

1 "Hello World!"


The simplest thing that does *something*



[Python](#) | [Java](#)

2 Work queues


Distributing tasks among workers



[Python](#) | [Java](#)

3 Publish/Subscribe

Sending messages to many consumers at once



[Python](#) | [Java](#)

4 Routing


Receiving messages selectively



[Python](#) | [Java](#)

5 Topics

Receiving messages based on a pattern



[Python](#) | [Java](#)

Publish/Subscribe

(using the [pika 0.9.8 Python client](#))

In the **previous tutorial** we created a work queue. The assumption behind a work queue is that each task is delivered to exactly one worker. In this part we'll do something completely different -- we'll deliver a message to multiple consumers. This pattern is known as "publish/subscribe".

Where to get help

If you're having trouble going through this tutorial you can **contact us** through the discussion list or directly.

To illustrate the pattern, we're going to build a simple logging system. It will consist of two programs -- the first will emit log messages and the second will receive and print them.

In our logging system every running copy of the receiver program will get the messages. That way we'll be able to run one receiver and direct the logs to disk; and at the same time we'll be able to run another receiver and see the logs on the screen.

Essentially, published log messages are going to be broadcast to all the receivers.

Exchanges

In previous parts of the tutorial we sent and received messages to and from a queue. Now it's time to introduce the full messaging model in Rabbit.

Let's quickly go over what we covered in the previous tutorials:


A *producer* is a user application that sends messages.

A *queue* is a buffer that stores messages.

A *consumer* is a user application that receives messages.

The core idea in the messaging model in RabbitMQ is that the producer never sends any messages directly to a queue. Actually, quite often the producer doesn't even know if a message will be delivered to any queue at all.

Instead, the producer can only send messages to an *exchange*. An exchange is a very simple thing. On one side it receives messages from producers and the other side it pushes them to queues. The exchange must know exactly what to do with a message it receives. Should it be appended to a particular queue? Should it be appended to many queues? Or should it get discarded. The rules for that are defined by the *exchange type*.



There are a few exchange types available: `direct`, `topic`, `headers` and `fanout`. We'll focus on the last one -- the fanout. Let's create an exchange of that type, and call it `logs`:

```
channel.exchange_declare(exchange='logs',
                        type='fanout')
```

The fanout exchange is very simple. As you can probably guess from the name, it just broadcasts all the messages it receives to all the queues it knows. And that's exactly what we need for our logger.

Listing exchanges

To list the exchanges on the server you can run the ever useful `rabbitmqctl`:

```
$ sudo rabbitmqctl list_exchanges
Listing exchanges ...
logs      fanout
amq.direct  direct
amq.topic   topic
amq.fanout  fanout
```

6 RPC

Remote procedure call implementation



Python | Java

```
amq.headers    headers
...done.
```

In this list there are some `amq.*` exchanges and the default (unnamed) exchange. These are created by default, but it is unlikely you'll need to use them at the moment.

Nameless exchange

In previous parts of the tutorial we knew nothing about exchanges, but still were able to send messages to queues. That was possible because we were using a default exchange, which we identify by the empty string (`""`).

Recall how we published a message before:

```
channel.basic_publish(exchange='',
                      routing_key='hello',
                      body=message)
```

The `exchange` parameter is the the name of the exchange. The empty string denotes the default or *nameless* exchange: messages are routed to the queue with the name specified by `routing_key`, if it exists.

Now, we can publish to our named exchange instead:

```
channel.basic_publish(exchange='logs',
                      routing_key='',
                      body=message)
```

Temporary queues

As you may remember previously we were using queues which had a specified name (remember `hello` and `task_queue`?). Being able to name a queue was crucial for us -- we needed to point the workers to the same queue. Giving a queue a name is important when you want to share the queue between producers and consumers.

But that's not the case for our logger. We want to hear about all log messages, not just a subset of them. We're also interested only in currently flowing messages not in the old ones. To solve that we need two things.

Firstly, whenever we connect to Rabbit we need a fresh, empty queue. To do it we could create a queue with a random name, or, even better - let the server choose a random queue name for us. We can do this by not supplying the `queue` parameter to `queue_declare`:


```
result = channel.queue_declare()
```

At this point `result.method.queue` contains a random queue name. For example it may look like `amq.gen-JzTY20BRgKO-HjmUJj0wLg`.

Secondly, once we disconnect the consumer the queue should be deleted. There's an `exclusive` flag for that:

```
result = channel.queue_declare(exclusive=True)
```

Bindings



We've already created a fanout exchange and a queue. Now we need to tell the exchange to send messages to our queue. That relationship between exchange and a queue is called a *binding*.


```
channel.queue_bind(exchange='logs',
                  queue=result.method.queue)
```

From now on the `logs` exchange will append messages to our queue.

Listing bindings

You can list existing bindings using, you guessed it, `rabbitmqctl list_bindings`.

Putting it all together



The producer program, which emits log messages, doesn't look much different from the previous tutorial. The most important change is that we now want to publish messages to our `logs` exchange instead of the nameless one. We need to supply a `routing_key` when sending, but its value is ignored for `fanout` exchanges. Here goes the code for `emit_log.py` script:

```
1  #!/usr/bin/env python
2  import pika
3  import sys
4
5  connection = pika.BlockingConnection(pika.ConnectionParameters(
6      host='localhost'))
7  channel = connection.channel()
8
9  channel.exchange_declare(exchange='logs',
10                          type='fanout')
11
12 message = ' '.join(sys.argv[1:]) or "info: Hello World!"
13 channel.basic_publish(exchange='logs',
14                      routing_key='',
15                      body=message)
16 print " [x] Sent %r" % (message,)
17 connection.close()
```

(emit_log.py source)

As you see, after establishing the connection we declared the exchange. This step is necessary as publishing to a non-existing exchange is forbidden.

The messages will be lost if no queue is bound to the exchange yet, but that's okay for us; if no consumer is listening yet we can safely discard the message.

The code for `receive_logs.py`:

```
1  #!/usr/bin/env python
2  import pika
3
4  connection = pika.BlockingConnection(pika.ConnectionParameters(
5      host='localhost'))
6  channel = connection.channel()
7
8  channel.exchange_declare(exchange='logs',
9                          type='fanout')
10
11 result = channel.queue_declare(exclusive=True)
12 queue_name = result.method.queue
13
14 channel.queue_bind(exchange='logs',
15                  queue=queue_name)
16
17 print ' [*] Waiting for logs. To exit press CTRL+C'
18
19 def callback(ch, method, properties, body):
20     print " [x] %r" % (body,)
21
22 channel.basic_consume(callback,
23                      queue=queue_name,
24                      no_ack=True)
25
26 channel.start_consuming()
```

(receive_logs.py source)

We're done. If you want to save logs to a file, just open a console and type:

```
$ python receive_logs.py > logs_from_rabbit.log
```

If you wish to see the logs on your screen, spawn a new terminal and run:

```
$ python receive_logs.py
```

And of course, to emit logs type:

```
$ python emit_log.py
```

Using `rabbitmqctl list_bindings` you can verify that the code actually creates bindings and queues as we want. With two `receive_logs.py` programs running you should see something like:

```
$ sudo rabbitmqctl list_bindings
Listing bindings ...
logs    exchange      amq.gen-JzTY20BRgK0-HjmUJj0wLg  queue    []
logs    exchange      amq.gen-vso0PVvyiRIL2WoV3i48Yg  queue    []
...done.
```

The interpretation of the result is straightforward: data from exchange `logs` goes to two queues with server-assigned names. And that's exactly what we intended.

To find out how to listen for a subset of messages, let's move on to **tutorial 4**