMQP message will be locally created and passed into the consumer. With using the default values for the method, if you pass in a JSON serializable object, the message body will automatically be JSON serialized.

If an exception is not raised, a `Measurement` instance is returned that will contain all of the measurements collected during the processing of the message.

Parameters

- **`message_body`** (*`any`*) – the body of the message to create

- **content_type** (*str*) – The mime type
- **message_type** (*str*) – identifies the type of message to create
- **properties** (*dict*) – AMQP message properties
- **exchange** (*str*) – The exchange the message should appear to be from
- **routing_key** (*str*) – The message's routing key

Raises *rejected.consumer.ConsumerException*

Raises *rejected.consumer.MessageException*

Raises *rejected.consumer.ProcessingException*

Return type `rejected.data.Measurement`

`rejected.testing.`**gen_test** (*func=None*, *timeout=None*)

Testing equivalent of `tornado.gen.coroutine()`, to be applied to test methods.

CHAPTER 5

Issues

Please report any issues to the Github repo at https://github.com/gmr/rejected/issues

CHAPTER 6

Source

rejected source is available on Github at https://github.com/gmr/rejected

CHAPTER 7

Version History

See history

CHAPTER 8

Indices and tables

- genindex
- modindex
- search

Python Module Index

r

rejected.testing, 39

Index

A

app_id (rejected.consumer.Consumer attribute), 16
app_id (rejected.consumer.PublishingConsumer attribute), 22
app_id (rejected.consumer.SmartConsumer attribute), 27
app_id (rejected.consumer.SmartPublishingConsumer attribute), 33
AsyncTestCase (class in rejected.testing), 40

B

body (rejected.consumer.Consumer attribute), 16
body (rejected.consumer.PublishingConsumer attribute), 22
body (rejected.consumer.SmartConsumer attribute), 27
body (rejected.consumer.SmartPublishingConsumer attribute), 33

C

Consumer (class in rejected.consumer), 16
ConsumerException (class in rejected.consumer), 39
content_encoding (rejected.consumer.Consumer attribute), 17
content_encoding (rejected.consumer.PublishingConsumer attribute), 22
content_encoding (rejected.consumer.SmartConsumer attribute), 28
content_encoding (rejected.consumer.SmartPublishingConsumer attribute), 33
content_type (rejected.consumer.Consumer attribute), 17
content_type (rejected.consumer.PublishingConsumer attribute), 22
content_type (rejected.consumer.SmartConsumer attribute), 28
content_type (rejected.consumer.SmartPublishingConsumer attribute), 33
correlation_id (rejected.consumer.Consumer attribute), 17
correlation_id (rejected.consumer.PublishingConsumer attribute), 22
correlation_id (rejected.consumer.SmartConsumer attribute), 28
correlation_id (rejected.consumer.SmartPublishingConsumer attribute), 33
create_message() (rejected.testing.AsyncTestCase method), 40

E

exchange (rejected.consumer.Consumer attribute), 17
exchange (rejected.consumer.PublishingConsumer attribute), 22
exchange (rejected.consumer.SmartConsumer attribute), 28
exchange (rejected.consumer.SmartPublishingConsumer attribute), 33
expiration (rejected.consumer.Consumer attribute), 17
expiration (rejected.consumer.PublishingConsumer attribute), 22
expiration (rejected.consumer.SmartConsumer attribute), 28
expiration (rejected.consumer.SmartPublishingConsumer attribute), 33

F

finish() (rejected.consumer.Consumer method), 17
finish() (rejected.consumer.PublishingConsumer method), 22
finish() (rejected.consumer.SmartConsumer method), 28
finish() (rejected.consumer.SmartPublishingConsumer method), 33

G

gen_test() (in module rejected.testing), 41
get_consumer() (rejected.testing.AsyncTestCase method), 40
get_settings() (rejected.testing.AsyncTestCase method), 40

H

headers (rejected.consumer.Consumer attribute), 17

headers (rejected.consumer.PublishingConsumer attribute), 22
headers (rejected.consumer.SmartConsumer attribute), 28
headers (rejected.consumer.SmartPublishingConsumer attribute), 34

I

initialize() (rejected.consumer.Consumer method), 17
initialize() (rejected.consumer.PublishingConsumer method), 22
initialize() (rejected.consumer.SmartConsumer method), 28
initialize() (rejected.consumer.SmartPublishingConsumer method), 34

M

message_id (rejected.consumer.Consumer attribute), 17
message_id (rejected.consumer.PublishingConsumer attribute), 22
message_id (rejected.consumer.SmartConsumer attribute), 28
message_id (rejected.consumer.SmartPublishingConsumer attribute), 34
message_type (rejected.consumer.Consumer attribute), 17
message_type (rejected.consumer.PublishingConsumer attribute), 22
message_type (rejected.consumer.SmartConsumer attribute), 28
message_type (rejected.consumer.SmartPublishingConsumer attribute), 34
MessageException (class in rejected.consumer), 39

N

name (rejected.consumer.Consumer attribute), 17
name (rejected.consumer.PublishingConsumer attribute), 23
name (rejected.consumer.SmartConsumer attribute), 28
name (rejected.consumer.SmartPublishingConsumer attribute), 34

O

on_blocked() (rejected.consumer.Consumer method), 17
on_blocked() (rejected.consumer.PublishingConsumer method), 23
on_blocked() (rejected.consumer.SmartConsumer method), 28
on_blocked() (rejected.consumer.SmartPublishingConsumer method), 34
on_finish() (rejected.consumer.Consumer method), 18
on_finish() (rejected.consumer.PublishingConsumer method), 23
on_finish() (rejected.consumer.SmartConsumer method), 29
on_finish() (rejected.consumer.SmartPublishingConsumer method), 34
on_unblocked() (rejected.consumer.Consumer method), 18
on_unblocked() (rejected.consumer.PublishingConsumer method), 23
on_unblocked() (rejected.consumer.SmartConsumer method), 29
on_unblocked() (rejected.consumer.SmartPublishingConsumer method), 34

P

prepare() (rejected.consumer.Consumer method), 18
prepare() (rejected.consumer.PublishingConsumer method), 23
prepare() (rejected.consumer.SmartConsumer method), 29
prepare() (rejected.consumer.SmartPublishingConsumer method), 34
priority (rejected.consumer.Consumer attribute), 18
priority (rejected.consumer.PublishingConsumer attribute), 23
priority (rejected.consumer.SmartConsumer attribute), 29
priority (rejected.consumer.SmartPublishingConsumer attribute), 35
process() (rejected.consumer.Consumer method), 18
process() (rejected.consumer.PublishingConsumer method), 23
process() (rejected.consumer.SmartConsumer method), 29
process() (rejected.consumer.SmartPublishingConsumer method), 35
process_message() (rejected.testing.AsyncTestCase method), 40
ProcessingException (class in rejected.consumer), 39
properties (rejected.consumer.Consumer attribute), 18
properties (rejected.consumer.PublishingConsumer attribute), 24
properties (rejected.consumer.SmartConsumer attribute), 29
properties (rejected.consumer.SmartPublishingConsumer attribute), 35
publish_message() (rejected.consumer.Consumer method), 18
publish_message() (rejected.consumer.PublishingConsumer method), 24
publish_message() (rejected.consumer.SmartConsumer method), 30
publish_message() (rejected.consumer.SmartPublishingConsumer method), 35
PublishingConsumer (class in rejected.consumer), 21

R

redelivered (rejected.consumer.Consumer attribute), 19

redelivered (rejected.consumer.PublishingConsumer attribute), 24
redelivered (rejected.consumer.SmartConsumer attribute), 30
redelivered (rejected.consumer.SmartPublishingConsumer attribute), 35
rejected.testing (module), 39
reply() (rejected.consumer.Consumer method), 19
reply() (rejected.consumer.PublishingConsumer method), 24
reply() (rejected.consumer.SmartConsumer method), 30
reply() (rejected.consumer.SmartPublishingConsumer method), 36
reply_to (rejected.consumer.Consumer attribute), 19
reply_to (rejected.consumer.PublishingConsumer attribute), 24
reply_to (rejected.consumer.SmartConsumer attribute), 31
reply_to (rejected.consumer.SmartPublishingConsumer attribute), 36
returned (rejected.consumer.Consumer attribute), 19
returned (rejected.consumer.PublishingConsumer attribute), 24
returned (rejected.consumer.SmartConsumer attribute), 31
returned (rejected.consumer.SmartPublishingConsumer attribute), 36
routing_key (rejected.consumer.Consumer attribute), 19
routing_key (rejected.consumer.PublishingConsumer attribute), 25
routing_key (rejected.consumer.SmartConsumer attribute), 31
routing_key (rejected.consumer.SmartPublishingConsumer attribute), 36

S

send_exception_to_sentry() (rejected.consumer.Consumer method), 19
send_exception_to_sentry() (rejected.consumer.PublishingConsumer method), 25
send_exception_to_sentry() (rejected.consumer.SmartConsumer method), 31
send_exception_to_sentry() (rejected.consumer.SmartPublishingConsumer method), 36
sentry_client (rejected.consumer.Consumer attribute), 20
sentry_client (rejected.consumer.PublishingConsumer attribute), 25
sentry_client (rejected.consumer.SmartConsumer attribute), 31
sentry_client (rejected.consumer.SmartPublishingConsumer attribute), 36
set_sentry_context() (rejected.consumer.Consumer method), 20
set_sentry_context() (rejected.consumer.PublishingConsumer method), 25
set_sentry_context() (rejected.consumer.SmartConsumer method), 31
set_sentry_context() (rejected.consumer.SmartPublishingConsumer method), 36
settings (rejected.consumer.Consumer attribute), 20
settings (rejected.consumer.PublishingConsumer attribute), 25
settings (rejected.consumer.SmartConsumer attribute), 31
settings (rejected.consumer.SmartPublishingConsumer attribute), 37
shutdown() (rejected.consumer.Consumer method), 20
shutdown() (rejected.consumer.PublishingConsumer method), 25
shutdown() (rejected.consumer.SmartConsumer method), 31
shutdown() (rejected.consumer.SmartPublishingConsumer method), 37
SmartConsumer (class in rejected.consumer), 27
SmartPublishingConsumer (class in rejected.consumer), 33
stats_add_timing() (rejected.consumer.Consumer method), 20
stats_add_timing() (rejected.consumer.PublishingConsumer method), 25
stats_add_timing() (rejected.consumer.SmartConsumer method), 31
stats_add_timing() (rejected.consumer.SmartPublishingConsumer method), 37
stats_incr() (rejected.consumer.Consumer method), 20
stats_incr() (rejected.consumer.PublishingConsumer method), 25
stats_incr() (rejected.consumer.SmartConsumer method), 31
stats_incr() (rejected.consumer.SmartPublishingConsumer method), 37
stats_set_tag() (rejected.consumer.Consumer method), 20
stats_set_tag() (rejected.consumer.PublishingConsumer method), 26
stats_set_tag() (rejected.consumer.SmartConsumer method), 32
stats_set_tag() (rejected.consumer.SmartPublishingConsumer method), 37
stats_set_value() (rejected.consumer.Consumer method), 20
stats_set_value() (rejected.consumer.PublishingConsumer method), 26
stats_set_value() (rejected.consumer.SmartConsumer method), 32

stats_set_value() (rejected.consumer.SmartPublishingConsumer method), 37
stats_track_duration() (rejected.consumer.Consumer method), 21
stats_track_duration() (rejected.consumer.PublishingConsumer method), 26
stats_track_duration() (rejected.consumer.SmartConsumer method), 32
stats_track_duration() (rejected.consumer.SmartPublishingConsumer method), 37
statsd_add_timing() (rejected.consumer.Consumer method), 21
statsd_add_timing() (rejected.consumer.PublishingConsumer method), 26
statsd_add_timing() (rejected.consumer.SmartConsumer method), 32
statsd_add_timing() (rejected.consumer.SmartPublishingConsumer method), 38
statsd_incr() (rejected.consumer.Consumer method), 21
statsd_incr() (rejected.consumer.PublishingConsumer method), 26
statsd_incr() (rejected.consumer.SmartConsumer method), 32
statsd_incr() (rejected.consumer.SmartPublishingConsumer method), 38
statsd_track_duration() (rejected.consumer.Consumer method), 21
statsd_track_duration() (rejected.consumer.PublishingConsumer method), 26
statsd_track_duration() (rejected.consumer.SmartConsumer method), 32
statsd_track_duration() (rejected.consumer.SmartPublishingConsumer method), 38

T

timestamp (rejected.consumer.Consumer attribute), 21
timestamp (rejected.consumer.PublishingConsumer attribute), 26
timestamp (rejected.consumer.SmartConsumer attribute), 32
timestamp (rejected.consumer.SmartPublishingConsumer attribute), 38

U

unset_sentry_context() (rejected.consumer.Consumer method), 21
unset_sentry_context() (rejected.consumer.PublishingConsumer method), 26
unset_sentry_context() (rejected.consumer.SmartConsumer method), 33
unset_sentry_context() (rejected.consumer.SmartPublishingConsumer method), 38
user_id (rejected.consumer.Consumer attribute), 21
user_id (rejected.consumer.PublishingConsumer attribute), 26
user_id (rejected.consumer.SmartConsumer attribute), 33
user_id (rejected.consumer.SmartPublishingConsumer attribute), 38

Y

yield_to_ioloop() (rejected.consumer.Consumer method), 21
yield_to_ioloop() (rejected.consumer.PublishingConsumer method), 27
yield_to_ioloop() (rejected.consumer.SmartConsumer method), 33
yield_to_ioloop() (rejected.consumer.SmartPublishingConsumer method), 38