
Fedora Messaging

Release 3.5.0

Jeremy Cline

Apr 11, 2024

USER GUIDE

1	Installation	3
1.1	PyPI	3
1.2	Fedora	3
2	Quick Start	5
2.1	Local Broker	5
2.2	Fedora's Public Broker	6
2.3	Fedora's Restricted Broker	9
3	Configuration	11
3.1	Generic Options	13
3.2	Publisher Options	15
3.3	Consumer Options	16
4	Publishing	19
4.1	Overview	19
4.2	Introduction	19
4.3	Handling Errors	20
5	Message Schemas	21
5.1	Schema	21
5.2	Message Conventions	26
5.3	Packaging	27
5.4	Upgrade and deprecation	27
6	Consumers	29
6.1	Introduction	29
6.2	Command Line Interface	30
6.3	Consumer API	30
6.4	systemd Service	32
7	Available Schemas	33
7.1	anitya	33
7.2	bodhi	34
7.3	Copr	35
7.4	fedocal	35
7.5	elections	36
7.6	Git	36
7.7	The New Hotness	36
7.8	planet	36
7.9	ansible	37

7.10	FMN	37
7.11	kerneltest	37
7.12	Koji	37
7.13	Koschei	38
7.14	mdapi	38
7.15	Wiki	38
7.16	meetbot	38
7.17	FAS	38
7.18	nuancier	39
7.19	Pagure	39
7.20	tahrir	41
8	Testing	43
9	Command Line Interface Manuals	45
9.1	fedora-messaging	45
10	Installation	51
10.1	Installing the library	51
10.2	Setting up RabbitMQ	51
10.3	Configuration	52
11	Using the API	53
11.1	Publishing	53
11.2	Listening	54
12	JSON schemas	57
12.1	Creating the schema package	57
12.2	Writing the schema	57
12.3	Testing it	59
12.4	Using it	59
12.5	Updating it	60
13	Handling exceptions	61
13.1	When publishing	61
13.2	When consuming	61
14	Converting a fedmsg application	63
14.1	Converting publishers	63
14.2	Converting consumers	68
15	Developer Interface	71
15.1	Publishing	72
15.2	Subscribing	73
15.3	Signals	75
15.4	Message Schemas	76
15.5	Utilities	83
15.6	Exceptions	84
15.7	Configuration	86
16	Message Format	87
16.1	Content Type	87
16.2	Content Encoding	87
16.3	Message ID	87
16.4	Delivery Mode	87

16.5	Headers	88
16.6	Body	89
17	Contributor guide	91
17.1	Quickstart	91
17.2	Python Support	91
17.3	Code Style	91
17.4	Tests	92
17.5	Release notes	92
17.6	Licensing	92
17.7	Releasing	93
18	Release Notes	95
18.1	3.5.0 (2024-03-20)	95
18.2	3.4.1 (2023-05-26)	96
18.3	3.4.0 (2023-05-26)	96
18.4	3.3.0 (2023-03-31)	96
18.5	3.2.0 (2022-10-17)	97
18.6	3.1.0 (2022-09-13)	97
18.7	3.0.2 (2022-05-19)	97
18.8	3.0.1 (2022-05-12)	98
18.9	3.0.0 (2021-12-14)	98
18.10	2.1.0 (2021-05-12)	99
18.11	2.0.2 (2020-08-04)	100
18.12	2.0.1 (2020-01-02)	101
18.13	2.0.0 (2019-12-03)	101
18.14	1.7.2 (2019-08-02)	102
18.15	v1.7.1 (2019-06-24)	103
18.16	v1.7.0 (2019-05-21)	104
18.17	v1.6.1 (2019-04-17)	104
18.18	v1.6.0 (2019-04-04)	105
18.19	v1.5.0 (2019-02-28)	106
18.20	v1.4.0 (2019-02-07)	106
18.21	v1.3.0 (2019-01-24)	107
18.22	v1.2.0 (2019-01-21)	107
18.23	v1.1.0 (2018-11-13)	108
18.24	v1.0.1 (2018-10-10)	108
18.25	v1.0.0 (2018-10-10)	108
18.26	v1.0.0b1	109
18.27	v1.0.0a1	110
	Python Module Index	111
	Index	113

This package provides tools and APIs to make using Fedora’s messaging infrastructure easier. These include a framework for declaring message schemas, a set of synchronous APIs to publish messages to AMQP brokers, a set of asynchronous APIs to consume messages, and services to easily run consumers.

This library is designed to be a replacement for the [PyZMQ](#)-backed [fedmsg](#) library.

INSTALLATION

1.1 PyPI

The Python package is available on the [Python Package Index \(PyPI\)](#) as `fedora-messaging`:

```
$ pip install --user fedora-messaging
```

It is, of course, recommended that you install it in a Python virtual environment.

1.2 Fedora

The library is available in Fedora 29 and greater as `fedora-messaging`:

```
$ sudo dnf install fedora-messaging
```


QUICK START

This is a quick-start guide that covers a few common use-cases and contains pointers to more in-depth documentation for the curious.

2.1 Local Broker

To publish and consume messages locally can be a useful way to learn about the library, and is also helpful during development of your application or service.

To install the message broker on Fedora:

```
$ sudo dnf install rabbitmq-server
```

RabbitMQ is also available in EPEL7, although it is quite old and the library is not regularly tested against it. You can also install the broker from RabbitMQ directly if you are not using Fedora.

Next, it's recommended that you enable the management interface:

```
$ sudo rabbitmq-plugins enable rabbitmq_management
```

This provides an HTTP interface and API, available at <http://localhost:15672/> by default. The “guest” user with the password “guest” is created by default.

Finally, start the broker:

```
$ sudo systemctl start rabbitmq-server
```

You should now be able to consume messages with the following Python script:

```
from fedora_messaging import api, config

config.conf.setup_logging()
api.consume(lambda message: print(message))
```

To learn more about consuming messages, check out the *Consumers* documentation.

You can publish messages with:

```
from fedora_messaging import api, config

config.conf.setup_logging()
api.publish(api.Message(topic="hello", body={"Hello": "world!"}))
```

To learn more about publishing messages, check out the *Publishing* documentation.

2.2 Fedora's Public Broker

Fedora's message broker has a publicly accessible virtual host located at `amqps://rabbitmq.fedoraproject.org/%2Fpublic_pubsub`. This virtual host mirrors all messages published to the restricted `/pubsub` virtual host and allows anyone to consume messages being published by the various Fedora services.

These public queues have some restrictions applied to them. Firstly, they are limited to about 50 megabytes in size, so if your application cannot handle the message throughput messages will be automatically discarded once you hit this limit. Secondly, queues that are set to be durable (in other words, not exclusive or auto-deleted) are automatically deleted if they have no consumers after approximately an hour.

If you need more robust guarantees about message delivery, or if you need to publish messages into Fedora's message broker, contact the Fedora Infrastructure team about getting access to the private virtual host.

2.2.1 Getting Connected

The public virtual host still requires users to authenticate when connecting, so a public user has been created and its private key and x509 certificate are distributed with `fedora-messaging`.

If `fedora-messaging` was installed via RPM, they should be in `/etc/fedora-messaging/` along with a configuration file called `fedora.toml`. If it's been installed via pip, it's easiest to get the [key](#), [certificate](#), and the [CA certificate](#) from the upstream git repository and start with the following configuration file:

```
# A basic configuration for Fedora's message broker, using the example callback
# which simply prints messages to standard output.
#
# This file is in the TOML format.
amqp_url = "amqps://fedora:~@rabbitmq.fedoraproject.org/%2Fpublic_pubsub"
callback = "fedora_messaging.example:printer"

[tls]
ca_cert = "/etc/fedora-messaging/cacert.pem"
keyfile = "/etc/fedora-messaging/fedora-key.pem"
certfile = "/etc/fedora-messaging/fedora-cert.pem"

[client_properties]
app = "Example Application"
# Some suggested extra fields:
# URL of the project that provides this consumer
app_url = "https://github.com/fedora-infra/fedora-messaging"
# Contact emails for the maintainer(s) of the consumer - in case the
# broker admin needs to contact them, for e.g.
app_contacts_email = ["admin@fedoraproject.org"]

[exchanges."amq.topic"]
type = "topic"
durable = true
auto_delete = false
arguments = {}

# Queue names *must* be in the normal UUID format: run "uuidgen" and use the
# output as your queue name. If you don't define a queue here, the server will
# generate a queue name for you. This queue will be non-durable, auto-deleted and
```

(continues on next page)

(continued from previous page)

```

# exclusive.
# If your queue is not exclusive, anyone can connect and consume from it, causing
# you to miss messages, so do not share your queue name. Any queues that are not
# auto-deleted on disconnect are garbage-collected after approximately one hour.
#
# If you require a stronger guarantee about delivery, please talk to Fedora's
# Infrastructure team.
#
# [queues.000000000-0000-0000-0000-000000000000]
# durable = false
# auto_delete = true
# exclusive = true
# arguments = {}

# If you use the server-generated queue names, you can leave out the "queue"
# parameter in the bindings definition.
[[bindings]]
# queue = "000000000-0000-0000-0000-000000000000"
exchange = "amq.topic"
routing_keys = ["#"] # Set this to the specific topics you are interested in.

[consumer_config]
example_key = "for my consumer"

[qos]
prefetch_size = 0
prefetch_count = 25

[log_config]
version = 1
disable_existing_loggers = true

[log_config.formatters.simple]
format = "[%{levelname}s %(name)s] %(message)s"

[log_config.handlers.console]
class = "logging.StreamHandler"
formatter = "simple"
stream = "ext://sys.stdout"

[log_config.loggers.fedora_messaging]
level = "INFO"
propagate = false
handlers = ["console"]

[log_config.loggers.twisted]
level = "INFO"
propagate = false
handlers = ["console"]

[log_config.loggers.pika]
level = "WARNING"

```

(continues on next page)

(continued from previous page)

```

propagate = false
handlers = ["console"]

# If your consumer sets up a logger, you must add a configuration for it
# here in order for the messages to show up. e.g. if it set up a logger
# called 'example_printer', you could do:
#[log_config.loggers.example_printer]
#level = "INFO"
#propagate = false
#handlers = ["console"]

[log_config.root]
level = "ERROR"
handlers = ["console"]

```

Assuming the `/etc/fedora-messaging/fedora.toml`, `/etc/fedora-messaging/cacert.pem`, `/etc/fedora-messaging/fedora-key.pem`, and `/etc/fedora-messaging/fedora-cert.pem` files exist, the following command will create a configuration file called `my_config.toml` with a unique queue name for your consumer:

```

$ sed -e "s/[0-9a-f]\{8\}-[0-9a-f]\{4\}-[0-9a-f]\{4\}-[0-9a-f]\{4\}-[0-9a-f]\{12\}/
→$(uuidgen)/g" \
    /etc/fedora-messaging/fedora.toml > my_config.toml

```

Warning: Do not skip the step above. This is important because if there are multiple consumers on a queue the broker delivers messages to them in a round-robin fashion. In other words, you'll only get some of the messages being sent.

Run a quick test to make sure you can connect to the broker. The configuration file comes with an example consumer which simply prints the message to standard output:

```
$ fedora-messaging --conf my_config.toml consume
```

Alternatively, you can start a Python shell and use the API:

```

$ FEDORA_MESSAGING_CONF=my_config.toml python
>>> from fedora_messaging import api, config
>>> config.conf.setup_logging()
>>> api.consume(lambda message: print(message))

```

If all goes well, you'll see a log entry similar to:

```

Successfully registered AMQP consumer Consumer(queue=af0f78d2-159e-4279-b404-
→7b8c1b4649cc, callback=<function printer at 0x7f9a59e077b8>)

```

This will be followed by the messages being sent inside Fedora's Infrastructure. All that's left to do is change the callback in the configuration to use your consumer *callback* and adjusting the routing keys in your *bindings* to receive only the messages your consumer is interested in.

2.3 Fedora's Restricted Broker

Connecting the Fedora's private virtual host requires working with the Fedora infrastructure team. The current process and configuration for this is documented in the [infrastructure team's development guide](#).

Note: It is essential to configure the `passive_declares` option correctly in the `/etc/fedora-messaging/config.toml` file. This setting is mandatory for users of Fedora's `/pubsub` vhost and should be set to `true`. It controls the declaration of queues and exchanges.

CONFIGURATION

fedora-messaging can be configured with the `/etc/fedora-messaging/config.toml` file or by setting the `FEDORA_MESSAGING_CONF` environment variable to the path of the configuration file.

Each configuration option has a default value.

Table of Configuration Options

- *Generic Options*
 - *amqp_url*
 - *passive_declares*
 - *tls*
 - *client_properties*
 - *exchanges*
 - *log_config*
- *Publisher Options*
 - *publish_exchange*
 - *topic_prefix*
 - *publish_priority*
- *Consumer Options*
 - *queues*
 - *bindings*
 - *callback*
 - *consumer_config*
 - *qos*

A complete example TOML configuration:

```
# A sample configuration for fedora-messaging. This file is in the TOML format.
amqp_url = "amqp://"
callback = "fedora_messaging.example:printer"
passive_declares = false
publish_exchange = "amq.topic"
```

(continues on next page)

(continued from previous page)

```
topic_prefix = ""

[tls]
ca_cert = "/etc/fedora-messaging/cacert.pem"
keyfile = "/etc/fedora-messaging/fedora-key.pem"
certfile = "/etc/fedora-messaging/fedora-cert.pem"

[client_properties]
app = "Example Application"

# If the exchange or queue name has a "." in it, use quotes as seen here.
[exchanges."amq.topic"]
type = "topic"
durable = true
auto_delete = false
arguments = {}

[queues.my_queue]
durable = true
auto_delete = false
exclusive = false
arguments = {}

# Note the double brackets below. To add another binding, add another
# [[bindings]] section. To use multiple routing keys, just expand the list here.
[[bindings]]
queue = "my_queue"
exchange = "amq.topic"
routing_keys = ["#"]

[consumer_config]
example_key = "for my consumer"

[qos]
prefetch_size = 0
prefetch_count = 25

[log_config]
version = 1
disable_existing_loggers = true

[log_config.formatters.simple]
format = "[%{levelname}s %{name}s] %(message)s"

[log_config.handlers.console]
class = "logging.StreamHandler"
formatter = "simple"
stream = "ext://sys.stdout"

[log_config.loggers.fedora_messaging]
level = "INFO"
propagate = false
```

(continues on next page)

(continued from previous page)

```
handlers = ["console"]

# Twisted is the asynchronous framework that manages the TCP/TLS connection, as well
# as the consumer event loop. When debugging you may want to lower this log level.
[log_config.loggers.twisted]
level = "INFO"
propagate = false
handlers = ["console"]

# Pika is the underlying AMQP client library. When debugging you may want to
# lower this log level.
[log_config.loggers.pika]
level = "WARNING"
propagate = false
handlers = ["console"]

[log_config.root]
level = "ERROR"
handlers = ["console"]
```

3.1 Generic Options

These options apply to both consumers and publishers.

3.1.1 amqp_url

The AMQP broker to connect to. This URL should be in the format described by the [pika.connection.URLParameters](#) documentation. This defaults to 'amqp:///connection_attempts=3&retry_delay=5'.

Note: When using the Twisted consumer API, which the CLI does by default, any connection-related setting won't apply as Twisted manages the TCP/TLS connection.

3.1.2 `passive_declares`

A boolean to specify if queues and exchanges should be declared passively (i.e checked, but not actually created on the server). Defaults to `False`.

3.1.3 `tls`

A dictionary of the TLS settings to use when connecting to the AMQP broker. The default is:

```
{
  'ca_cert': '/etc/pki/tls/certs/ca-bundle.crt',
  'keyfile': None,
  'certfile': None,
}
```

The value of `ca_cert` should be the path to a bundle of CA certificates used to validate the certificate presented by the server. The `'keyfile'` and `'certfile'` values should be to the client key and client certificate to use when authenticating with the broker.

Note: The broker URL must use the `amqps` scheme. It is also possible to provide these setting via the `amqp_url` setting using a URL-encoded JSON object. This setting is provided as a convenient way to avoid that.

3.1.4 `client_properties`

A dictionary that describes the client to the AMQP broker. This makes it easy to identify the application using a connection. The dictionary can contain arbitrary string keys and values. The default is:

```
{
  'app': 'Unknown',
  'product': 'Fedora Messaging with Pika',
  'information': 'https://fedora-messaging.readthedocs.io/en/stable/',
  'version': 'fedora_messaging-<version> with pika-<version>',
}
```

Apps should set the `app` along with any additional keys they feel will help administrators when debugging application connections. At a minimum, the recommended fields are:

- `app_url`: The value of this key should be a URL to the upstream project for the client.
- `app_contacts_email`: One or more emails of maintainers to contact with questions (if, for example, a client is misbehaving, or a service disruption is about to occur).

Do not use the `product`, `information`, and `version` keys as these will be set automatically.

3.1.5 exchanges

A dictionary of exchanges that should be present in the broker. Each key should be an exchange name, and the value should be a dictionary with the exchange's configuration. Options are:

- `type` - the type of exchange to create.
- `durable` - whether or not the exchange should survive a broker restart.
- `auto_delete` - whether or not the exchange should be deleted once no queues are bound to it.
- `arguments` - dictionary of arbitrary keyword arguments for the exchange, which depends on the broker in use and its extensions.

For example:

```
{
  'my_exchange': {
    'type': 'fanout',
    'durable': True,
    'auto_delete': False,
    'arguments': {},
  },
}
```

The default is to ensure the `'amq.topic'` topic exchange exists which should be sufficient for most use cases.

3.1.6 log_config

A dictionary describing the logging configuration to use, in a format accepted by `logging.config.dictConfig()`.

Note: Logging is only configured for consumers, not for producers.

3.2 Publisher Options

The following configuration options are publisher-related.

3.2.1 publish_exchange

A string that identifies the exchange to publish to. The default is `amq.topic`.

3.2.2 topic_prefix

A string that will be prepended to topics on sent messages. This is useful to migrate from `fedmsg`, but should not be used otherwise. The default is an empty string.

3.2.3 publish_priority

A number that will be set as the priority for the messages. The range of possible priorities depends on the `x-max-priority` argument of the destination queue, as described in [RabbitMQ's priority documentation](#). The default is `None`, which RabbitMQ will interpret as zero.

3.3 Consumer Options

The following configuration options are consumer-related.

3.3.1 queues

A dictionary of queues that should be present in the broker. Each key should be a queue name, and the value should be a dictionary with the queue's configuration. Options are:

- `durable` - whether or not the queue should survive a broker restart. This is set to `False` for the default queue.
- `auto_delete` - whether or not the queue should be deleted once the consumer disconnects. This is set to `True` for the default queue.
- `exclusive` - whether or not the queue is exclusive to the current connection. This is set to `False` for the default queue.
- `arguments` - dictionary of arbitrary keyword arguments for the queue, which depends on the broker in use and its extensions. This is set to `{}` for the default queue

For example:

```
{
  'my_queue': {
    'durable': True,
    'auto_delete': True,
    'exclusive': False,
    'arguments': {},
  },
}
```

3.3.2 bindings

A list of dictionaries that define queue bindings to exchanges that consumers will subscribe to. The `queue` key is the queue's name. The `exchange` key should be the exchange name and the `routing_keys` key should be a list of routing keys. For example:

```
[
  {
    'queue': 'my_queue',
    'exchange': 'amq.topic',
    'routing_keys': ['topic1', 'topic2.#'],
  },
]
```

This would create two bindings for the `my_queue` queue, both to the `amq.topic` exchange. Consumers will consume from both queues.

3.3.3 callback

The Python path of the callback. This should be in the format `<module>:<object>`. For example, if the callback was called “my_callback” and was located in the “my_module” module of the “my_package” package, the path would be defined as `my_package.my_module:my_callback`. The default is `None`.

Consult the *Consumers* documentation for details on implementing a callback.

3.3.4 consumer_config

A dictionary for the consumer to use as configuration. The consumer should access this key in its callback for any configuration it needs. Defaults to an empty dictionary. If, for example, this dictionary contains the `print_messages` key, the callback can access this configuration with:

```
from fedora_messaging import config

def callback(message):
    if config.conf["consumer_config"]["print_messages"]:
        print(message)
```

3.3.5 qos

The quality of service settings to use for consumers. This setting is a dictionary with two keys. `prefetch_count` specifies the number of messages to pre-fetch from the server. Pre-fetching messages improves performance by reducing the amount of back-and-forth between client and server. The downside is if the consumer encounters an unexpected problem, messages won't be returned to the queue and sent to a different consumer until the consumer times out. `prefetch_size` limits the size of pre-fetched messages (in bytes), with 0 meaning there is no limit. The default settings are:

```
{
    'prefetch_count': 10,
    'prefetch_size': 0,
}
```


PUBLISHING

4.1 Overview

Publishing messages is simple. Messages are made up of a topic, some optional headers, and a body. Messages are encapsulated in a `fedora_messaging.message.Message` object. For details on defining messages, see the *Message Schemas* documentation. For details on the publishing API, see the *Publishing* API documentation.

4.1.1 Topics

Topics are strings of words separated by the `.` character, up to 255 characters. Topics are used by clients to filter messages, so choosing a good topic helps reduce the number of messages sent to a client. Topics should start broadly and become more specific.

4.1.2 Headers

Headers are key-value pairs attached that are useful for storing information about the message itself. This library adds a header to every message with the `fedora_messaging_schema` key, pointing to the message schema used.

You should not use any key starting with `fedora_messaging` for yourself.

You can write *Header Schema* for your messages to enforce a particular schema.

4.1.3 Body

The only restrictions on the message body is that it must be serializable to a JSON object. You should write a *Body Schema* for your messages to ensure you don't change your message format unintentionally.

4.2 Introduction

To publish a message, first create a `fedora_messaging.message.Message` object, then pass it to the `fedora_messaging.api.publish()` function:

```
from fedora_messaging import api, message

msg = message.Message(topic=u'nice.message', headers={u'niceness': u'very'},
                      body={u'encouragement': u"You're doing great!"})
api.publish(msg)
```

The API relies on the *Configuration* you've provided to connect to the message broker and publish the message to an exchange.

4.3 Handling Errors

Your message might fail to publish for a number of reasons, so you should be prepared to see (and potentially handle) some errors.

4.3.1 Validation

The message you create may not be successfully validated against its schema. This is not an error you should catch, since it must be fixed by the developer and cannot be recovered from.

4.3.2 Connection Errors

The publish API will attempt to reconnect to the broker several times before an exception is raised. Once this occurs it is up to the application to decide what to do.

4.3.3 Rejected Messages

The broker may reject a message. This could occur because the message is too large, or because the publisher does not have permission to publish messages with a particular topic, or some other reason.

MESSAGE SCHEMAS

Before you release your application, you should create a subclass of *fedora_messaging.message.Message*, define a schema, define a default severity, and implement some methods.

5.1 Schema

Defining a message schema is important for several reasons.

First and foremost, it will help you (the developer) ensure you don't accidentally change your message's format. When messages are being generated from, say, a database object, it's easy to make a schema change to the database and unintentionally alter your message, which breaks consumers of your message. Without a schema, you might not catch this until you deploy your application and consumers start crashing. With a schema, you'll get an error as you develop!

Secondly, it allows you to change your message format in a controlled fashion by versioning your schema. You can then choose to implement methods one way or another based on the version of the schema used by a message. For details on how to deprecate and upgrade message schemas, see *Upgrade and deprecation*.

Message schema are defined using *JSON Schema*. The complete API can be found in the *Message Schemas* API documentation.

5.1.1 Header Schema

The default header schema declares that the header field must be a JSON object with several expected keys. You can leave the schema as-is when you define your own message, or you can refine it. The base schema will always be enforced in addition to your custom schema.

5.1.2 Body Schema

The default body schema simply declares that the header field must be a JSON object.

5.1.3 Example Schema

```
# Copyright (C) 2018 Red Hat, Inc.
#
# This program is free software; you can redistribute it and/or modify
# it under the terms of the GNU General Public License as published by
# the Free Software Foundation; either version 2 of the License, or
# (at your option) any later version.
#
# This program is distributed in the hope that it will be useful,
# but WITHOUT ANY WARRANTY; without even the implied warranty of
# MERCHANTABILITY or FITNESS FOR A PARTICULAR PURPOSE. See the
# GNU General Public License for more details.
#
# You should have received a copy of the GNU General Public License along
# with this program; if not, write to the Free Software Foundation, Inc.,
# 51 Franklin Street, Fifth Floor, Boston, MA 02110-1301 USA.
"""This is an example of a message schema."""

from email.utils import parseaddr

from fedora_messaging import message

class BaseMessage(message.Message):
    """
    You should create a super class that each schema version inherits from.
    This lets consumers perform ``isinstance(msg, BaseMessage)`` if they are
    receiving multiple message types and allows the publisher to change the
    schema as long as they preserve the Python API.
    """

    def __str__(self):
        """Return a complete human-readable representation of the message."""
        return f"Subject: {self.subject}\n{self.email_body}\n"

    @property
    def summary(self):
        """Return a summary of the message.

        By convention, in Fedora all schemas should provide this property.
        """
        return self.subject

    @property
    def subject(self):
        """The email's subject."""
        return 'Message did not implement "subject" property'

    @property
    def email_body(self):
        """The email message body."""
        return 'Message did not implement "email_body" property'
```

(continues on next page)

(continued from previous page)

```

@property
def url(self):
    """An URL to the email in HyperKitty

    By convention, in Fedora all schemas should provide this property.

    Returns:
        str or None: A relevant URL.
    """
    base_url = "https://lists.fedoraproject.org/archives"
    archived_at = self._get_archived_at()
    if archived_at and archived_at.startswith("<"):
        archived_at = archived_at[1:]
    if archived_at and archived_at.endswith(">"):
        archived_at = archived_at[:-1]
    if archived_at and archived_at.startswith("http"):
        return archived_at
    elif archived_at:
        return base_url + archived_at
    else:
        return None

@property
def app_name(self):
    """The name of the application that generated the message.

    By convention, in Fedora all schemas should provide this property.
    """
    return "Mailman"

@property
def app_icon(self):
    """A URL to the icon of the application that generated the message.

    By convention, in Fedora all schemas should provide this property.
    """
    return "https://apps.fedoraproject.org/img/icons/hyperkitty.png"

@property
def usernames(self):
    """List of users affected by the action that generated this message."""
    return []

@property
def packages(self):
    """List of packages affected by the action that generated this message."""
    return []

def _get_username_from_from_header(self, from_header):
    """Converts a From email header to a username."""
    # Extract the username

```

(continues on next page)

(continued from previous page)

```

        addr = parseaddr(from_header)[1]
        return addr.split("@")[0]

class MessageV1(BaseMessage):
    """
    A sub-class of a Fedora message that defines a message schema for messages
    published by Mailman when it receives mail to send out.
    """

    body_schema = {
        "id": "http://fedoraproject.org/message-schema/mailman#",
        "$schema": "http://json-schema.org/draft-04/schema#",
        "description": "Schema for message sent to mailman",
        "type": "object",
        "properties": {
            "mlist": {
                "type": "object",
                "properties": {
                    "list_name": {
                        "type": "string",
                        "description": "The name of the mailing list",
                    }
                }
            },
            "msg": {
                "description": "An object representing the email",
                "type": "object",
                "properties": {
                    "delivered-to": {"type": "string"},
                    "from": {"type": "string"},
                    "cc": {"type": "string"},
                    "to": {"type": "string"},
                    "x-mailman-rule-hits": {"type": "string"},
                    "x-mailman-rule-misses": {"type": "string"},
                    "x-message-id-hash": {"type": "string"},
                    "references": {"type": "string"},
                    "in-reply-to": {"type": "string"},
                    "message-id": {"type": "string"},
                    "archived-at": {"type": "string"},
                    "subject": {"type": "string"},
                    "body": {"type": "string"},
                },
                "required": ["from", "to", "subject", "body"],
            },
        },
        "required": ["mlist", "msg"],
    }

    @property
    def subject(self):
        """The email's subject."""

```

(continues on next page)

(continued from previous page)

```

        return self.body["msg"]["subject"]

@property
def email_body(self):
    """The email message body."""
    return self.body["msg"]["body"]

@property
def agent_name(self):
    """The username of the user who caused the action."""
    return self._get_username_from_from_header(self.body["msg"]["from"])

def _get_archived_at(self):
    return self.body["msg"]["archived-at"]

class MessageV2(BaseMessage):
    """
    This is a revision from the MessageV1 schema which flattens the message
    structure into a single object, but is backwards compatible for any users
    that make use of the properties (`subject` and `body`).
    """

    body_schema = {
        "id": "http://fedoraproject.org/message-schema/mailman#",
        "$schema": "http://json-schema.org/draft-04/schema#",
        "description": "Schema for message sent to mailman",
        "type": "object",
        "required": ["mailing_list", "from", "to", "subject", "body"],
        "properties": {
            "mailing_list": {
                "type": "string",
                "description": "The name of the mailing list",
            },
            "delivered-to": {"type": "string"},
            "from": {"type": "string"},
            "cc": {"type": "string"},
            "to": {"type": "string"},
            "x-mailman-rule-hits": {"type": "string"},
            "x-mailman-rule-misses": {"type": "string"},
            "x-message-id-hash": {"type": "string"},
            "references": {"type": "string"},
            "in-reply-to": {"type": "string"},
            "message-id": {"type": "string"},
            "archived-at": {"type": "string"},
            "subject": {"type": "string"},
            "body": {"type": "string"},
        },
    }

@property
def subject(self):

```

(continues on next page)

(continued from previous page)

```
        """The email's subject."""
        return self.body["subject"]

    @property
    def email_body(self):
        """The email message body."""
        return self.body["body"]

    @property
    def agent_name(self):
        """The username of the user who caused the action."""
        return self._get_username_from_from_header(self.body["from"])

    def _get_archived_at(self):
        return self.body["archived-at"]
```

Note that message schema can be composed of other message schema, and validation of fields can be much more detailed than just a simple type check. Consult the [JSON Schema](#) documentation for complete details.

5.2 Message Conventions

5.2.1 Schema are Immutable

Message schema should be treated as immutable. Once defined, they should not be altered. Instead, define a new schema class, mark the old one as deprecated, and remove it after an appropriate transition period.

5.2.2 Provide Accessors

The JSON schema ensures the message sent “on the wire” conforms to a particular format. Messages should provide Python properties to access the deserialized JSON object. This Python API should maintain backwards compatibility between schema. This shields consumers from changes in schema.

Useful Accessors

All available accessors are described in the [Message Schemas](#) API documentation ; here is a list of those we recommend implementing to allow users to get notifications for your messages:

- `__str__()`: A human-readable representation of this message. This can be a multi-line string that forms the body of an email notification.
- `summary`: A short, single-line, human-readable summary of the message, much like the subject line of an email.
- `agent_name`: The username of the user who caused the action.
- `app_name`: The name of the application that generated the message. This can be implemented as a class attribute or as a property.
- `app_icon`: A URL to the icon of the application that generated the message. This can be implemented as a class attribute or as a property.
- `packages`: A list of RPM packages affected by the action that generated this message, if any.
- `flatpaks`: A list of flatpaks affected by the action that generated this message, if any.

- **modules**: A list of modules affected by the action that generated this message, if any.
- **containers**: A list of containers affected by the action that generated this message, if any.
- **usernames**: A list of usernames affected by the action that generated this message. This may contain the `agent_name`.
- **groups**: A list of group names affected by the action that generated this message.
- **url**: A URL to the action that caused this message to be emitted, if any.
- **severity**: An integer that indicates the severity of the message. This is used to determine what messages to notify end users about and should be `DEBUG`, `INFO`, `WARNING`, or `ERROR`. The default is `INFO`, and can be set as a class attribute or on an instance-by-instance basis.

5.3 Packaging

Finally, you must distribute your schema to clients. It is recommended that you maintain your message schema in your application's git repository in a separate Python package. The package name should be `<your-app-name>-messages`.

A complete sample schema package can be found in [the fedora-messaging repository](#). This includes unit tests, the schema classes, and a `setup.py`. You can adapt this boilerplate with the following steps:

- Change the package name from `mailman_messages` to `<your-app-name>_messages` in `setup.py`.
- Rename the `mailman_messages` directory to `<your-app-name>_messages`.
- Add your schema classes to `messages.py` and tests to `tests/test_messages.py`.
- Update the `README` file.
- Build the distribution with `python setup.py sdist bdist_wheel`.
- Upload the sdist and wheel to PyPI with `twine`.
- Submit an RPM package for it to Fedora and EPEL.

If you prefer [CookieCutter](#), there is a [template repository](#) that you can use with the command:

```
cookiecutter gh:fedora-infra/cookiecutter-message-schemas
```

It will ask you for the application name and some other variables, and will create the package structure for you.

5.4 Upgrade and deprecation

Message schema classes should not be modified in a backwards-incompatible fashion. To facilitate the evolution of schemas, we recommend including the schema version in the topic itself, such as `myservice.myevent.v1`.

When a backwards-incompatible change is required, create a new class with the topic ending in `.v2`, set the `Message.deprecated` attribute to `True` on the old class, and send both versions for a reasonable period of time. Note that you need to add the new class to the schema package's entry points as well.

We leave the duration to the developer's appreciation, since it depends on how many different consumers they expect to have, whether they are only inside the Fedora infrastructure or outside too, etc. This duration can range from weeks to months, possibly a year. At the time of this writing, Fedora's message bus is very far from being overwhelmed by messages, so you don't need to worry about that.

Proceeding this way ensures that consumers subscribing to `.v1` will not break when `.v2` arrives, and can choose to subscribe to the `.v2` topic when they are ready to handle the new format. They will get a warning in their logs when they receive deprecated messages, prompting them to upgrade.

When you add the new version, please upgrade the major version number of your schema package, and communicate clearly that the old version is deprecated, including for how long you have decided to send both versions.

CONSUMERS

This library is aimed at making implementing a message consumer as simple as possible by implementing common boilerplate code and offering a command line interface to easily start a consumer as a service under init systems like systemd.

6.1 Introduction

AMQP consumers configure a queue for their use in the message broker. When a message is published to an exchange and matches the bindings the consumer has declared, the message is placed in the queue and eventually delivered to the consumer. Fedora uses a topic exchange for general-purpose messages.

Fortunately, you don't need to manage the connection to the broker or configure the queue. All you need to do is to implement some code to run when a message is received. The API expects a callable object that accepts a single positional argument:

```
from fedora_messaging import api, config

# The fedora_messaging API does not automatically configure logging so as
# to not destroy application logging setup. This is a convenience method
# to configure the Python logger with the fedora-messaging logging config.
config.conf.setup_logging()

# First, define a function to be used as our callback. This will be called
# whenever a message is received from the server.
def printer_callback(message):
    """
    Print the message to standard output.

    Args:
        message (fedora_messaging.message.Message): The message we received
            from the queue.
    """
    print(str(message))

# Next, we need a queue to consume messages from. We can define
# the queue and binding configurations in these dictionaries:
queues = {
    'demo': {
        'durable': False, # Delete the queue on broker restart
        'auto_delete': True, # Delete the queue when the client terminates
```

(continues on next page)

(continued from previous page)

```
'exclusive': False, # Allow multiple simultaneous consumers
'arguments': {},
},
}
binding = {
    'exchange': 'amq.topic', # The AMQP exchange to bind our queue to
    'queue': 'demo', # The unique name of our queue on the AMQP broker
    'routing_keys': ['#'], # The topics that should be delivered to the queue
}

# Start consuming messages using our callback. This call will block until
# a KeyboardInterrupt is raised, or the process receives a SIGINT or SIGTERM
# signal.
api.consume(printer_callback, bindings=binding, queues=queues)
```

In this example, there's one queue and the queue only has one binding, but it's possible to consume from multiple queues and each queue can have multiple bindings.

6.2 Command Line Interface

A command line interface, *fedora-messaging*, is included to make running consumers easier. It's not necessary to write any boilerplate code calling the API, just run `fedora-messaging consume` and provide it the Python path to your callback:

```
$ fedora-messaging consume --callback=fedora_messaging.example:printer
```

Consult the manual page for complete details on this command line interface.

Note: For users of `fedmsg`, this is roughly equivalent to `fedmsg-hub`

6.3 Consumer API

The introduction contains a very minimal callback. This section covers the complete API for consumers.

6.3.1 The Callback

The callback provided to `fedora_messaging.api.consume()` or the command-line interface can be any callable Python object, so long as it accepts the message object as a single positional argument.

The API will also accept a Python class, which it will instantiate before using as a callable object. This allows you to write a callback with easy one-time initialization or a callback that maintains state between calls:

```
import os

from fedora_messaging import config
```

(continues on next page)

(continued from previous page)

```

class SaveMessage:
    """
    A fedora-messaging consumer that saves the message to a file.

    A single configuration key is used from fedora-messaging's
    "consumer_config" key, "path", which is where the consumer will save
    the messages::

        [consumer_config]
        path = "/tmp/fedora-messaging/messages.txt"
    """

    def __init__(self):
        """Perform some one-time initialization for the consumer."""
        self.path = config.conf["consumer_config"]["path"]

        # Ensure the path exists before the consumer starts
        if not os.path.exists(os.path.dirname(self.path)):
            os.mkdir(os.path.dirname(self.path))

    def __call__(self, message):
        """
        Invoked when a message is received by the consumer.

        Args:
            message (fedora_messaging.api.Message): The message from AMQP.
        """
        with open(self.path, "a") as fd:
            fd.write(str(message))

```

When running this type of callback from the command-line interface, specify the Python path to the class object, not the `__call__` method:

```
$ fedora-messaging consume --callback=package_name.module:SaveMessage
```

6.3.2 Exceptions

- Consumers should raise the `fedora_messaging.exceptions.Nack` exception if the consumer cannot handle the message at this time. The message will be re-queued, and the server will attempt to re-deliver it at a later time.
- Consumers should raise the `fedora_messaging.exceptions.Drop` exception when they wish to explicitly indicate they do not want handle the message. This is similar to simply calling `return`, but the server is informed the client dropped the message. What the server does depends on configuration.
- Consumers should raise the `fedora_messaging.exceptions.HaltConsumer` exception if they wish to stop consuming messages.

If a consumer raises any other exception, a traceback will be logged at the error level, the message being processed and any pre-fetched messages will be returned to the queue for later delivery, and the consumer will be canceled.

If the CLI is being used, it will halt with a non-zero exit code. If the API is being used directly, consult the API documentation for exact results, as the synchronous and asynchronous APIs communicate failures differently.

6.3.3 Synchronous and Asynchronous Calls

The AMQP consumer runs in a Twisted event loop. When a message arrives, it calls the callback in a separate Python thread to avoid blocking vital operations like the connection heartbeat. The callback is free to use any blocking (synchronous) calls it likes.

Note: Your callback does not need to be thread-safe. By default, messages are processed serially.

It is safe to start threads to perform IO-blocking work concurrently. If you wish to make use of a Twisted API, you must use the `twisted.internet.threads.blockingCallFromThread()` or `twisted.internet.interfaces.IReactorFromThreads` APIs.

You may also use asyncio-based asynchronous callbacks, either via an `async` function or via an object that has an `async __call__()` method. In this case, the callback will not be run in a separate thread, it will instead be scheduled as a regular asyncio task.

6.3.4 Consumer Configuration

A special section of the fedora-messaging configuration will be available for consumers to use if they need configuration options. Refer to the *consumer_config* in the Configuration documentation for details.

6.4 systemd Service

A systemd service file is also included in the Python package for your convenience. It is called `fm-consumer@.service` and simply runs `fedora-messaging consume` with a configuration file from `/etc/fedora-messaging/` that matches the service name:

```
$ systemctl start fm-consumer@sample.service # uses /etc/fedora-messaging/sample.toml
```

AVAILABLE SCHEMAS

These are the topics that you can expect to see on Fedora’s message bus, sorted by the python package that contains their schema. Install the corresponding python package if you want to make use of the schema and access additional information on the message you’re receiving.

In the Fedora Infrastructure, some of those topics will be prefixed by `org.fedoraproject.stg.` in staging and `org.fedoraproject.prod.` in production.

7.1 anitya

You can view the history of [all anitya messages](#) in datagrepper.

- `org.release-monitoring.prod.anitya.distro.add`: Message sent by Anitya to the “anitya.distro.add” topic when a new distribution is added. ([history](#))
- `org.release-monitoring.prod.anitya.distro.edit`: Message sent by Anitya when a distribution is edited. ([history](#))
- `org.release-monitoring.prod.anitya.distro.remove`: Message sent by Anitya when a distribution is removed. ([history](#))
- `org.release-monitoring.prod.anitya.project.add`: The message sent when a new project is created in Anitya. ([history](#))
- `org.release-monitoring.prod.anitya.project.edit`: The message sent when a project is edited in Anitya. ([history](#))
- `org.release-monitoring.prod.anitya.project.flag`: Sent when a new flag is created for a project. ([history](#))
- `org.release-monitoring.prod.anitya.project.flag.set`: Sent when a flag is closed for a project. ([history](#))
- `org.release-monitoring.prod.anitya.project.map.new`: Sent when new distribution mapping is created in Anitya. ([history](#))
- `org.release-monitoring.prod.anitya.project.map.remove`: Sent when distribution mapping is deleted in Anitya. ([history](#))
- `org.release-monitoring.prod.anitya.project.map.update`: Sent when distribution mapping is edited in Anitya. ([history](#))
- `org.release-monitoring.prod.anitya.project.remove`: The message sent when a project is deleted in Anitya. ([history](#))
- `org.release-monitoring.prod.anitya.project.version.remove`: Sent when version is deleted in Anitya. ([history](#))

- `org.release-monitoring.prod.anitya.project.version.remove.v2`: Sent when version is deleted in Anitya. ([history](#))
- `org.release-monitoring.prod.anitya.project.version.update`: Sent when new version is discovered by Anitya. This message will be deprecated in future. ([history](#))
- `org.release-monitoring.prod.anitya.project.version.update.v2`: Sent when new versions are discovered by Anitya. ([history](#))

7.2 bodhi

You can view the history of [all bodhi messages](#) in datagrepper.

- `bodhi.buildroot.override.tag`: Sent when a buildroot override is added and tagged into the build root. ([history](#))
- `bodhi.buildroot.override.untag`: Sent when a buildroot override is untagged from the build root. ([history](#))
- `bodhi.compose.complete`: Sent when a compose task completes. ([history](#))
- `bodhi.compose.composing`: Sent when the compose task composes. ([history](#))
- `bodhi.compose.start`: Sent when a compose task starts. ([history](#))
- `bodhi.compose.sync.done`: Sent when a compose task sync is done. ([history](#))
- `bodhi.compose.sync.wait`: Sent when a compose task sync is waiting. ([history](#))
- `bodhi.errata.publish`: Sent when an errata is published. ([history](#))
- `bodhi.repo.done`: Sent when a repo is created and ready to be signed or otherwise processed. ([history](#))
- `bodhi.update.comment`: Sent when a comment is made on an update. ([history](#))
- `bodhi.update.complete.stable`: Sent when an update is available in the stable repository. ([history](#))
- `bodhi.update.complete.testing`: Sent when an update is available in the testing repository. ([history](#))
- `bodhi.update.edit`: Sent when an update is edited. ([history](#))
- `bodhi.update.edit`: Sent when an update is edited. Newer version. Has ‘new_builds’ and ‘removed_builds’ properties. ([history](#))
- `bodhi.update.eject`: Sent when an update is ejected from the push. ([history](#))
- `bodhi.update.karma.threshold.reach`: Sent when an update reaches its karma threshold. ([history](#))
- `bodhi.update.request.obsolete`: Sent when an update is requested to be obsoleted. ([history](#))
- `bodhi.update.request.revoke`: Sent when an update is revoked. ([history](#))
- `bodhi.update.request.stable`: Sent when an update is submitted as a stable candidate. ([history](#))
- `bodhi.update.request.testing`: Sent when an update is submitted as a testing candidate. ([history](#))
- `bodhi.update.request.unpush`: Sent when an update is requested to be unpushed. ([history](#))
- `bodhi.update.requirements_met.stable`: Sent when all the update requirements are met for stable. ([history](#))
- `bodhi.update.status.testing.koji-build-group.build.complete`: Sent when an update is ready to be tested. Original version. Does not have ‘update’ property or inherit from UpdateMessage. ([history](#))

- `bodhi.update.status.testing.koji-build-group.build.complete`: Sent when an update is ready to be tested. Newer version. Has ‘update’ property, like other update messages. ([history](#))
- `bodhi.update.status.testing.koji-build-group.build.complete`: Sent when an update is ready to be tested. Simplified version. Specifically, this message is sent: * When an update is created * When an update is edited and its builds change * When a “re-trigger tests” request is made via the web UI or API These are the points where we expect that automated systems will test the update. Inherits from `UpdateMessage` and only contains as much extra information (in the ‘artifact’ dict) as the Fedora CI schedulers actually need. ([history](#))

7.3 Copr

You can view the history of [all Copr messages](#) in datagrepper.

- `copr.build.end`: schema for the old fedmsg-era ‘copr.build.end’ message ([history](#))
- `copr.build.start`: schema for the old fedmsg-era ‘copr.build.start’ message ([history](#))
- `copr.chroot.start`: Schema for the old fedmsg-era ‘copr.chroot.start’ message, this message duplicated the ‘copr.build.start’ message, so you should never use this. ([history](#))
- `build.end` ([history](#))
- `build.start` ([history](#))
- `chroot.start` ([history](#))

7.4 fedocal

You can view the history of [all fedocal messages](#) in datagrepper.

- `fedocal.calendar.clear`: A sub-class of a Fedora message that defines a message schema for messages published by fedocal when a calendar is cleared. ([history](#))
- `fedocal.calendar.delete`: A sub-class of a Fedora message that defines a message schema for messages published by fedocal when a calendar is deleted. ([history](#))
- `fedocal.calendar.new`: A sub-class of a Fedora message that defines a message schema for messages published by fedocal when a calendar is created. ([history](#))
- `fedocal.calendar.update`: A sub-class of a Fedora message that defines a message schema for messages published by fedocal when a calendar is updated. ([history](#))
- `fedocal.calendar.upload`: A sub-class of a Fedora message that defines a message schema for messages published by fedocal when meetings have been uploaded into the calendar. ([history](#))
- `fedocal.meeting.delete`: A sub-class of a Fedora message that defines a message schema for messages published by fedocal when a meeting is deleted. ([history](#))
- `fedocal.meeting.new`: A sub-class of a Fedora message that defines a message schema for messages published by fedocal when a meeting is created. ([history](#))
- `fedocal.meeting.reminder`: A sub-class of a Fedora message that defines a message schema for messages published by fedocal when a reminder is sent. ([history](#))
- `fedocal.meeting.update`: A sub-class of a Fedora message that defines a message schema for messages published by fedocal when a meeting is updated. ([history](#))

7.5 elections

You can view the history of [all elections messages](#) in datagrepper.

- `fedora_elections.candidate.delete`: A sub-class of a Fedora message that defines a message schema for messages published by Elections when a candidate is deleted. ([history](#))
- `fedora_elections.candidate.edit`: A sub-class of a Fedora message that defines a message schema for messages published by Elections when a candidate is edited. ([history](#))
- `fedora_elections.candidate.new`: A sub-class of a Fedora message that defines a message schema for messages published by Elections when a new candidate is added. ([history](#))
- `fedora_elections.election.edit`: A sub-class of a Fedora message that defines a message schema for messages published by Elections when an election is edited. ([history](#))
- `fedora_elections.election.new`: A sub-class of a Fedora message that defines a message schema for messages published by Elections when a new election is created. ([history](#))

7.6 Git

You can view the history of [all Git messages](#) in datagrepper.

- `git.receive`: A sub-class of a Fedora message that defines a message schema for messages published by Fedora Messaging Git Hook when a new commit is received. ([history](#))

7.7 The New Hotness

You can view the history of [all The New Hotness messages](#) in datagrepper.

- `org.fedoraproject.prod.hotness.update.bug.file`: Message sent by the-new-hotness to “hotness.update.bug.file” topic when bugzilla issue is filled. ([history](#))
- `org.fedoraproject.prod.hotness.update.drop`: Message sent by the-new-hotness to “hotness.update.drop” topic when update is dropped. ([history](#))

7.8 planet

You can view the history of [all planet messages](#) in datagrepper.

- `org.fedoraproject.prod.planet.post.new`: The message sent when a new post is published in planet. ([history](#))

7.9 ansible

You can view the history of [all ansible messages](#) in datagrepper.

- `ansible.playbook.complete`: Defines the message that is sent when an Ansible Playbook completes ([history](#))
- `ansible.playbook.start`: Defines the message that is sent when an Ansible Playbook starts ([history](#))
- `git.receive`: Defines the message that is sent when an Ansible Playbook starts ([history](#))

7.10 FMN

You can view the history of [all FMN messages](#) in datagrepper.

- `fmn.rule.update.v1` ([history](#))
- `fmn.rule.delete.v1` ([history](#))
- `fmn.rule.update.v1` ([history](#))

7.11 kerneltest

You can view the history of [all kerneltest messages](#) in datagrepper.

- `kerneltest.release.edit`: The message sent when an admin creates a new release ([history](#))
- `kerneltest.release.new`: The message sent when an admin creates a new release ([history](#))
- `kerneltest.upload.new`: The message sent when a user uploads a new kerneltest ([history](#))

7.12 Koji

You can view the history of [all Koji messages](#) in datagrepper.

- `buildsys.build.state.change`: This message is sent when a build state changes. ([history](#))
- `buildsys.package.list.change`: This message is sent when a package list changes. ([history](#))
- `buildsys.repo.done`: This message is sent when a package repo is done. ([history](#))
- `buildsys.repo.init`: This message is sent when a package repo is initialized. ([history](#))
- `buildsys.rpm.sign`: This message is sent when a rpm is signed. ([history](#))
- `buildsys.tag`: This message is sent when a package is tagged. ([history](#))
- `buildsys.untag`: This message is sent when a package is untagged. ([history](#))
- `buildsys.task.state.change`: This message is sent when a task state changes. ([history](#))

7.13 Koschei

You can view the history of [all Koschei messages](#) in datagrepper.

- `koschei.collection.state.change`: Messages published by Koschei when a collection state changes. For example when collection buildroot becomes unresolvable (broken) or when it is fixed. ([history](#))
- `koschei.package.state.change`: Messages published by Koschei when a package state changes. For example when package starts to fail to build, package dependencies become unresolved or when package is fixed. ([history](#))

7.14 mdapi

You can view the history of [all mdapi messages](#) in datagrepper.

- `mdapi.repo.update`: A sub-class of a Fedora message that defines a message schema for messages published by mdapi when a repo's info is updated. ([history](#))

7.15 Wiki

You can view the history of [all Wiki messages](#) in datagrepper.

- `wiki.article.edit`: A sub-class of a Fedora message that defines a message schema for messages published by Mediawiki when a new thing is created. ([history](#))

7.16 meetbot

You can view the history of [all meetbot messages](#) in datagrepper.

- `meetbot.meeting.complete` ([history](#))
- `meetbot.meeting.start` ([history](#))

7.17 FAS

You can view the history of [all FAS messages](#) in datagrepper.

- `fas.group.member.sponsor`: The message sent when a user is added to a group by a sponsor ([history](#))
- `fas.user.create`: The message sent when a user is created ([history](#))
- `fas.user.update`: The message sent when a user is updated ([history](#))

7.18 nuancier

You can view the history of [all nuancier messages](#) in datagrepper.

- `nuancier.new`: A sub-class of a Fedora message that defines a message schema for messages published by nuancier when a new thing is created. ([history](#))

7.19 Pagure

You can view the history of [all Pagure messages](#) in datagrepper.

- `pagure.Test.notification`: A sub-class of a Fedora message that defines a message schema for messages published by pagure when a new thing is created. ([history](#))
- `pagure.commit.flag.added`: A sub-class of a Fedora message that defines a message schema for messages published by pagure when a new thing is created. ([history](#))
- `pagure.commit.flag.updated`: A sub-class of a Fedora message that defines a message schema for messages published by pagure when a new thing is created. ([history](#))
- `pagure.git.branch.creation`: A sub-class of a Fedora message that defines a message schema for messages published by pagure when a new thing is created. ([history](#))
- `pagure.git.branch.deletion`: A sub-class of a Fedora message that defines a message schema for messages published by pagure when a new thing is created. ([history](#))
- `pagure.git.receive`: A sub-class of a Fedora message that defines a message schema for messages published by pagure when a new thing is created. ([history](#))
- `pagure.git.tag.creation`: A sub-class of a Fedora message that defines a message schema for messages published by pagure when a new thing is created. ([history](#))
- `pagure.git.tag.deletion`: A sub-class of a Fedora message that defines a message schema for messages published by pagure when a new thing is created. ([history](#))
- `pagure.group.edit`: A sub-class of a Fedora message that defines a message schema for messages published by pagure when a new thing is created. ([history](#))
- `pagure.issue.assigned.added`: A sub-class of a Fedora message that defines a message schema for messages published by pagure when an issue is assigned. ([history](#))
- `pagure.issue.assigned.reset`: A sub-class of a Fedora message that defines a message schema for messages published by pagure when an issue is un-assigned. ([history](#))
- `pagure.issue.comment.added`: A sub-class of a Fedora message that defines a message schema for messages published by pagure when a comment is added to an issue. ([history](#))
- `pagure.issue.dependency.added`: A sub-class of a Fedora message that defines a message schema for messages published by pagure when a dependency is added to an issue. ([history](#))
- `pagure.issue.dependency.removed`: A sub-class of a Fedora message that defines a message schema for messages published by pagure when an issue is deleted. ([history](#))
- `pagure.issue.drop`: A sub-class of a Fedora message that defines a message schema for messages published by pagure when an issue is deleted. ([history](#))
- `pagure.issue.edit`: A sub-class of a Fedora message that defines a message schema for messages published by pagure when an issue is updated. ([history](#))
- `pagure.issue.new`: A sub-class of a Fedora message that defines a message schema for messages published by pagure when a new thing is created. ([history](#))

- `pagure.issue.tag.added`: A sub-class of a Fedora message that defines a message schema for messages published by pagure when an issue is deleted. ([history](#))
- `pagure.issue.tag.removed`: A sub-class of a Fedora message that defines a message schema for messages published by pagure when an issue is deleted. ([history](#))
- `pagure.project.deleted`: A sub-class of a Fedora message that defines a message schema for messages published by pagure when a new thing is created. ([history](#))
- `pagure.project.edit`: A sub-class of a Fedora message that defines a message schema for messages published by pagure when a new thing is created. ([history](#))
- `pagure.project.forked`: A sub-class of a Fedora message that defines a message schema for messages published by pagure when a new thing is created. ([history](#))
- `pagure.project.group.access.updated`: A sub-class of a Fedora message that defines a message schema for messages published by pagure when a new thing is created. ([history](#))
- `pagure.project.group.added`: A sub-class of a Fedora message that defines a message schema for messages published by pagure when a new thing is created. ([history](#))
- `pagure.project.group.removed`: A sub-class of a Fedora message that defines a message schema for messages published by pagure when a new thing is created. ([history](#))
- `pagure.project.new`: A sub-class of a Fedora message that defines a message schema for messages published by pagure when a new thing is created. ([history](#))
- `pagure.project.tag.edited`: A sub-class of a Fedora message that defines a message schema for messages published by pagure when a new thing is created. ([history](#))
- `pagure.project.tag.removed`: A sub-class of a Fedora message that defines a message schema for messages published by pagure when a new thing is created. ([history](#))
- `pagure.project.user.access.updated`: A sub-class of a Fedora message that defines a message schema for messages published by pagure when a new thing is created. ([history](#))
- `pagure.project.user.added`: A sub-class of a Fedora message that defines a message schema for messages published by pagure when a new thing is created. ([history](#))
- `pagure.project.user.removed`: A sub-class of a Fedora message that defines a message schema for messages published by pagure when a new thing is created. ([history](#))
- `pagure.pull-request.assigned.added`: A sub-class of a Fedora message that defines a message schema for messages published by pagure when a pull request is assigned. ([history](#))
- `pagure.pull-request.assigned.reset`: A sub-class of a Fedora message that defines a message schema for messages published by pagure when a pull request is un-assigned. ([history](#))
- `pagure.pull-request.closed`: A sub-class of a Fedora message that defines a message schema for messages published by pagure when a pull request is closed. ([history](#))
- `pagure.pull-request.comment.added`: A sub-class of a Fedora message that defines a message schema for messages published by pagure when a comment is added to a PR. ([history](#))
- `pagure.pull-request.comment.edited`: A sub-class of a Fedora message that defines a message schema for messages published by pagure when a comment is edited on a PR. ([history](#))
- `pagure.pull-request.flag.added`: A sub-class of a Fedora message that defines a message schema for messages published by pagure when a flag is added on a PR. ([history](#))
- `pagure.pull-request.flag.updated`: A sub-class of a Fedora message that defines a message schema for messages published by pagure when a flag is updated on a PR ([history](#))

- `pagure.pull-request.initial_comment.edited`: A sub-class of a Fedora message that defines a message schema for messages published by pagure when an initial PR comment is edited. ([history](#))
- `pagure.pull-request.new`: A sub-class of a Fedora message that defines a message schema for messages published by pagure when a pull request is created. ([history](#))
- `pagure.pull-request.rebased`: A sub-class of a Fedora message that defines a message schema for messages published by pagure when a PR is rebased. ([history](#))
- `pagure.pull-request.reopened`: A sub-class of a Fedora message that defines a message schema for messages published by pagure when a PR is reopened. ([history](#))
- `pagure.pull-request.tag.added`: A sub-class of a Fedora message that defines a message schema for messages published by pagure when a tag is added on a PR. ([history](#))
- `pagure.pull-request.tag.removed`: A sub-class of a Fedora message that defines a message schema for messages published by pagure when a tag is removed on a PR. ([history](#))
- `pagure.pull-request.updated`: A sub-class of a Fedora message that defines a message schema for messages published by pagure when a PR is updated. ([history](#))

7.20 tahrir

You can view the history of [all tahrir messages](#) in datagrepper.

- `person.login.first`: The message sent when a user logs into tahrir for the first time ([history](#))

TESTING

Once you've written code to publish or consume messages, you'll probably want to test it. The `fedora_messaging.testing` module has utilities for common test patterns.

If you find yourself implementing a pattern over and over in your test code, consider contributing it here!

`fedora_messaging.testing.mock_sends(*expected_messages)`

Assert a block of code results in the provided messages being sent without actually sending them.

This is intended for unit tests. The call to publish is mocked out and messages are captured and checked at the end of the `with`.

For example:

```
>>> from fedora_messaging import api, testing
>>> def publishes():
...     api.publish(api.Message(body={"Hello": "world"}))
...
>>> with testing.mock_sends(api.Message, api.Message(body={"Hello": "world"})):
...     publishes()
...     publishes()
...
>>> with testing.mock_sends(api.Message(body={"Goodbye": "everybody"})):
...     publishes()
...
AssertionError
```

Parameters

***expected_messages** – The messages you expect to be sent. These can be classes instances of classes derived from `fedora_messaging.message.Message`. If the class is provided, the message is checked to make sure it is an instance of that class and that it passes schema validation. If an instance is provided, it is checked for equality with the sent message.

Raises

AssertionError – If the messages published don't match the messages asserted.

COMMAND LINE INTERFACE MANUALS

9.1 fedora-messaging

9.1.1 Synopsis

`fedora-messaging COMMAND [OPTIONS] [ARGS]...`

9.1.2 Description

`fedora-messaging` can be used to work with AMQP message brokers using the `fedora-messaging` library to start message consumers.

9.1.3 Options

`--help`

Show help text and exit.

`--conf`

Path to a valid configuration file to use in place of the configuration in `/etc/fedora-messaging/config.toml`.

9.1.4 Commands

There are three sub-commands, `consume`, `publish` and `record`, described in detail in their own sections below.

`fedora-messaging consume [OPTIONS]`

Starts a consumer process with a user-provided callback function to execute when a message arrives.

`fedora-messaging publish [OPTIONS] FILE`

Loads serialized messages from a file and publishes them to the specified exchange.

`fedora-messaging record [OPTIONS] FILE`

Records messages arrived from AMQP queue and saves them to file with specified name.

consume

All options below correspond to settings in the configuration file. However, not all available configuration keys can be overridden with options, so it is recommended that for complex setups and production environments you use the configuration file and no options on the command line.

--app-name

The name of the application, used by the AMQP client to identify itself to the broker. This is purely for administrator convenience to determine what applications are connected and own particular resources.

This option is equivalent to the `app` setting in the `client_properties` section of the configuration file.

--callback

The Python path to the callable object to execute when a message arrives. The Python path should be in the format `module.path:object_in_module` and should point to either a function or a class. Consult the API documentation for the interface required for these objects.

This option is equivalent to the `callback` setting in the configuration file.

--routing-key

The AMQP routing key to use with the queue. This controls what messages are delivered to the consumer. Can be specified multiple times; any message that matches at least one will be placed in the message queue.

Setting this option is equivalent to setting the `routing_keys` setting in *all* `bindings` entries in the configuration file.

--queue-name

The name of the message queue in AMQP. Can contain ASCII letters, digits, hyphen, underscore, period, or colon. If one is not specified, a unique name will be created for you.

Setting this option is equivalent to setting the `queue` setting in *all* `bindings` entries and creating a `queue.<queue-name>` section in the configuration file.

--exchange

The name of the exchange to bind the queue to. Can contain ASCII letters, digits, hyphen, underscore, period, or colon.

Setting this option is equivalent to setting the `exchange` setting in *all* `bindings` entries in the configuration file.

publish

The `publish` command expects the message or messages provided in `FILE` to be JSON objects with each message separated by a newline character. The JSON object is described by the following JSON schema:

```
{
  "$schema": "http://json-schema.org/draft-04/schema#",
  "description": "Schema for the JSON object used to represent messages in a file",
  "type": "object",
  "properties": {
    "topic": {"type": "string", "description": "The message topic"},
    "headers": {
      "type": "object",
      "properties": Message.headers_schema["properties"],
      "description": "The message headers",

```

(continues on next page)

(continued from previous page)

```

    },
    "id": {"type": "string", "description": "The message's UUID."},
    "body": {"type": "object", "description": "The message body."},
    "queue": {
        "type": "string",
        "description": "The queue the message arrived on, if any.",
    },
},
"required": ["topic", "headers", "id", "body", "queue"],
}

```

They can be produced from Message objects by the `fedora_messaging.api.dumps()` API. `stdin` can be used instead of a file by providing `-` as an argument:

```
$ fedora-messaging publish -
```

Options

`--exchange`

The name of the exchange to publish to. Can contain ASCII letters, digits, hyphen, underscore, period, or colon.

`record`

`--limit`

The maximum number of messages to record.

`--app-name`

The name of the application, used by the AMQP client to identify itself to the broker. This is purely for administrator convenience to determine what applications are connected and own particular resources. If not specified, the default is `recorder`.

This option is equivalent to the `app` setting in the `client_properties` section of the configuration file.

`--routing-key`

The AMQP routing key to use with the queue. This controls what messages are delivered to the consumer. Can be specified multiple times; any message that matches at least one will be placed in the message queue.

Setting this option is equivalent to setting the `routing_keys` setting in *all* `bindings` entries in the configuration file.

`--queue-name`

The name of the message queue in AMQP. Can contain ASCII letters, digits, hyphen, underscore, period, or colon. If one is not specified, a unique name will be created for you.

Setting this option is equivalent to setting the `queue` setting in *all* `bindings` entries and creating a queue. `<queue-name>` section in the configuration file.

`--exchange`

The name of the exchange to bind the queue to. Can contain ASCII letters, digits, hyphen, underscore, period, or colon.

Setting this option is equivalent to setting the `exchange` setting in *all* `bindings` entries in the configuration file.

9.1.5 Exit codes

consume

The `consume` command can exit for a number of reasons:

0

The consumer intentionally halted by raising a `HaltConsumer` exception.

2

The argument or option provided is invalid.

10

The consumer was unable to declare an exchange, queue, or binding in the message broker. This occurs with the user does not have permission on the broker to create the object *or* the object already exists, but does not have the attributes the consumer expects (e.g. the consumer expects it to be a durable queue, but it is transient).

11

The consumer encounters an unexpected error while registering the consumer with the broker. This is a bug in `fedora-messaging` and should be reported.

12

The consumer is canceled by the message broker. The consumer is typically canceled when the queue it is subscribed to is deleted on the broker, but other exceptional cases could result in this. The broker administrators should be consulted in this case.

13

An unexpected general exception is raised by your consumer callback.

Additionally, consumer callbacks can cause the command to exit with a custom exit code. Consult the consumer's documentation to see what error codes it uses.

publish

0

The messages were all successfully published.

1

A general, unexpected exception occurred and the message was not successfully published.

121

The message broker rejected the message, likely due to resource constraints.

111

A connection to the broker could not be established.

9.1.6 Signals

consume

The `consume` command handles the `SIGTERM` and `SIGINT` signals by allowing any consumers which are currently processing a message to finish, acknowledging the message to the message broker, and then shutting down. Repeated `SIGTERM` or `SIGINT` signals are ignored. To halt immediately, send the `SIGKILL` signal; messages that are partially processed will be re-delivered when the consumer restarts.

9.1.7 Systemd service

The `consume` subcommand can be started as a system service, and Fedora Messaging provides a dynamic systemd service file.

First, create a valid Fedora Messaging configuration file in `/etc/fedora-messaging/foo.toml`, with the `callback` parameter pointing to your consuming function or class. Remember that you can use the `consumer_config` section for your own configuration.

Enable and start the service in systemd with the following commands:

```
systemctl enable fm-consumer@foo.service
systemctl start fm-consumer@foo.service
```

The service name after the `@` and before the `.service` must match your filename in `/etc/fedora-messaging` (without the `.toml` suffix).

9.1.8 Help

If you find bugs in `fedora-messaging` or its man page, please file a bug report or a pull request:

```
https://github.com/fedora-infra/fedora-messaging
```

Or, if you prefer, send an email to infrastructure@fedoraproject.org with bug reports or patches.

`fedora-messaging`'s documentation is available online:

```
https://fedora-messaging.readthedocs.io/
```


INSTALLATION

10.1 Installing the library

Create a Python virtual environment:

```
mkdir fedora-messaging-tutorial
cd fedora-messaging-tutorial
mkvirtualenv -p python3 -a `pwd` fedora-messaging-tutorial
workon fedora-messaging-tutorial
```

Install the library and its dependencies:

```
pip install fedora-messaging
# Alternatively, install it directly from the git repository
git clone https://github.com/fedora-infra/fedora-messaging.git
cd fedora-messaging
pip install -e .
```

Make sure it is available and working:

```
fedora-messaging --help
```

10.2 Setting up RabbitMQ

Install RabbitMQ and start it:

```
dnf install rabbitmq-server
systemctl start rabbitmq-server
```

Enable RabbitMQ web admin interface:

```
rabbitmq-plugins enable rabbitmq_management
```

RabbitMQ has a web admin interface that you can access at: <http://localhost:15672/>. The username is `guest` and the password is `guest`. This interface lets you change the configuration, send messages and read the messages in the queues. Keep it open in a browser tab, we'll need it later.

If your project uses containers, consult the [RabbitMQ documentation](#) about containers.

10.3 Configuration

An example of the library configuration file is provided in the `config.toml.example` file. Copy that file to `/etc/fedora-messaging/config.toml` to make it available system-wide. Alternatively, you can copy it to `config.toml` anywhere and set the `FEDORA_MESSAGING_CONF` environment variable to that file's path.

Refer to [the documentation](#) for a complete description of the configuration options.

Comment out the callback and bindings options, and all the `[exchanges.custom_exchange]` and `[queues.my_queue]` sections.

In the `[client_properties]` section, change the `app` value to `Fedora Messaging tutorial`.

USING THE API

We will be creating some scripts to publish and subscribe to the bus. First, create a directory to hold the code you will write, then change to this directory.

11.1 Publishing

To publish on the Fedora Messaging bus, you just need to use the `fedora_messaging.api.publish()` function, passing it an instance of the `fedora_messaging.message.Message` class that represents the message you want to publish.

A message has a schema, a topic, a severity, a body, and a set of headers. We'll cover the schema later in this tutorial. The headers and the body are Python dictionaries with JSON-serializable values. The topic is a string containing elements separated by dots that will be used to route messages.

Create a publishing script called `publish.py`:

```
#!/usr/bin/env python3

from fedora_messaging.api import publish, Message
from fedora_messaging.config import conf

conf.setup_logging()
message = Message(
    topic="tutorial.topic",
    body={"reason": "test message"}
)
publish(message)
```

Of course, you can make a smarter script that will use command-line arguments, this is left as an exercise to the reader. Now run it:

```
chmod +x publish.py
./publish.py
```

The script should complete without error. If you go to RabbitMQ's web interface, you'll see that a message has been sent to the `amq.topic` exchange. However, since no one is listening to this topic, the message has been discarded. Now, we'll setup listeners.

11.2 Listening

Clients listen on the Fedora Messaging bus by subscribing to a topic or a topic pattern using the hash (#) symbol as a wildcard. For example you can subscribe to `bodhi.updates.kernel` but also to `bodhi.updates.#`. In the former case you'll get kernel updates, in the latter case you'll get all Bodhi updates.

After subscription, all messages with a topic matching the pattern will be routed to a queue on the server, and clients will consume messages from this queue. In the AMQP language, this is called *binding* a queue to an exchange, and the topic pattern is called the *routing_key*.

In the configuration file, the `bindings` section controls which queues will be subscribed to which topic patterns. Edit the file so the option looks like this:

```
[[bindings]]
queue = "tutorial"
exchange = "amq.topic"
routing_keys = ["tutorial.#"]
```

This means that the queue named `tutorial` will be created and subscribed to the `amq.topic` exchange using the `tutorial.#` pattern. All messages with a topic starting with `tutorial.` will end up in this queue, and no other.

Now configure this new queue's properties in the file using a snippet that looks like this:

```
[[queues.tutorial]]
durable = true
auto_delete = false
exclusive = false
arguments = {}
```

This means that messages in this queue will survive a client's disconnection and a server restart, and that multiple client can connect to it simultaneously to consume messages in a round-robin fashion.

11.2.1 Python script

Now create the following script, called `consume.py`:

```
#!/usr/bin/env python3

from fedora_messaging.api import consume
from fedora_messaging.config import conf

conf.setup_logging()

def print_message(message):
    print(message)

if __name__ == "__main__":
    conf.setup_logging()
    consume(print_message)
```

The script should run and wait for new messages. Now run the `publish.py` script again in another terminal (remember to activate the virtualenv with `workon fedora-messaging-tutorial`). You should see the message being printed where the `consume.py` script is running.

11.2.2 Python callback

You can also just define the callback function and use the `fedora-messaging` command-line tool to do the listening:

```
fedora-messaging consume --callback="consume:print_message"
```

This should behave identically.

11.2.3 Round robin

When multiple programs are simultaneously consuming from the same queue, they get the messages in a round-robin fashion. Try running another instance of the `consume.py` script, and run the `publish.py` script multiple times. You'll see that `consume.py` instances get a message one after the other.

JSON SCHEMAS

Message bodies are JSON objects, that adhere to a schema. Message schemas live in their own Python package, so they can be installed on the producer and on the consumer.

In Fedora Messaging, we follow the [JSON Schema](#) standard, and use the [jsonschema](#) library.

12.1 Creating the schema package

Copy the `docs/sample_schema_package/` directory from the `fedora-messaging` git clone to your app directory.

Edit the `setup.py` file to change the package metadata. Rename the `mailman_messages` directory to something relevant to your app, like `yourapp_messages`. This is the naming convention. Edit the `README` file too.

If you prefer [CookieCutter](#), there is a [template repository](#) that you can use with the command:

```
cookiecutter gh:fedora-infra/cookiecutter-message-schemas
```

12.2 Writing the schema

JSON objects are converted to dictionaries in Python. Writing a JSON schema with the [jsonschema](#) library means writing a Python dictionary that will describe the message's JSON object body. Read up on the [jsonschema](#) library documentation if you have questions about the format.

Open the `messages.py` file, it contains an example schema for Mailman-originating messages on the bus. The schema is a Python class containing an important dictionary attribute: `body_schema`. This is where the JSON schema lives.

For clarity, edit the `setup.py` file and in the entry points list change the `mailman.messageV1` name to something more relevant to your app, like `yourapp.my_messageV1`. The entry point name needs to be unique to your application, so it's best to prefix it with your package or application name.

12.2.1 Schema format

This dictionary describes the possible keys and types in the JSON object being validated, using the following reserved keys:

- `id` (or `$id`): an URI identifying this schema. Change the last part of the example URL to use your app's name.
- `$schema`: an URI describing the validator to use, you can leave that one as it is. It is only present at the root of the dictionary.
- `description`: a fulltext description of the key.

- **type:** the value type for this key. You can choose among: - **null:** equivalent to `None` - **boolean:** equivalent to `True` or `False` - **object:** a Python dictionary - **array:** a Python list - **number:** an `int` or a `float` - **string:** a Python string
- **properties:** a dictionary describing the possible keys contained in the JSON object, where keys are possible key names, and values are JSON schemas. Those schemas can also have **properties** keys to describe all the possible nested keys.
- **required:** a list of keys that must be present in the JSON object.
- **format:** a format validation type. You can choose among: - **hostname** - **ipv4** - **ipv6** - **email** - **uri** (requires the `rfc3987` package) - **date** - **time** - **date-time** (requires the `strict-rfc3339` package) - **regex** - **color** (requires the `webcolors` package)

For information on creating JSON schemas to validate your data, there is a good introduction to JSON Schema fundamentals underway at [Understanding JSON Schema](#).

12.2.2 Example

Now edit the `body_schema` key to use the following schema:

```
{
  'id': 'http://fedoraproject.org/message-schema/fedora-messaging-tutorial#',
  '$schema': 'http://json-schema.org/draft-04/schema#',
  'description': 'Schema for the Fedora Messaging tutorial',
  'type': 'object',
  'properties': {
    'package': {
      'type': 'object',
      'properties': {
        'name': {
          'type': 'string',
          'description': 'The name of the package',
        },
      },
      'version': {'type': 'string'},
    },
    'required': ['name'],
  },
  'owner': {
    'description': 'The owner of the package',
    'type': 'string',
  },
  'required': ['package', 'owner'],
}
```


12.2.3 Human readable representation

The schema class also contains a few methods to extract relevant information from the message, or to create a human-readable representation.

Change the `__str__()` method to use the expected items from the message body. For example:

```
return '{owner} did something to the {package} package'.format(
    owner=self.body['owner'], package=self.body['package']['name'])
```

Also edit the `summary` property to return something relevant.

12.2.4 Severity

Messages can also have a severity level. This is used by consumers to determine the importance of a message to an end user. The possibly severity levels are defined in the [Message Severity](#) API documentation.

You should set a reasonable default for your messages.

12.3 Testing it

JSON schemas can also be unit-tested. Check out the `tests/test_messages.py` file and write the unit tests that are appropriate for the message schema and the methods you just wrote. Use the example tests for inspiration.

12.4 Using it

To use your new JSON schema, its Python distribution must be available on the system. Run `python setup.py develop` in the schema directory to install it.

Now you can use the `yourapp_messages.messages.Message` class (or however you named the package) to construct your message instances and call `fedora_messaging.api.publish` on them. Edit the `publish.py` script to read:

```
#!/usr/bin/env python3

from fedora_messaging.api import publish
from fedora_messaging.config import conf
from yourapp_messages.messages import Message

conf.setup_logging()
message = Message(
    topic="tutorial.topic",
    body={
        "owner": "fedorauser",
        "package": {
            "name": "foobar",
            "version": "1.0",
        }
    }
)
publish(message)
```

Start a consumer, and send the message. Try to comment out the “owner” key and see what happens when you try to send a message that is not valid according to the schema.

12.5 Updating it

Message formats can change over time, and the schema must change to reflect that. When that happens, you need to copy the old class to a new class in the schemas package, make the changes you need to do, and import the new one in your publisher. You must also add a new entry in the `entry_points` argument in the schema package’s `setup.py` file. The name of the entry point is currently unused, only the class path matters.

However, be warned that messages published with the new class may be dropped by the receivers if they don’t have the new schema available locally. Therefore, you should publish the schema package with the new schema, update it on all the receivers, restart them, and then start using the new version in the publishers.

You should keep the old schema versions in the schemas package for a reasonable amount of time, long enough to make sure all receivers are up-to-date. To avoid clutter, we recommend you use a separate module per schema version (`yourapp_messages.v1:Message`, `yourapp_messages.v2:Message`, etc)

Now create a new version and use it in the `publish.py` script. Send a message before restarting the `consume.py` script to see what happens when a message with an unknown schema is received. Now restart the `consume.py` script and re-send the message.

HANDLING EXCEPTIONS

All exceptions are located in the *fedora_messaging.exceptions* module.

13.1 When publishing

When calling *fedora_messaging.api.publish()*, the following exceptions can be raised:

- **ValidationError**: raised if the message fails validation with its JSON schema. This only depends on the message you are trying to send, the AMQP server is not involved.
- **PublishReturned**: raised if the broker rejects the message.
- **PublishForbidden**: raised if the broker rejects the message because of permissions issues.
- **ConnectionException**: raised if a connection error occurred before the publish confirmation arrived.

The **ValidationError** exception means you should fix either the schema (and maybe make a new version) or the message. No need to catch it, this should crash your app during development and testing.

Your app may handle the other two exceptions in whichever way is relevant. It should involve logging, and sending again or discarding may be valid options.

You already noticed the **ValidationError** being raised when you tried sending an invalid message in the previous chapter.

13.2 When consuming

Invalid messages according to the JSON schema are automatically rejected by the client.

The callback function can raise the following exceptions:

- **Nack**: raise this to return the message to the queue
- **Drop**: raise this to drop the message
- **HaltConsumer**: raise this to shutdown the consumer and return the message to the queue.

Any other exception will bubble up in the consumer, shut it down, and return pending messages to the queue. Your app will have to handle the exception.

Modify the callback function to raise those exceptions and see what happens.

When returning **Nack** systematically, the consumer will just loop on that one message, as it is put back in the queue and delivered again forever.

Notice how raising `HaltConsumer` or another exception stops the consumer, but does not consume the message: it will be re-delivered on the next startup.

CONVERTING A FEDMSG APPLICATION

14.1 Converting publishers

14.1.1 Converting a Flask app

Let's use the `elections` app as an example. Clone the code using the following command:

```
git clone https://pagure.io/elections.git
```

And change to this directory.

In the `elections` app, all calls to publish messages on `fedmsg` are going through the `fedora_elections.fedmsgshim.publish` wrapper function. We can thus modify this function to make it call Fedora Messaging instead of `fedmsg`.

JSON schema

First, you will need a Message schema. To write this schema you must know what kind of messages are sent on the bus. A `git grep` command will reveal that all calls are made from the `admin.py` file. Open that file and examine those calls.

In parallel, copy the `docs/sample_schema_package/` directory from the `fedora-messaging` git clone to your app directory. Rename it to `elections-messages`. Edit the `setup.py` file like you did before, to change the package metadata (including the entry point). Use `fedora_elections_messages` for the name. Rename the `mailman_messages` directory to `fedora_elections_messages` and adapt the `setup.py` metadata.

Edit the `messages.py` file and write the basic structure for the elections message schema. According to the different calls in `admin.py`, it could be something like:

```
{
    'id': 'http://fedoraproject.org/message-schema/elections#',
    '$schema': 'http://json-schema.org/draft-04/schema#',
    'description': 'Schema for Fedora Elections',
    'type': 'object',
    'properties': {
        'agent': {'type': 'string'},
        'election': {'type': 'object'},
        'candidate': {'type': 'object'},
    },
    'required': ['agent', 'election'],
}
```

This could be sufficient, but it would be best to list what properties are available in the `election` and `candidate` keys. Unfortunately, those are just JSON dumps of the database model, so you'll have to look further to know the structure.

Examining the `to_json()` methods in `models.py` shows which keys are dumped to JSON. The schema could be written as:

```
{
  'id': 'http://fedoraproject.org/message-schema/elections#',
  '$schema': 'http://json-schema.org/draft-04/schema#',
  'description': 'Schema for Fedora Elections',
  'type': 'object',
  'properties': {
    'agent': {'type': 'string'},
    'election': {
      'type': 'object',
      'properties': {
        'shortdesc': {'type': 'string'},
        'alias': {'type': 'string'},
        'description': {'type': 'string'},
        'url': {'type': 'string', 'format': 'uri'},
        'start_date': {'type': 'string'},
        'end_date': {'type': 'string'},
        'embargoed': {'type': 'number'},
        'voting_type': {'type': 'string'},
      },
      'required': [
        'shortdesc', 'alias', 'description', 'url',
        'start_date', 'end_date', 'embargoed', 'voting_type',
      ],
    },
    'candidate': {
      'type': 'object',
      'properties': {
        'name': {'type': 'string'},
        'url': {'type': 'string', 'format': 'uri'},
      },
      'required': ['name', 'url'],
    },
  },
  'required': ['agent', 'election'],
}
```

Use this schema and adapt the `__str__()` method and the `summary` property.

Since the schema is distributed in a separate python package, it must be added to the `election` app's dependencies in `requirements.txt`.

Wrapper function

Now you can import this class in `fedora_elections/fedmsgshim.py` and use it to encapsulate the messages. The wrapper could look like:

```
import logging

from fedora_elections_messages.schema import Message
from fedora_messaging.api import publish as fm_publish
from fedora_messaging.exceptions import PublishReturned, PublishForbidden, \
    ConnectionException

LOGGER = logging.getLogger(__name__)

def publish(topic, msg):
    try:
        fm_publish(Message(
            topic="fedora.elections." + topic,
            body=msg,
        ))
    except (PublishReturned, PublishForbidden) as e:
        LOGGER.warning(
            "Fedora Messaging broker rejected message %s: %s",
            msg.id, e
        )
    except ConnectionException as e:
        LOGGER.warning("Error sending the message %s: %s", msg.id, e)
```

With this you'll get a couple of nice features over the previous state of things:

- the message format is validated, so it's your responsibility to update the schema when you decide to change the format, and not the receiver's responsibility to handle any database schema changes you may make that may bleed into the message dictionary. And you'll know during development if you break compatibility.
- you may handle messaging errors in anyway you deem relevant. Here we're just logging them but you could choose to re-send the messages, store them for further analysis, etc.
- when there are no exceptions, you know that the message has reached the broker and has been distributed.

Testing

Let's start the election app and make sure messages are properly sent on the bus. First, we'll create a virtualenv, and install election and fedora-messaging with the following commands:

```
virtualenv venv
source ./venv/bin/activate
pushd elections-message-schemas
python setup.py develop
popd
pip install -r requirements.txt
python setup.py develop
```

Make sure the Fedora Messaging configuration file is correct in `/etc/fedora-messaging/config.toml`. We will add a queue binding to route messages with the `fedora.elections` topic to the tutorial queue. Add this entry in the bindings list:

```
[[bindings]]
queue = "tutorial"
exchange = "amq.topic"
routing_keys = ["fedora.elections.#"]
```

You could also add "fedora.elections.#" to the "routing_keys" value in the existing entry.

Now make sure that RabbitMQ is still running, and run the `consume.py` script *we used before*. Make sure it is not systematically raising exceptions in the callback function (as we did before).

Now we'll run the election app, but first we need to create a configuration file. Create a file called `config.py` with the following content:

```
FEDORA_ELECTIONS_ADMIN_GROUP = ""
```

This will allow any Fedora account to be an admin on your instance, which is good enough for this tutorial. Now start the app with:

```
python createdb.py
python runserver.py -c config.py
```

Open your browser to <http://localhost:5000/admin/new>. Login with FAS, then create an election. Check the terminal where the `consume.py` script is running. You should see the message that the `elections` app has sent on election creation. Edit the election, and you should see the corresponding message in the terminal where `consume.py` is running.

14.1.2 Converting a Pyramid app

Let's use the `github2fedmsg` app as an example. It is a Pyramid webapp that registers a webhook with Github on all subscribed projects, and then broadcasts actions (commits, pull-request, tickets) received on this webhook to the message bus.

Clone the code using the following command:

```
git clone git@github.com:fedora-infra/github2fedmsg.git
```

And change to this directory.

JSON Schema

The only call to `fedmsg` is in `github2fedmsg/views/webhooks.py`. Since the app transmits the webhook payload almost transparently to the message bus, the structure isn't obvious, so it's harder to define a schema. Fortunately, the Github documentation has a [comprehensive list](#) of payload formats.

It would be too long to define precise JSON schemas for each event type, so we'll just use the generic schema.

Sending the messages

Now you can replace the current call to `fedmsg` with a call to `fedora_messaging.api.publish`. Add these lines in the `github2fedmsg.views.webhook` module:

```
import logging
from fedora_messaging.api import Message, publish
from fedora_messaging.exceptions import PublishReturned, PublishForbidden,
↳ ConnectionException

LOGGER = logging.getLogger(__name__)
```

And replace the call to `fedmsg.publish` with:

```
try:
    msg = Message(
        topic="github." + event_type,
        body=payload,
    )
    publish(msg)
except (PublishReturned, PublishForbidden) as e:
    LOGGER.warning(
        "Fedora Messaging broker rejected message %s: %s",
        msg.id, e
    )
except ConnectionException as e:
    LOGGER.warning("Error sending message %s: %s", msg.id, e)
```

Testing it

Make sure the Fedora Messaging configuration file is correct in `/etc/fedora-messaging/config.toml`. We will add a queue binding to route messages with the `github` topic to the `tutorial` queue. Add this entry in the `bindings` list:

```
[[bindings]]
queue = "tutorial"
exchange = "amq.topic"
routing_keys = ["github.#"]
```

You could also add `"github.#"` to the `"routing_keys"` value in the existing entry.

Now make sure that RabbitMQ is still running, and run the `consume.py` script *we used before*. Make sure it is not systematically raising exceptions in the callback function (as we did before).

To setup the `github2fedmsg` application, follow the `README.rst` file:

```
virtualenv venv
source ./venv/bin/activate
python setup.py develop
pip install waitress
```

Go off and [register your development application with GitHub](#). Save the oauth tokens and add the secret one to a new file you create called `secret.ini`. Use the example `secret.ini.example` file.

Create the database and start the application:

```
initialize_github2fedmsg_db development.ini
pserve development.ini --reload
```

14.2 Converting consumers

Let's use [the-new-hotness](#) app as an example. Clone the code and switch to state before conversion by using the following commands:

```
git clone https://github.com/fedora-infra/the-new-hotness.git
git checkout 0.10.1
```

And change to this directory.

In the-new-hotness app, all calls to consume messages on fedmsg are going through the `hotness.consumers.BugzillaTicketFiler.consume` method. We can thus modify this function to make it use Fedora Messaging instead of fedmsg.

14.2.1 Configuration

First we need to convert configuration file from fedmsg format to Fedora Messaging. Unlike fedmsg, fedora-messaging does not allow for arbitrary configuration keys.

The converted configuration `config.toml` could look like following:

```
# Define the callback function
# This will allow you to call only ``fedora-messaging consume`` without explicitly
# specifying the callback every time you starting ``the-new-hotness``.
callback = "hotness.consumers:BugzillaTicketFiler"

# In case of the-new-hotness we are listening to three topics, so we
# create a new binding for them
[[bindings]]
queue = "the-new-hotness"
exchange = "amq.topic"
routing_keys = [
    "org.release-monitoring.prod.anitya.project.version.update",
    "org.release-monitoring.prod.anitya.project.map.new",
    "org.fedoraproject.prod.buildsys.task.state.change",
]

# Define a queue
[queues.the-new-hotness]
durable = true
auto_delete = false
exclusive = false
arguments = {}
```

Any application specific configuration should go to `consumer_config` section [Configuration](#).

14.2.2 Init method

The `BugzillaTicketFiler` class in `consumers.py` is doing all the consuming work. First we need to change the inheritance of this class.

Then we need to modify the `__init__` method and use the `fedora_messaging.config.conf` dictionary instead of the `fedmsg` configuration. The `__init__` method could look something like this after the change:

```
from fedora_messaging.config import conf

class BugzillaTicketFiler:
    """
    A fedora-messaging consumer that is the heart of the-new-hotness.

    This consumer subscribes to the following topics:

    * 'org.fedoraproject.prod.buildsys.task.state.change'
      handled by :method:`BugzillaTicketFiler.handle_buildsys_scratch`

    * 'org.release-monitoring.prodanitya.project.version.update'
      handled by :method:`BugzillaTicketFiler.handle_anitya_version_update`

    * 'org.release-monitoring.prodanitya.project.map.new'
      handled by :method:`BugzillaTicketFiler.handle_anitya_map_new`
    """

    def __init__(self):
        # This is just convenient.
        self.config = conf["consumer_config"]

        ...
```

Note: Unrelated code was deleted from the example.

14.2.3 Wrapper function

The next step is to change `consume` method to `__call__` method. This is pretty straightforward. After this modification `__call__` method should look like this:

```
def __call__(self, msg):
    """
    Called when a message is received from queue.

    Params:
        msg (fedora_messaging.message.Message) The message we received
        from the queue.
    """
    topic, body, msg_id = msg.topic, msg.body, msg.id
    _log.debug("Received %r" % msg_id)
```

(continues on next page)

(continued from previous page)

```
if topic.endswith("anitya.project.version.update"):
    self.handle_anitya_version_update(msg)
elif topic.endswith("anitya.project.map.new"):
    self.handle_anitya_map_new(msg)
elif topic.endswith("buildsys.task.state.change"):
    self.handle_buildsys_scratch(msg)
else:
    _log.debug("Dropping %r %r" % (topic, body))
    pass
```

In this case we are working with the message using the standard `fedora_messaging.message.Message` methods. It is always better to use schema specific methods for any topic you are receiving.

14.2.4 Testing

To prepare the-new-hotness for testing checkout the `requirements.txt` file and `devel` folder from master branch:

```
git checkout master devel requirements.txt
```

This will convert development environment to the state when it's ready for Fedora Messaging. In a [configured development environment](#) we can easily test our conversion.

Start app by using alias `hotstart`, this will start the systemd service with following command `fedora-messaging consume`. The systemd unit could be find in `.config/systemd/user/`.

For testing you can use any message from [datagrepper](#). Just add `/raw?category=<application name>&delta=259200` to URL and pick any message. For example category for Anitya is `anitya`.

To send the message you need simple publisher. One is created for the new hotness in `devel/fedora_messaging_replay.py`. To send the message you can use any message id found in [datagrepper](#):

```
python3 devel/fedora_messaging_replay.py <msg_id>
```

And now you can check if the message was received using `hotlog` alias, which shows the journal log for the-new-hotness.

DEVELOPER INTERFACE

This documentation covers the public interfaces `fedora_messaging` provides.

Note: Documented interfaces follow [Semantic Versioning 2.0.0](#). Any interface not documented here may change at any time without warning.

API Table of Contents

- *Publishing*
 - *publish*
- *Subscribing*
 - *twisted_consume*
 - *Consumer*
 - *consume*
- *Signals*
 - *pre_publish_signal*
 - *publish_signal*
 - *publish_failed_signal*
- *Message Schemas*
 - *Message*
 - *Message Severity*
 - * *DEBUG*
 - * *INFO*
 - * *WARNING*
 - * *ERROR*
 - * *SEVERITIES*
 - *dumps*
 - *loads*
 - *SERIALIZED_MESSAGE_SCHEMA*

- *Utilities*
 - *libravatar_url*
- *Exceptions*
- *Configuration*
 - *conf*
 - *DEFAULTS*

15.1 Publishing

15.1.1 publish

`fedora_messaging.api.publish(message, exchange=None, timeout=30)`

Publish a message to an exchange.

This is a synchronous call, meaning that when this function returns, an acknowledgment has been received from the message broker and you can be certain the message was published successfully.

There are some cases where an error occurs despite your message being successfully published. For example, if a network partition occurs after the message is received by the broker. Therefore, you may publish duplicate messages. For complete details, see the [Publishing](#) documentation.

```
>>> from fedora_messaging import api
>>> message = api.Message(body={'Hello': 'world'}, topic='Hi')
>>> api.publish(message)
```

If an attempt to publish fails because the broker rejects the message, it is not retried. Connection attempts to the broker can be configured using the “connection_attempts” and “retry_delay” options in the broker URL. See [pika.connection.URLParameters](#) for details.

Parameters

- **message** (`message.Message`) – The message to publish.
- **exchange** (`str`) – The name of the AMQP exchange to publish to; defaults to *publish_exchange*
- **timeout** (`int`) – The maximum time in seconds to wait before giving up attempting to publish the message. If the timeout is reached, a `PublishTimeout` exception is raised.

Raises

- *fedora_messaging.exceptions.PublishReturned* – Raised if the broker rejects the message.
- *fedora_messaging.exceptions.PublishTimeout* – Raised if the broker could not be contacted in the given timeout time.
- *fedora_messaging.exceptions.PublishForbidden* – Raised if the broker rejects the message because of permission issues.
- *fedora_messaging.exceptions.ValidationError* – Raised if the message fails validation with its JSON schema. This only depends on the message you are trying to send, the AMQP server is not involved.

15.2 Subscribing

15.2.1 `twisted_consume`

`fedora_messaging.api.twisted_consume(callback, bindings=None, queues=None)`

Start a consumer using the provided callback and run it using the Twisted event loop (reactor).

Note: Callbacks run in a Twisted-managed thread pool using the `twisted.internet.threads.deferToThread()` API to avoid them blocking the event loop. If you wish to use Twisted APIs in your callback you must use the `twisted.internet.threads.blockingCallFromThread()` or `twisted.internet.interfaces.IReactorFromThreads` APIs.

This API expects the caller to start the reactor.

Parameters

- **callback** (*callable*) – A callable object that accepts one positional argument, a `Message` or a class object that implements the `__call__` method. The class will be instantiated before use.
- **bindings** (*dict or list of dict*) – Bindings to declare before consuming. This should be the same format as the `bindings` configuration.
- **queues** (*dict*) – The queue to declare and consume from. Each key in this dictionary should be a queue name to declare, and each value should be a dictionary with the “durable”, “auto_delete”, “exclusive”, and “arguments” keys.

Raises

ValueError – If the callback, bindings, or queues are invalid.

Returns

A deferred that fires with the list of one or more `Consumer` objects. Each consumer object has a `Consumer.result` instance variable that is a Deferred that fires or errors when the consumer halts. Note that this API is meant to survive network problems, so consuming will continue until `Consumer.cancel()` is called or a fatal server error occurs. The deferred returned by this function may error back with a `fedora_messaging.exceptions.BadDeclaration` if queues or bindings cannot be declared on the broker, a `fedora_messaging.exceptions.PermissionException` if the user doesn’t have access to the queue, or `fedora_messaging.exceptions.ConnectionException` if the TLS or AMQP handshake fails.

Return type

`twisted.internet.defer.Deferred`

15.2.2 Consumer

`class fedora_messaging.api.Consumer(queue=None, callback=None)`

Represents a Twisted AMQP consumer and is returned from the call to `fedora_messaging.api.twisted_consume()`.

queue

The AMQP queue this consumer is subscribed to.

Type

`str`

callback

The callback to run when a message arrives.

Type

callable

result

A deferred that runs the callbacks if the consumer exits gracefully after being canceled by a call to `Consumer.cancel()` and errbacks if the consumer stops for any other reason. The reasons a consumer could stop are: a `fedora_messaging.exceptions.PermissionException` if the consumer does not have permissions to read from the queue it is subscribed to, a `HaltConsumer` is raised by the consumer indicating it wishes to halt, an unexpected `Exception` is raised by the consumer, or if the consumer is canceled by the server which happens if the queue is deleted by an administrator or if the node the queue lives on fails.

Type

`twisted.internet.defer.Deferred`

cancel()

Cancel the consumer and clean up resources associated with it. Consumers that are canceled are allowed to finish processing any messages before halting.

Returns

A deferred that fires when the consumer has finished processing any message it was in the middle of and has been successfully canceled.

Return type

`defer.Deferred`

15.2.3 consume

`fedora_messaging.api.consume(callback, bindings=None, queues=None)`

Start a message consumer that executes the provided callback when messages are received.

This API is blocking and will not return until the process receives a signal from the operating system.

Warning: This API runs the callback in the IO loop thread. This means if your callback could run for a length of time near the heartbeat interval, which is likely on the order of 60 seconds, the broker will kill the TCP connection and the message will be re-delivered on start-up.

For now, use the `twisted_consume()` API which runs the callback in a thread and continues to handle AMQP events while the callback runs if you have a long-running callback.

The callback receives a single positional argument, the message:

```
>>> from fedora_messaging import api
>>> def my_callback(message):
...     print(message)
>>> bindings = [{'exchange': 'amq.topic', 'queue': 'demo', 'routing_keys': ['#']}]
>>> queues = {
...     "demo": {"durable": False, "auto_delete": True, "exclusive": True,
↪ "arguments": {}}
... }
>>> api.consume(my_callback, bindings=bindings, queues=queues)
```


If the bindings and queue arguments are not provided, they will be loaded from the configuration.

For complete documentation on writing consumers, see the *Consumers* documentation.

Parameters

- **callback** (*callable*) – A callable object that accepts one positional argument, a *Message* or a class object that implements the `__call__` method. The class will be instantiated before use.
- **bindings** (*dict or list of dict*) – Bindings to declare before consuming. This should be the same format as the *bindings* configuration.
- **queues** (*dict*) – The queue or queues to declare and consume from. This should be in the same format as the *queues* configuration dictionary where each key is a queue name and each value is a dictionary of settings for that queue.

Raises

- *fedora_messaging.exceptions.HaltConsumer* – If the consumer requests that it be stopped.
- *ValueError* – If the consumer provides a callback that is not a class that implements `__call__` and is not a function, if the bindings argument is not a dict or list of dicts with the proper keys, or if the queues argument isn't a dict with the proper keys.

15.3 Signals

Signals sent by fedora_messaging APIs using `blinker.base.Signal` signals.

15.3.1 pre_publish_signal

```
fedora_messaging.api.pre_publish_signal = <blinker.base.NamedSignal object at 0x7f64018e7620; 'pre_publish'>
```

A signal triggered before the message is published. The signal handler should accept a single keyword argument, `message`, which is the instance of the *fedora_messaging.message.Message* being sent. It is acceptable to mutate the message, but the `validate` method will be called on it after this signal.

15.3.2 publish_signal

```
fedora_messaging.api.publish_signal = <blinker.base.NamedSignal object at 0x7f64028f7110; 'publish_success'>
```

A signal triggered after a message is published successfully. The signal handler should accept a single keyword argument, `message`, which is the instance of the *fedora_messaging.message.Message* that was sent.

15.3.3 publish_failed_signal

`fedora_messaging.api.publish_failed_signal = <blinker.base.NamedSignal object at 0x7f6401661430; 'publish_failed_signal'>`

A signal triggered after a message fails to publish for some reason. The signal handler should accept two keyword argument, `message`, which is the instance of the `fedora_messaging.message.Message` that failed to be sent, and `error`, the exception that was raised.

15.4 Message Schemas

This module defines the base class of message objects and keeps a registry of known message implementations. This registry is populated from Python entry points in the “fedora.messages” group.

To implement your own message schema, simply create a class that inherits the `Message` class, and add an entry point in your Python package under the “fedora.messages” group. For example, an entry point for the `Message` schema would be:

```
entry_points = {
    'fedora.messages': [
        'base.message=fedora_messaging.message:Message'
    ]
}
```

The entry point name must be unique to your application and is used to map messages to your message class, so it’s best to prefix it with your application name (e.g. `bodhi.new_update_messageV1`). When publishing, the Fedora Messaging library will add a header with the entry point name of the class used so the consumer can locate the correct schema.

Since every client needs to have the message schema installed, you should define this class in a small Python package of its own.

15.4.1 Message

```
class fedora_messaging.message.Message(body=None, headers=None, topic=None, properties=None,
                                       severity=None)
```

Messages are simply JSON-encoded objects. This allows message authors to define a schema and implement Python methods to abstract the raw message from the user. This allows the schema to change and evolve without breaking the user-facing API.

There are a number of properties that are intended to be overridden by users. These fields are used to sort messages for notifications or are used to create human-readable versions of the messages. Properties that are intended for this purpose are noted in their attribute documentation below.

Parameters

- **headers** (*dict*) – A set of message headers. Consult the headers schema for expected keys and values.
- **body** (*dict*) – The message body. Consult the body schema for expected keys and values. This dictionary must be JSON-serializable by the default serializer.
- **topic** (*str*) – The message topic as a unicode string. If this is not provided, the default topic for the class is used. See the attribute documentation below for details.

- **properties** (*pika.BasicProperties*) – The AMQP properties. If this is not provided, they will be generated. Most users should not need to provide this, but it can be useful in testing scenarios.
- **severity** (*int*) – An integer that indicates the severity of the message. This is used to determine what messages to notify end users about and should be *DEBUG*, *INFO*, *WARNING*, or *ERROR*. The default is *INFO*, and can be set as a class attribute or on an instance-by-instance basis.

id

The message id as a unicode string. This attribute is automatically generated and set by the library and users should only set it themselves in testing scenarios.

Type

str

topic

The message topic as a unicode string. The topic is used by message consumers to filter what messages they receive. Topics should be a string of words separated by ‘.’ characters, with a length limit of 255 bytes. Because of this byte limit, it is best to avoid non-ASCII character. Topics should start general and get more specific each word. For example: “bodhi.update.kernel” is a possible topic. “bodhi” identifies the application, “update” identifies the message, and “kernel” identifies the package in the update. This can be set at a class level or on an instance level. Dynamic, specific topics that allow for fine-grain filtering are preferred.

Type

str

headers_schema

A JSON schema to be used with `jsonschema.validate()` to validate the message headers. For most users, the default definition should suffice.

Type

dict

body_schema

A JSON schema to be used with `jsonschema.validate()` to validate the message body. The `body_schema` is retrieved on a message instance so it is not required to be a class attribute, although this is a convenient approach. Users are also free to write the JSON schema as a file and load the file from the filesystem or network if they prefer.

Type

dict

body

The message body as a Python dictionary. This is validated by the body schema before publishing and before consuming.

Type

dict

severity

An integer that indicates the severity of the message. This is used to determine what messages to notify end users about and should be *DEBUG*, *INFO*, *WARNING*, or *ERROR*. The default is *INFO*, and can be set as a class attribute or on an instance-by-instance basis.

Type

int

queue

The name of the queue this message arrived through. This attribute is set automatically by the library and users should never set it themselves.

Type

str

deprecated

Whether this message schema has been deprecated by a more recent version. Emits a warning when a message of this class is received, to let consumers know that they should plan to upgrade. Defaults to False.

Type

bool

priority

The priority for the message, if the destination queue supports it. Defaults to zero (lowest priority).

This value is taken into account in queues that have the `x-max-priority` argument set. Most queues in Fedora don't support priorities, in which case the value will be ignored.

Larger numbers indicate higher priority, you can read more about it in [RabbitMQ's documentation on priority](#).

Type

int

__str__()

A human-readable representation of this message.

This should provide a detailed, long-form representation of the message. The default implementation is to format the raw message id, topic, headers, and body.

Note: Sub-classes should override this method. It is used to create the body of email notifications and by other tools to display messages to humans.

property agent_avatar

An URL to the avatar of the user who caused the action.

Note: Sub-classes should override this method if the default Libravatar and OpenID-based URL generator is not appropriate.

Returns

The URL to the user's avatar.

Return type

str or None

property agent_name

The username of the user who caused the action.

Note: Sub-classes should override this method if the message was triggered by a particular user.

Returns

The agent's username.

Return type

`str` or `None`

property app_icon

An URL to the icon of the application that generated the message.

Note: Sub-classes should override this method if their application has an icon and they wish that image to appear in applications that consume messages.

Returns

The URL to the app's icon.

Return type

`str` or `None`

property app_name

The name of the application that generated the message.

Note: Sub-classes should override this method.

Returns

The name of the application.

Return type

`str` or `None`

property containers

List of containers affected by the action that generated this message.

Note: Sub-classes should override this method if the message pertains to one or more container images. The data returned from this property is used to filter notifications.

Returns

A list of affected container names.

Return type

`list(str)`

property flatpaks

List of flatpaks affected by the action that generated this message.

Note: Sub-classes should override this method if the message pertains to one or more flatpaks. The data returned from this property is used to filter notifications.

Returns

A list of affected flatpaks names.

Return type

`list(str)`

property groups

List of groups affected by the action that generated this message.

Note: Sub-classes should override this method if the message pertains to a group or groups. The data returned from this property is used to filter notifications.

Returns

A list of affected groups.

Return type

`list(str)`

property modules

List of modules affected by the action that generated this message.

Note: Sub-classes should override this method if the message pertains to one or more modules. The data returned from this property is used to filter notifications.

Returns

A list of affected module names.

Return type

`list(str)`

property packages

List of RPM packages affected by the action that generated this message.

Note: Sub-classes should override this method if the message pertains to one or more RPM packages. The data returned from this property is used to filter notifications.

Returns

A list of affected package names.

Return type

`list(str)`

property summary

A short, human-readable representation of this message.

This should provide a short summary of the message, much like the subject line of an email.

Note: Sub-classes should override this method. It is used to create the subject of email notifications, IRC notification, and by other tools to display messages to humans in short form.

The default implementation is to simply return the message topic.

property url

An URL to the action that caused this message to be emitted.

Note: Sub-classes should override this method if there is a URL associated with message.

Returns

A relevant URL.

Return type

`str` or `None`

property usernames

List of users affected by the action that generated this message.

Note: Sub-classes should override this method if the message pertains to a user or users. The data returned from this property is used to filter notifications.

Returns

A list of affected usernames.

Return type

`list(str)`

validate()

Validate the headers and body with the message schema, if any.

In addition to the user-provided schema, all messages are checked against the base schema which requires certain message headers and the that body be a JSON object.

<p>Warning: This method should not be overridden by sub-classes.</p>

Raises

- **`jsonschema.ValidationError`** – If either the message headers or the message body are invalid.
- **`jsonschema.SchemaError`** – If either the message header schema or the message body schema are invalid.

15.4.2 Message Severity

Each message can have a severity associated with it. The severity is used by applications like the notification service to determine what messages to send to users. The severity can be set at the class level, or on a message-by-message basis. The following are valid severity levels:

DEBUG

`fedora_messaging.message.DEBUG = 10`

Indicates the message is for debugging or is otherwise very low priority. Users will not be notified unless they've explicitly requested DEBUG level messages.

INFO

`fedora_messaging.message.INFO = 20`

Indicates the message is informational. End users will not receive notifications for these messages by default. For example, automated tests passed for their package.

WARNING

`fedora_messaging.message.WARNING = 30`

Indicates a problem or an otherwise important problem. Users are notified of these messages when they pertain to packages they are associated with by default. For example, one or more automated tests failed against their package.

ERROR

`fedora_messaging.message.ERROR = 40`

Indicates a critically important message that users should act upon as soon as possible. For example, their package no longer builds.

SEVERITIES

`fedora_messaging.message.SEVERITIES = (10, 20, 30, 40)`

A tuple of all valid severity levels

15.4.3 dumps

`fedora_messaging.message.dumps(messages)`

Serialize messages to a file format acceptable for `loads()` or for the publish CLI command. The format is a string where each line is a JSON object that conforms to the `SERIALIZED_MESSAGE_SCHEMA` format.

Parameters

messages (*list* or *Message*) – The messages to serialize. Each message in the messages is subclass of Message.

Returns

Serialized messages.

Return type

`str`

Raises

ValidationError – If one of the messages provided doesn't conform to its schema.

15.4.4 loads

`fedora_messaging.message.loads(serialized_messages)`

Deserialize messages from a file format produced by `dumps()`. The format is a string where each line is a JSON object that conforms to the `SERIALIZED_MESSAGE_SCHEMA` format.

Parameters

serialized_messages (*str*) – A string made up of a JSON object per line.

Returns

Deserialized message objects.

Return type

list

Raises

ValidationError – If the string isn't formatted properly or message doesn't pass the message schema validation

15.4.5 SERIALIZED_MESSAGE_SCHEMA

```
fedora_messaging.message.SERIALIZED_MESSAGE_SCHEMA = {'$schema':
'http://json-schema.org/draft-04/schema#', 'description': 'Schema for the JSON object
used to represent messages in a file', 'properties': {'body': {'description': 'The
message body.', 'type': 'object'}, 'headers': {'description': 'The message headers',
'properties': {'fedora_messaging_schema': {'type': 'string'},
'fedora_messaging_severity': {'enum': [10, 20, 30, 40], 'type': 'number'}, 'sent-at':
{'type': 'string'}}, 'type': 'object'}, 'id': {'description': 'The message's UUID.',
'type': 'string'}, 'priority': {'description': 'The priority that the message has been
sent with.', 'type': ['integer', 'null']}, 'queue': {'description': 'The queue the
message arrived on, if any.', 'type': ['string', 'null']}, 'topic': {'description':
'The message topic', 'type': 'string'}}, 'required': ['topic', 'body'], 'type':
'object'}
```

The schema for each JSON object produced by `dumps()`, consumed by `loads()`, and expected by CLI commands like “fedora-messaging publish”.

15.5 Utilities

The `schema_utils` module contains utilities that may be useful when writing the Python API of your message schemas.

15.5.1 libravator_url

`fedora_messaging.schema_utils.libravator_url(email=None, openid=None, size=64, default='retro')`

Get the URL to an avatar from libravator.

Either the user's email or openid must be provided.

If you want to use Libravator federation (through DNS), you should install and use the libravator library instead. Check out the `libravator.libravator_url()` function.

Parameters

- **email** (*str*) – The user's email

- **openid** (*str*) – The user’s OpenID
- **size** (*int*) – Size of the avatar in pixels (it’s a square).
- **default** (*str*) – Default avatar to return if not found.

Returns

The URL to the avatar image.

Return type

str

Raises

ValueError – If neither email nor openid are provided.

15.6 Exceptions

Exceptions raised by Fedora Messaging.

exception `fedora_messaging.exceptions.BadDeclaration`(*obj_type=None, description=None, reason=None*)

Raised when declaring an object in AMQP fails.

Parameters

- **obj_type** (*str*) – The type of object being declared. One of “binding”, “queue”, or “exchange”.
- **description** (*dict*) – The description of the object.
- **reason** (*str*) – The reason the server gave for rejecting the declaration.

exception `fedora_messaging.exceptions.BaseException`

The base class for all exceptions raised by `fedora_messaging`.

exception `fedora_messaging.exceptions.ConfigurationException`(*message*)

Raised when there’s an invalid configuration setting

Parameters

message (*str*) – A detailed description of the configuration problem which is presented to the user.

exception `fedora_messaging.exceptions.ConnectionException`(*args, **kwargs)

Raised if a general connection error occurred.

You may handle this exception by logging it and resending or discarding the message.

exception `fedora_messaging.exceptions.ConsumeException`

Base class for exceptions related to consuming.

exception `fedora_messaging.exceptions.ConsumerCanceled`

Raised when the server has canceled the consumer.

This can happen when the queue the consumer is subscribed to is deleted, or when the node the queue is located on fails.

exception `fedora_messaging.exceptions.Drop`

Consumer callbacks should raise this to indicate they wish the message they are currently processing to be dropped.

exception `fedora_messaging.exceptions.HaltConsumer`(*exit_code=0, reason=None, requeue=False, **kwargs*)

Consumer callbacks should raise this exception if they wish the consumer to be shut down.

Parameters

- **exit_code** (*int*) – The exit code to use when halting.
- **reason** (*str*) – A reason for halting, presented to the user.
- **requeue** (*bool*) – If true, the message is re-queued for later processing.

exception `fedora_messaging.exceptions.Nack`

Consumer callbacks should raise this to indicate they wish the message they are currently processing to be re-queued.

exception `fedora_messaging.exceptions.NoFreeChannels`

Raised when a connection has reached its channel limit

exception `fedora_messaging.exceptions.PermissionException`(*obj_type=None, description=None, reason=None*)

Generic permissions exception.

Parameters

- **obj_type** (*str*) – The type of object being accessed that caused the permission error. May be None if the cause is unknown.
- **description** (*object*) – The description of the object, if any. May be None.
- **reason** (*str*) – The reason the server gave for the permission error, if any. If no reason is supplied by the server, this should be the best guess for what caused the error.

exception `fedora_messaging.exceptions.PublishException`(*reason=None, **kwargs*)

Base class for exceptions related to publishing.

exception `fedora_messaging.exceptions.PublishForbidden`(*reason=None, **kwargs*)

Raised when the broker rejects the message due to permission errors.

You may handle this exception by logging it and discarding the message, as it is likely a permanent error.

exception `fedora_messaging.exceptions.PublishReturned`(*reason=None, **kwargs*)

Raised when the broker rejects and returns the message to the publisher.

You may handle this exception by logging it and resending or discarding the message.

exception `fedora_messaging.exceptions.PublishTimeout`(*reason=None, **kwargs*)

Raised when the message could not be published in the given timeout.

This means the message was either never delivered to the broker, or that it was delivered, but never acknowledged by the broker.

exception `fedora_messaging.exceptions.ValidationError`

This error is raised when a message fails validation with its JSON schema

This exception can be raised on an incoming or outgoing message. No need to catch this exception when publishing, it should warn you during development and testing that you're trying to publish a message with a different format, and that you should either fix it or update the schema.

property summary

A short summary of the error.

15.7 Configuration

15.7.1 conf

```
fedora_messaging.config.conf = {}
```

The configuration dictionary used by fedora-messaging and consumers.

15.7.2 DEFAULTS

```
fedora_messaging.config.DEFAULTS = {'amqp_url':
'amqp://?connection_attempts=3&retry_delay=5', 'bindings': [{'exchange': 'amq.topic',
'queue': '', 'routing_keys': ['#']}], 'callback': None, 'client_properties': {'app':
'Unknown', 'information': 'https://fedora-messaging.readthedocs.io/en/stable/',
'product': 'Fedora Messaging with Pika', 'version': 'fedora_messaging-3.5.0 with
pika-1.3.2'}, 'consumer_config': {}, 'exchanges': {'amq.topic': {'arguments': {}},
'auto_delete': False, 'durable': True, 'type': 'topic'}}, 'log_config':
{'disable_existing_loggers': False, 'formatters': {'simple': {'format': '%(name)s
%(levelname)s] %(message)s'}}, 'handlers': {'console': {'class':
'logging.StreamHandler', 'formatter': 'simple', 'stream': 'ext://sys.stdout'}},
'loggers': {'fedora_messaging': {'handlers': ['console'], 'level': 'INFO',
'propagate': False}}, 'root': {'handlers': ['console'], 'level': 'WARNING'},
'version': 1}, 'passive_declares': False, 'publish_exchange': 'amq.topic',
'publish_priority': None, 'qos': {'prefetch_count': 10, 'prefetch_size': 0},
'queues': {'': {'arguments': {}}, 'auto_delete': True, 'durable': False, 'exclusive':
True}}, 'tls': {'ca_cert': None, 'certfile': None, 'keyfile': None}, 'topic_prefix':
''}
```

The default configuration settings for fedora-messaging. This should not be modified and should be copied with `copy.deepcopy()`.

MESSAGE FORMAT

This documentation covers the format of AMQP messages sent by this library. If you are interested in using a language other than Python to send or receive messages sent by Fedora applications, this document is for you.

Messages are AMQP [Basic](#) content. Basic messages have the content type, content encoding, a table of headers, delivery mode, priority, correlation ID, reply-to, expiration, message ID, timestamp, type, user ID, and app ID fields.

16.1 Content Type

Your messages *MUST* have a content-type of `application/json` and they must be JSON objects. Consult the [Message Schemas](#) documentation for details on message format.

16.2 Content Encoding

Your message *MUST* have the content-encoding property set to `utf-8` and they must be encoding with UTF-8.

16.3 Message ID

The message ID field *MUST* be a [version 4 UUID](#) as a standard hexadecimal digit string (e.g. `f81d4fae-7dec-11d0-a765-00a0c91e6bf6`).

16.4 Delivery Mode

The delivery mode of your message *SHOULD* be 2 (persistent) unless you know what you are doing and have a very good reason for setting it to 1 (transient).

16.5 Headers

The headers field of AMQP message allows you to set a dictionary (map) of arbitrary strings. Several header keys are used by Fedora's applications to determine the message schema, the importance of the message for human beings, when it was originally sent by the application, what packages or users it relates to, and more.

16.5.1 Required

Messages must have, at a minimum, the `fedora_messaging_severity`, `fedora_messaging_schema`, and `sent-at` keys.

The `fedora_messaging_severity` key should be set to an integer that indicates the importance of the message to an end user, with 10 being debug-level information, 20 being informational, 30 being warning-level, and 40 being critically important.

The `fedora_messaging_schema` key should be set to a string that uniquely identifies the type of message. In the Python library this is the entry point name, which is mapped to a class containing the schema and a Python API to interact with the message object.

The `sent-at` key should be a ISO8601 date time that should include the UTC offset and should *not* include microseconds. For example: `2019-07-30T19:12:22+00:00`.

The header's json-schema is:

```
{
  "$schema": "http://json-schema.org/draft-04/schema#",
  "description": "Schema for message headers",
  "type": "object",
  "properties": {
    "fedora_messaging_severity": {
      "type": "number",
      "enum": [10, 20, 30, 40],
    },
    "fedora_messaging_schema": {"type": "string"},
    "sent-at": {"type": "string"},
  },
}
```

16.5.2 Optional

In addition to the required headers, there are a number of optional headers you can set that have special meaning. The general format of these headers is `fedora_messaging_<object>_<id>` where the `<object>` is one of `user`, `rpm`, `container`, `module`, or `flatpak` and `<id>` uniquely identifies the object. Set these headers when the message pertains to the referenced object.

For example, if the user `jcline` submitted a build for the `python-requests` RPM, the message about that event would have `fedora_messaging_user_jcline` and `fedora_messaging_rpm_python-requests` set.

At this time the value of the header key is not used and should always be set to a Boolean value of `true`.

16.6 Body

The message body must match the content-type and content-encoding. That is, it must be UTF-8 encoded JSON. Additionally, it must be a JSON Object. Beyond that, there are no restrictions. Messages should be validated using their JSON schema. If you are publishing a new message type, please write a json-schema for it and provide it to the Fedora infrastructure team. It will be distributed to applications that wish to consume the message.

CONTRIBUTOR GUIDE

Thanks for considering contributing to fedora-messaging, we really appreciate it!

17.1 Quickstart

1. Look for an [existing issue](#) about the bug or feature you're interested in. If you can't find an existing issue, create a [new one](#).
2. Fork the [repository on GitHub](#).
3. Fix the bug or add the feature, and then write one or more tests which show the bug is fixed or the feature works.
4. Add a [news fragment](#) with a summary of the change to include in the upcoming release notes.
5. Submit a pull request and wait for a maintainer to review it.

More detailed guidelines to help ensure your submission goes smoothly are below.

Note: If you do not wish to use GitHub, please send patches to infrastructure@lists.fedoraproject.org.

17.2 Python Support

fedora-messaging supports Python 3.6 or greater. This is automatically enforced by the continuous integration (CI) suite.

17.3 Code Style

We follow the [PEP8](#) style guide for Python. This is automatically enforced by the CI suite.

We are using *Black* <<https://github.com/ambv/black>> to automatically format the source code. It is also checked in CI. The Black webpage contains instructions to configure your editor to run it on the files you edit.

We use [pre-commit](#) to run a set of linters and formatters upon commit. To setup to hook for your repo clone, install `pre-commit` and run `pre-commit install`.

17.4 Tests

The test suites can be run using `tox` by simply running `tox` from the repository root. All code must have test coverage or be explicitly marked as not covered using the `# no-qa` comment. This should only be done if there is a good reason to not write tests.

Your pull request should contain tests for your new feature or bug fix. If you're not certain how to write tests, we will be happy to help you.

17.5 Release notes

To add entries to the release notes, run `towncrier create <source.type>` to create a news fragment file in the `news` directory, where `type` is one of:

- `feature`: for new features
- `bug`: for bug fixes
- `api`: for API changes
- `dev`: for development-related changes
- `docs`: for documentation changes
- `author`: for contributor names
- `other`: for other changes

And where the `source` part of the filename is:

- `42` when the change is described in issue 42
- `PR42` when the change has been implemented in pull request 42, and there is no associated issue
- `Cabcdef` when the change has been implemented in changeset `abcdef`, and there is no associated issue or pull request.
- `username` for contributors (author extension). It should be the username part of their commits' email address.

For example,

```
towncrier create PR42.feature
```

The contents of the news fragment must be written in RST format. See the [towncrier documentation for more information](#).

A preview of the release notes can be generated with `towncrier --draft`.

17.6 Licensing

Your commit messages must include a Signed-off-by tag with your name and e-mail address, indicating that you agree to the [Developer Certificate of Origin](#) version 1.1:

```
Developer Certificate of Origin
Version 1.1

Copyright (C) 2004, 2006 The Linux Foundation and its contributors.
```

(continues on next page)

(continued from previous page)

1 Letterman Drive

Suite D4700

San Francisco, CA, 94129

Everyone **is** permitted to copy **and** distribute verbatim copies of this license document, but changing it **is not** allowed.

Developer's Certificate of Origin 1.1

By making a contribution to this project, I certify that:

- (a) The contribution was created **in** whole **or in** part by me **and** I have the right to submit it under the **open** source license indicated **in** the file; **or**
- (b) The contribution **is** based upon previous work that, to the best of my knowledge, **is** covered under an appropriate **open** source license **and** I have the right under that license to submit that work **with** modifications, whether created **in** whole **or in** part by me, under the same **open** source license (unless I am permitted to submit under a different license), **as** indicated **in** the file; **or**
- (c) The contribution was provided directly to me by some other person who certified (a), (b) **or** (c) **and** I have **not** modified it.
- (d) I understand **and** agree that this project **and** the contribution are public **and** that a record of the contribution (including **all** personal information I submit **with** it, including my sign-off) **is** maintained indefinitely **and** may be redistributed consistent **with** this project **or** the **open** source license(s) involved.

Use `git commit -s` to add the Signed-off-by tag.

17.7 Releasing

When cutting a new release, follow these steps:

- update the version in `pyproject.toml`
- add missing authors to the release notes fragments by changing to the `news` directory and running the `get-authors.py` script, but check for duplicates and errors
- generate the changelog by running `poetry run towncrier build`
- adjust the release notes in `docs/changelog.md`
- generate the docs with `tox -e docs` and check them in `docs/_build/html`
- change the Development Status classifier in `pyproject.toml` if necessary
- commit the changes

- push the commit to the upstream Github repository (via a PR or not).
- change to the stable branch and merge the `develop` branch
- run the checks with `tox` one last time to be sure
- tag the commit with `-s` to generate a signed tag
- push the commit to the upstream Github repository with `git push` and the new tag with `git push --tags`
- generate a tarball and push to PyPI with the command `poetry publish --build`
- create [the release on GitHub](#) and copy the release notes in there
- deploy and announce

RELEASE NOTES

All notable changes to this project will be documented in this file.

The format is based on [Keep a Changelog](#), and this project adheres to [Semantic Versioning](#).

This project uses [towncrier](#) and the changes for the upcoming release can be found in [the news directory](#).

18.1 3.5.0 (2024-03-20)

18.1.1 Features

- Add a replay command ([#332](#))
- Add support Python 3.11 and 3.12, drop support for Python 3.6 and 3.7
- Better protection against invalid bodies breaking the headers generation and the instantiation of a message
- Testing framework: make the sent messages available in the context manager

18.1.2 Documentation Improvements

- Add SECURITY.md for project security policy (PR [#314](#))
- Add fedora-messaging-git-hook-messages to the known schema packages

18.1.3 Development Changes

- Make the tests use the pytest fixtures and assert system ([961b82d](#))
- Make fedora-messaging use poetry ([#294](#))
- Add some generic pre-commit checks
- Don't distribute the tests in the wheel

18.1.4 Contributors

Many thanks to the contributors of bug reports, pull requests, and pull request reviews for this release:

- Aurélien Bompard
- Thibaut Batale
- Khaled Achech
- Lenka Segura
- Ryan Lerch

18.2 3.4.1 (2023-05-26)

18.2.1 Bug Fixes

- Fix CI ([0f2e39c](#))

18.3 3.4.0 (2023-05-26)

18.3.1 Features

- Mirror the message priority in the headers ([eba336b](#))

18.4 3.3.0 (2023-03-31)

18.4.1 Features

- Add support for asyncio-based callbacks in the consumer. As a consequence, the Twisted reactor used by the CLI is now `asyncio`. (PR [#282](#))

18.4.2 Documentation Improvements

- Improve documentation layout, and add some documentation on the message schemas. ([1fa8998](#))
- Add `koji-fedoramessaging-messages` to the list of known schemas. ([ef12fa2](#))

18.4.3 Development Changes

- Update pre-commit linters. ([0efdde1](#))

18.5 3.2.0 (2022-10-17)

18.5.1 Features

- Add a priority property to messages, and a default priority in the configuration (PR #275)
- Add a message schema attribute and some documentation to help deprecate and upgrade message schemas (#227)

18.5.2 Other Changes

- Use tomllib from the standard library on Python 3.11 and above, fallback to tomli otherwise. (PR #274)

18.5.3 Contributors

Many thanks to the contributors of bug reports, pull requests, and pull request reviews for this release:

- Akashdeep Dhar
- Aurélien Bompard
- Erol Keskin
- Miro Hrončok
- Stephen Coady

18.6 3.1.0 (2022-09-13)

18.6.1 Features

- Add the app_name and agent_name properties to message schemas (PR #272)
- Added “groups” property to message schemas. This property can be used if an event affects a list of groups. (#252)

18.7 3.0.2 (2022-05-19)

18.7.1 Development Changes

- Fix CI in Github actions (6257100)
- Update pre-commit checkers (1d35a5d)
- Fix Packit configuration (d2ea85f)

18.7.2 Contributors

Many thanks to the contributors of bug reports, pull requests, and pull request reviews for this release:

- Akashdeep Dhar
- Aurélien Bompard

18.8 3.0.1 (2022-05-12)

18.8.1 Development Changes

- Add packit configuration allowing us to have automatic downstream RPM builds ([#259](#))
- Don't build universal wheels since we don't run on Python 2 anymore ([e8c5f4c](#))

18.8.2 Documentation Improvements

- Add some schema packages to the docs ([03e7f42](#))
- Change the example email addresses ([1555742](#))

18.8.3 Contributors

Many thanks to the contributors of bug reports, pull requests, and pull request reviews for this release:

- Akashdeep Dhar
- Aurélien Bompard

18.9 3.0.0 (2021-12-14)

18.9.1 API Changes

- Queues created by the CLI are now non-durable, auto-deleted and exclusive, as server-named queues are. (PR [#239](#))
- It is no longer necessary to declare a queue in the configuration file: a server-named queue will be created. Configured bindings which do not specify a queue name will be applied to the server-named queue. (PR [#239](#))
- Drop support for Python 2 (PR [#246](#))
- Drop the Twisted classes that had been flagged as deprecated. Drop the deprecated `Message._body` property. Refactor the consuming code into the `Consumer` class. (PR [#249](#))

18.9.2 Features

- Support anonymous (server-named) queues. (PR #239)
- Support Python 3.10 (PR #250)
- Raise `PublishForbidden` exception immediately if publishing to `virtual host` is denied rather than waiting until timeout occurs. (#203)

18.9.3 Bug Fixes

- Fixed validation exception of queue field on serialized schemas. (#240)

18.9.4 Documentation Improvements

- Improve release notes process documentation. (PR #238)
- Build a list of available topics in the documentation from known schema packages (PR #242)

18.9.5 Development Changes

- Start using pre-commit for linters and formatters (732c7fb)

18.9.6 Contributors

Many thanks to the contributors of bug reports, pull requests, and pull request reviews for this release:

- Aurélien Bompard
- David Jimenez
- Michal Konečný
- Onur Ozkan

18.10 2.1.0 (2021-05-12)

18.10.1 Features

- Improve the testing module to check message topics and bodies separately, and to use the rewritten assert that pytest provides (PR #230)
- Handle `topic authorization <https://www.rabbitmq.com/access-control.html#topic-authorisation>` by raising a `PublishForbidden` exception instead of being stuck in a retry loop (PR #235)
- Test on Python 3.8 and 3.9 (PR #237)

18.10.2 Bug Fixes

- Require `setuptools`, as `pkg_resources` is used (PR #233)

18.10.3 Development Changes

- Update test fixture keys to 4096 bits (PR #232)
- Use Github Actions for CI (PR #234)

18.10.4 Contributors

Many thanks to the contributors of bug reports, pull requests, and pull request reviews for this release:

- Aurélien Bompard
- Jeremy Cline
- Miro Hrončok
- Pierre-Yves Chibon

18.11 2.0.2 (2020-08-04)

18.11.1 Bug Fixes

- Set the QoS on the channel that is created for the consumer (#223)

18.11.2 Documentation Improvements

- When running `fedora-messaging consume`, the callback module should not contain a call to `api.consume()` or it will block. (df4055f)
- Update the schema docs (PR #219)
- Fix quickstart cert file links (PR #222)
- Fix the docs about exceptions being wrapped by `HaltConsumer` (#215)

18.11.3 Other Changes

- Only try to restart `fm-consumer@` services every 60 seconds (PR #214)

18.12 2.0.1 (2020-01-02)

18.12.1 Bug Fixes

- Fix handling of new connections after a publish timeout ([#212](#))

18.13 2.0.0 (2019-12-03)

18.13.1 Dependency Changes

- Drop official Python 3.4 and 3.5 support
- Bump the pika requirement to 1.0.1+
- New dependency: [Crochet](#)

18.13.2 API Changes

- Move all APIs to use the Twisted-managed connection. There are a few minor changes here which slightly change the APIs:
 1. Publishing now raises a `PublishTimeout` when the timeout is reached (30 seconds by default).
 2. Previously, the Twisted consume API did not validate arguments like the synchronous version did, so it now raises a `ValueError` on invalid arguments instead of crashing in some undefined way.
 3. Calling `publish` from the Twisted reactor thread now raises an exception instead of blocking the reactor thread.
 4. Consumer exceptions are not re-raised as `HaltingConsumer` exceptions anymore, the original exception bubbles up and has to be handled by the application.

18.13.3 Features

- The `fedora-messaging` cli now has 2 new sub-commands: `publish` and `record`. (PR [#43](#))
- Log the failure traceback on connection ready failures.

18.13.4 Bug Fixes

- Fix an issue where reconnection to the server would fail. ([#208](#))
- Don't declare exchanges when consuming. ([#171](#))
- Fix Twisted legacy logging (it does not accept format parameters).
- Handle `ConnectionLost` errors in the v2 Factory.

18.13.5 Development Changes

- Many Twisted-related tests were added.
- Include tests for sample schema package.
- Update the dumps and loads functions for a new message format.

18.13.6 Documentation Improvements

- Document that logging is only set up for consumers.
- Update the six intersphinx URL to fix the docs build.
- Add the “conf” and “DEFAULTS” variables to the API documentation.
- Update example config: extra properties, logging.
- Document a quick way to setup logging.
- Document the sent-at header in messages.
- Create a quick-start guide.
- Clarify queues are only deleted if unused.
- Wire-format: improve message properties documentation.
- Note the addition client properties in the config docs.

18.13.7 Contributors

Many thanks to the contributors of bug reports, pull requests, and pull request reviews for this release:

- Aurélien Bompard
- Adam Williamson
- dvejnz
- Jeremy Cline
- Randy Barlow
- Shraddha Agrawal
- Sebastian Wojciechowski

18.14 1.7.2 (2019-08-02)

18.14.1 Bug Fixes

- Fix variable substitution in log messages. (PR #200)
- Add MANIFEST.in and include tests for sample schema package. (PR #197)

18.14.2 Documentation Improvements

- Document the sent-at header in messages. (PR #199)
- Create a quick-start guide. (PR #196)

18.14.3 Contributors

Many thanks to the contributors of bug reports, pull requests, and pull request reviews for this release:

- Adam Williamson
- Aurélien Bompard
- Jeremy Cline
- Shraddha Agrawal

18.15 v1.7.1 (2019-06-24)

18.15.1 Bug Fixes

- Don't declare exchanges when consuming using the synchronous `fedora_messaging.api.consume()` API, which was causing consuming to fail from the Fedora broker (PR #191)

18.15.2 Contributors

Many thanks to the contributors of bug reports, pull requests, and pull request reviews for this release:

- Randy Barlow
- Aurélien Bompard
- Jeremy Cline
- Adam Williamson

18.15.3 Documentation Improvements

- Document some additional app properties and add a note about setting up logging in the `fedora.toml` and `stg.fedora.toml` configuration files (PR #188)
- Document how to setup logging in the consuming snippets so any problems are logged to stdout (PR #192)
- Document that logging is only set up for consumers (#181)
- Document the `fedora_messaging.config.conf` and `fedora_messaging.config.DEFAULTS` variables in the API documentation (#182)

18.16 v1.7.0 (2019-05-21)

18.16.1 Features

- “fedora-messaging consume” now accepts a “--callback-file” argument which will load a callback function from an arbitrary Python file. Previously, it was required that the callback be in the Python path (#159).

18.16.2 Bug Fixes

- Fix a bug where publishes that failed due to certain connection errors were not retried (#175).
- Fix a bug where AMQP protocol errors did not reset the connection used for publishing messages. This would result in publishes always failing with a `ConnectionError` (#178).

18.16.3 Documentation Improvements

- Document the `body` attribute on the `Message` class (#164).
- Clearly document what properties message schema classes should override (#166).
- Re-organize the consumer documentation to make the consuming API clearer (#168).

18.16.4 Contributors

Many thanks to the contributors of bug reports, pull requests, and pull request reviews for this release:

- Randy Barlow
- Aurélien Bompard
- Jeremy Cline
- Dusty Mabe

18.17 v1.6.1 (2019-04-17)

18.17.1 Bug Fixes

- Fix a bug in publishing where if the broker closed the connection, the client would not properly dispose of the connection object and publishing would fail forever (PR #157).
- Fix a bug in the `fedora_messaging.api.twisted_consume()` function where if the user did not have permissions to read from the specified queue which had already been declared, the `Deferred` that was returned never fired. It now errors back with a `fedora_messaging.exceptions.PermissionException` (PR #160).

18.17.2 Development Changes

- Stop pinning pytest to 4.0 or less as the incompatibility with pytest-twisted has been resolved (PR #158).

18.17.3 Other Changes

- Include commands to connect to the Fedora broker in the documentation (PR #154).

18.17.4 Contributors

Many thanks to the contributors of bug reports, pull requests, and pull request reviews for this release:

- Aurélien Bompard
- Jeremy Cline

18.18 v1.6.0 (2019-04-04)

18.18.1 Dependency Changes

- Twisted is no longer an optional dependency: fedora-messaging requires Twisted 12.2 or greater.

18.18.2 Features

- A new API, `fedora_messaging.api.twisted_consume()`, has been added to support consuming using the popular async framework Twisted. The fedora-messaging command-line interface has been switched to use this API. As a result, Twisted 12.2+ is now a dependency of fedora-messaging. Users of this new API are not affected by #130 (PR #139).

18.18.3 Bug Fixes

- Only prepend the `topic_prefix` on outgoing messages. Previously, the topic prefix was incorrectly applied to incoming messages (#143).

18.18.4 Documentation

- Add a note to the tutorial on how to instal the library and RabbitMQ in containers (PR #141).
- Document how to access the Fedora message broker from outside the Fedora infrastructure VPN. Users of fedmsg can now migrate to fedora-messaging for consumers outside Fedora's infrastructure. Consult the new documentation at `fedora-broker` for details (PR #149).

18.18.5 Contributors

Many thanks to the contributors of bug reports, pull requests, and pull request reviews for this release:

- Aurélien Bompard
- Jeremy Cline
- Shraddha Agrawal

18.19 v1.5.0 (2019-02-28)

18.19.1 Dependency Changes

- Replace the dependency on `pytoml` with `toml` (#132).

18.19.2 Features

- Support passive declarations for locked-down brokers (#136).

18.19.3 Bug Fixes

- Fix a bug in the sample schema package (#135).

18.19.4 Development Changes

- Switch to Mergify v2 (#129).

18.19.5 Contributors

Many thanks to the contributors of bug reports, pull requests, and pull request reviews for this release:

- Aurélien Bompard
- Jeremy Cline
- Michal Konečný
- Shraddha Agrawal

18.20 v1.4.0 (2019-02-07)

18.20.1 Features

- The `topic_prefix` configuration value has been added to automatically add a prefix to the topic of all outgoing messages. (#121)
- Support for Pika 0.13. (#126)
- Add a `systemd` service file for consumers.

18.20.2 Development Changes

- Use Bandit for security checking.

18.20.3 Contributors

Many thanks to the contributors of bug reports, pull requests, and pull request reviews for this release:

- Aurélien Bompard

18.21 v1.3.0 (2019-01-24)

18.21.1 API Changes

- The `Message._body` attribute is renamed to `body`, and is now part of the public API. (PR #119)

18.21.2 Contributors

Many thanks to the contributors of bug reports, pull requests, and pull request reviews for this release:

- Aurélien Bompard
- Jeremy Cline

18.22 v1.2.0 (2019-01-21)

18.22.1 Features

- The `fedora_messaging.api.consume()` API now accepts a “queues” keyword which specifies the queues to declare and consume from, and the “fedora-messaging” CLI makes use of this (PR #107)
- Utilities were added in the `schema_utils` module to help write the Python API of your message schemas (PR #108)
- No longer require “-exchange”, “-queue-name”, and “-routing-key” to all be specified when using “fedora-messaging consume”. If one is not supplied, a default is chosen. These defaults are documented in the command’s manual page (PR #117)

18.22.2 Bug Fixes

- Fix the “consumer” setting in `config.toml.example` to point to a real Python path (PR #104)
- `fedora-messaging consume` now actually uses the `-queue-name` and `-routing-key` parameter provided to it, and `-routing-key` can now be specified multiple times as was documented (PR #105)
- Fix the equality check on `fedora_messaging.message.Message` objects to exclude the ‘sent-at’ header (PR #109)
- Documentation for consumers indicated any callable object was acceptable to use as a callback as long as it accepted a single positional argument (the message). However, the implementation required that the callable be a function or a class, which it then instantiated. This has been fixed and you may now use any callable object, such as a method or an instance of a class that implements `__call__()` (PR #110)

- Fix an issue where the fedora-messaging CLI would only log if a configuration file was explicitly supplied (PR #113)

18.22.3 Contributors

Many thanks to the contributors of bug reports, pull requests, and pull request reviews for this release:

- Aurélien Bompard
- Jeremy Cline
- Sebastian Wojciechowski
- Tomas Tomecek

18.23 v1.1.0 (2018-11-13)

18.23.1 Features

- Initial work on a serialization format for `fedora_messaging.message.Message` and APIs for loading and storing messages. This is intended to make it easy to record and replay messages for testing purposes. (#84)
- Add a module, `fedora_messaging.testing`, to add useful test helpers. Check out the module documentation for details! (#100)

18.23.2 Contributors

Many thanks to the contributors of bug reports, pull requests, and pull request reviews for this release:

- Jeremy Cline
- Sebastian Wojciechowski

18.24 v1.0.1 (2018-10-10)

18.24.1 Bug Fixes

- Fix a compatibility issue in Twisted between pika 0.12 and 1.0. (#97)

18.25 v1.0.0 (2018-10-10)

18.25.1 API Changes

- The unused `exchange` parameter from the `PublisherSession` was removed (PR #56)
- The `setupRead` API in the Twisted protocol has been removed and replaced with `consume` and `cancel` APIs which allow for multiple consumers with multiple callbacks (PR #72)
- The name of the entry point is now used to identify the message type (PR #89)

18.25.2 Features

- Ensure proper TLS client cert checking with `service_identity` (PR #51)
- Support Python 3.7 (PR #53)
- Compatibility with `Click` 7.x (PR #86)
- The complete set of valid severity levels is now available at `fedora_messaging.api.SEVERITIES` (PR #60)
- A `queue` attribute is present on received messages with the name of the queue it arrived on (PR #65)
- The wire format of `fedora-messaging` is now documented (PR #88)

18.25.3 Development Changes

- Use `towncrier` to generate the release notes (PR #67)
- Check that our dependencies have Free licenses (PR #68)
- Test coverage is now at 97%.

18.25.4 Other Changes

- The library is available in Fedora as `fedora-messaging`.

18.25.5 Contributors

Many thanks to the contributors of bug reports, pull requests, and pull request reviews for this release:

- Aurélien Bompard
- Jeremy Cline
- Michal Konečný
- Sebastian Wojciechowski

18.26 v1.0.0b1

18.26.1 API Changes

- `fedora_messaging.message.Message.summary` is now a property rather than a method (#25).
- The non-functional `--amqp-url` parameter has been removed from the CLI (#49).

18.26.2 Features

- Configuration parsing failures now produce point to the line and column of the parsing error (#21).
- `fedora_messaging.message.Message` now come with a set of standard accessors (#32).
- Consumers can now specify whether a message should be re-queued when halting (#44).
- An example consumer that prints to standard output now ships with `fedora-messaging`. It can be used by running `fedora-messaging consume --callback="fedora_messaging.example:printer"` (#40).
- `fedora_messaging.message.Message` now have a `severity` associated with them (#48).

18.26.3 Bug Fixes

- Fix an issue where invalid or missing configuration files resulted in a traceback rather than a formatted error message from the CLI (#21).
- Client authentication with x509 now works with both the synchronous API and the Twisted API (#29, #35).
- `fedora_messaging.api.publish()` no longer raises a `pika.exceptions.ChannelClosed` exception. Instead, it raises a `fedora_messaging.exceptions.ConnectionException` (#31).
- `fedora_messaging.api.consume()` is now documented to raise a `ValueError` when the callback isn't callable (#47).

18.26.4 Development Features

- The `fedora-messaging` code base is now compliant with the `Black` Python formatter and this is enforced with continuous integration.
- Test coverage is moving up and to the right.

Many thanks to the contributors of bug reports, pull requests, and pull request reviews for this release:

- Aurélien Bompard
- Clement Verna
- Ken Dreyer
- Jeremy Cline
- Miroslav Suchý
- Patrick Uiterwijk
- Sebastian Wojciechowski

18.27 v1.0.0a1

The initial alpha release for `fedora-messaging` v1.0.0. The API is not expected to change significantly between this release and the final v1.0.0 release, but it may do so if serious flaws are discovered in it.

PYTHON MODULE INDEX

f

- `fedora_messaging.config`, [11](#)
- `fedora_messaging.exceptions`, [84](#)
- `fedora_messaging.message`, [76](#)
- `fedora_messaging.schema_utils`, [83](#)
- `fedora_messaging.signals`, [75](#)
- `fedora_messaging.testing`, [43](#)

Symbols

`__str__()` (*fedora_messaging.message.Message* method), 78

A

`agent_avatar` (*fedora_messaging.message.Message* property), 78

`agent_name` (*fedora_messaging.message.Message* property), 78

`app_icon` (*fedora_messaging.message.Message* property), 79

`app_name` (*fedora_messaging.message.Message* property), 79

B

`BadDeclaration`, 84

`BaseException`, 84

`body` (*fedora_messaging.message.Message* attribute), 77

`body_schema` (*fedora_messaging.message.Message* attribute), 77

C

`callback` (*fedora_messaging.api.Consumer* attribute), 73

`cancel()` (*fedora_messaging.api.Consumer* method), 74

`conf` (in module *fedora_messaging.config*), 86

`ConfigurationException`, 84

`ConnectionException`, 84

`consume()` (in module *fedora_messaging.api*), 74

`ConsumeException`, 84

`Consumer` (class in *fedora_messaging.api*), 73

`ConsumerCanceled`, 84

`containers` (*fedora_messaging.message.Message* property), 79

D

`DEBUG` (in module *fedora_messaging.message*), 82

`DEFAULTS` (in module *fedora_messaging.config*), 86

`deprecated` (*fedora_messaging.message.Message* attribute), 78

`Drop`, 84

`dumps()` (in module *fedora_messaging.message*), 82

E

`ERROR` (in module *fedora_messaging.message*), 82

F

`fedora_messaging.config`
module, 11

`fedora_messaging.exceptions`
module, 84

`fedora_messaging.message`
module, 76

`fedora_messaging.schema_utils`
module, 83

`fedora_messaging.signals`
module, 75

`fedora_messaging.testing`
module, 43

`flatpaks` (*fedora_messaging.message.Message* property), 79

G

`groups` (*fedora_messaging.message.Message* property), 80

H

`HaltConsumer`, 84

`headers_schema` (*fedora_messaging.message.Message* attribute), 77

I

`id` (*fedora_messaging.message.Message* attribute), 77

`INFO` (in module *fedora_messaging.message*), 82

L

`libravatar_url()` (in module *fedora_messaging.schema_utils*), 83

`loads()` (in module *fedora_messaging.message*), 83

M

`Message` (class in *fedora_messaging.message*), 76

`mock_sends()` (in module *fedora_messaging.testing*), 43
 module

- fedora_messaging.config*, 11
- fedora_messaging.exceptions*, 84
- fedora_messaging.message*, 76
- fedora_messaging.schema_utils*, 83
- fedora_messaging.signals*, 75
- fedora_messaging.testing*, 43

modules (*fedora_messaging.message.Message* property), 80

N

Nack, 85

NoFreeChannels, 85

P

packages (*fedora_messaging.message.Message* property), 80

PermissionException, 85

pre_publish_signal (in module *fedora_messaging.api*), 75

priority (*fedora_messaging.message.Message* attribute), 78

publish() (in module *fedora_messaging.api*), 72

publish_failed_signal (in module *fedora_messaging.api*), 76

publish_signal (in module *fedora_messaging.api*), 75

PublishException, 85

PublishForbidden, 85

PublishReturned, 85

PublishTimeout, 85

Q

queue (*fedora_messaging.api.Consumer* attribute), 73

queue (*fedora_messaging.message.Message* attribute), 77

R

result (*fedora_messaging.api.Consumer* attribute), 74

S

SERIALIZED_MESSAGE_SCHEMA (in module *fedora_messaging.message*), 83

SEVERITIES (in module *fedora_messaging.message*), 82

severity (*fedora_messaging.message.Message* attribute), 77

summary (*fedora_messaging.exceptions.ValidationError* property), 85

summary (*fedora_messaging.message.Message* property), 80

T

topic (*fedora_messaging.message.Message* attribute), 77

twisted_consume() (in module *fedora_messaging.api*), 73

U

url (*fedora_messaging.message.Message* property), 80

usernames (*fedora_messaging.message.Message* property), 81

V

validate() (*fedora_messaging.message.Message* method), 81

ValidationError, 85

W

WARNING (in module *fedora_messaging.message*), 82