HumusAmqp Documentation

Release 1.0.0

Sascha-Oliver Prolic

Nov 24, 2018

Contents

Documentation for HumusAMQP

CHAPTER 1

Overview

PHP 7 AMQP libray supporting multiple drivers and providing full-featured Consumer, Producer, and JSON-RPC Client / Server implementations.

The JSON-RPC part implements JSON-RPC 2.0 Specification.

Current supported drivers are: php-amqp and PhpAmqpLib.

This library ships with *container-interop* factories that help you setting up everything.

CHAPTER 2

Contents

2.1 Guides overview

We recommend that you read these guides, if possible, in this order:

2.1.1 Getting started

An overview of HumusAmqp with a quick tutorial that helps you to get started with it. It should take about 30 minutes to read and study the provided code examples.

2.1.2 AMQP 0.9.1 Model Concepts

This guide covers:

- AMQP 0.9.1 model overview
- What are channels
- What are vhosts
- What are queues
- What are exchanges
- What are bindings
- What are AMQP 0.9.1 classes and methods

2.1.3 Connecting to RabbitMQ with HumusAmqp

This guide covers:

- How to connect to RabbitMQ with HumusAmqp

- How to use connection URI to connect to RabbitMQ (also: in PaaS environments such as Heroku and CloudFoundry)
- How to open a channel
- How to close a channel
- How to disconnect

2.1.4 Exchanges and Publishing

This guide covers:

- Exchange types
- How to declare AMQP exchanges with HumusAmqp
- How to publish messages
- Exchange properties
- Fanout exchanges
- Direct exchanges
- Topic exchanges
- Default exchange
- Message and delivery properties
- Message routing
- Bindings
- How to delete exchanges
- Other topics related to exchanges and publishing

2.1.5 HumusAmqp Producer's

This guide covers:

- Producer types
- Custom producers types
- How to easily handle publishing messages with HumusAmqp

2.1.6 Queues

This guide covers:

- How to declare AMQP queues with HumusAmqp
- Queue properties
- How to declare server-named queues
- How to declare temporary exclusive queues
- How to consume messages ("push API")
- How to fetch messages ("pull API")

- Message and delivery properties
- Message acknowledgements
- How to purge queues
- How to delete queues
- Other topics related to queues

2.1.7 Bindings

This guide covers:

- How to bind exchanges to queues
- How to unbind exchanges from queues
- Other topics related to bindings

2.1.8 Consumers

This guide covers:

- How to declare AMQP queues with HumusAmqp
- Queue properties
- How to declare server-named queues
- How to declare temporary exclusive queues
- How to consume messages (“push API”)
- How to fetch messages (“pull API”)
- Message and delivery properties
- Message acknowledgements
- How to purge queues
- How to delete queues
- Other topics related to queues

2.1.9 Durability and Related Matters

This guide covers:

- Topics related to durability of exchanges and queues
- Durability of messages

JSON RPC Server & Client

This guide covers:

- How to set up a JSON RPC Server
- How to set up a JSON RPC Client

2.1.10 RabbitMQ Extensions to AMQP 0.9.1

This guide covers RabbitMQ extensions and how they are used in HumusAmqp:

- How to use exchange-to-exchange bindings
- How to the alternate exchange extension
- How to set per-queue message TTL
- How to set per-message TTL
- What are consumer cancellation notifications and how to use them
- Message *dead lettering* and the dead letter exchange

2.1.11 Error Handling and Recovery

This guide covers:

- AMQP 0.9.1 protocol exceptions
- How to deal with network failures
- Other things that may go wrong

2.1.12 Troubleshooting

This guide covers:

- What to check when your apps that use HumusAmqp and RabbitMQ misbehave

2.1.13 Deployment

This guide covers:

- What to check when your apps that use HumusAmqp and RabbitMQ misbehave

Tell Us What You Think!

Please take a moment to tell us what you think about this guide: Send an e-mail, say hello in the HumusAmqp gitter chat. or raise an issue on Github.

Let us know what was unclear or what has not been covered. Maybe you do not like the guide style or grammar or discover spelling mistakes. Reader feedback is key to making the documentation better.

2.2 Getting Started with HumusAmqp and RabbitMQ

2.2.1 About this guide

This guide is a quick tutorial that helps you to get started with RabbitMQ and HumusAmqp. It should take about 20 minutes to read and study the provided code examples. This guide covers:

- Installing RabbitMQ, a mature popular messaging broker server.
- Installing HumusAmqp via Composer.

- Producing and consuming messages from cli.
- Running a JSON-RPC server and client.

2.2.2 Installing RabbitMQ

The RabbitMQ site has a good installation guide that addresses many operating systems. On Mac OS X, the fastest way to install RabbitMQ is with Homebrew:

```
$ brew install rabbitmq
```

then run it:

```
$ rabbitmq-server
```

On Debian and Ubuntu, you can either download the RabbitMQ .deb package and install it with dpkg or make use of the apt repository that the RabbitMQ team provides.

For RPM-based distributions like RedHat or CentOS, the RabbitMQ team provides an RPM package.

2.2.3 Installing HumusAmqp

1. You can use composer to install HumusAmqp

```
$ php composer.phar require prolic/humus-amqp ^1.0
```

2. Select a driver to use, currently only php-amqplib and php-amqp (PHP extension) are supported.

Configure your container to return the DriverFactory for the Driver class. If your container supports configuration by array like ZendServiceManager f.e., this should look similar to this. For integration with the Symfony Framework you can use the Symfony container-interop bundle.

```
return [
    'dependencies' => [
        'factories' => [
            Driver::class => Container\DriverFactory::class,
        ]
    ]
];
```

3. Configure your application to use the desired driver.

```
return [
    'humus' => [
        'amqp' => [
            'driver' => 'amqp-extension',
        ]
    ]
];
```

4. Notes about drivers:

1) The PHP Extension (php-amqp) is the fastest one, we strongly recommend using 1.7.1 or building yourself from master to be able to use all features.

There is currently a bug in PhpAmqpLib, see this pull-request. As long as this is not merged and release, you have to manually apply the patch, sorry!

You can do this from the command-line with:

sed -i '/$message = $this->get_and_unset_message($delivery_tag);/a $message->delivery_info["delivery_tag"] = $delivery_tag;' vendor/php-amqplib/php-amqplib/PhpAmqpLib/Channel/AMQPChannel.php

2) When using php-amqplib as driver, it's worth pointing out, that a StreamConnection (same goes for SSLConnection) does not have the possibility to timeout. If you want to let the consumer timeout, when no more messages are received, you should use the SocketConnection instead (assuming you don't need an SSL connection).

2.2.4 Sample-Configuration

A sample configuration might look like this, more details an explanation will be in the coming chapters.

```
return [
    'dependencies' => [
        'factories' => [
            Driver::class => Container\DriverFactory::class,
            'default-amqp-connection' => [Container\ConnectionFactory::class,
↪'default-amqp-connection'],
            'demo-producer' => [Container\ProducerFactory::class, 'demo-producer'],
            'topic-producer' => [Container\ProducerFactory::class, 'topic-producer'],
            'demo-consumer' => [Container\CallbackConsumerFactory::class, 'demo-
↪consumer'],
            'topic-consumer-error' => [Container\CallbackConsumerFactory::class,
↪'topic-consumer-error'],
            'demo-rpc-server' => [Container\JsonRpcServerFactory::class, 'demo-rpc-
↪server'],
            'demo-rpc-server2' => [Container\JsonRpcServerFactory::class, 'demo-rpc-
↪server2'],
            'demo-rpc-client' => [Container\JsonRpcClientFactory::class, 'demo-rpc-
↪client'],
            'my_callback' => $my_callback_factory,
        ],
    ],
    'humus' => [
        'amqp' => [
            'driver' => 'amqp-extension',
            'exchange' => [
                'demo' => [
                    'name' => 'demo',
                    'type' => 'direct',
                    'connection' => 'default-amqp-connection',
                ],
                'demo.error' => [
                    'name' => 'demo.error',
                    'type' => 'direct',
                    'connection' => 'default-amqp-connection',
                ],
                'topic-exchange' => [
                    'name' => 'topic-exchange',
                    'type' => 'topic',
                    'connection' => 'default-amqp-connection',
                ],
                'demo-rpc-client' => [
                    'name' => 'demo-rpc-client',
                    'type' => 'direct',
                    'connection' => 'default-amqp-connection',
```

(continues on next page)

(continued from previous page)

```
                ],
                'demo-rpc-server' => [
                    'name' => 'demo-rpc-server',
                    'type' => 'direct',
                    'connection' => 'default-amqp-connection',
                ],
                'demo-rpc-server2' => [
                    'name' => 'demo-rpc-server2',
                    'type' => 'direct',
                    'connection' => 'default-amqp-connection',
                ],
            ],
            'queue' => [
                'foo' => [
                    'name' => 'foo',
                    'exchanges' => [
                        'demo' => [
                            [
                                'arguments' => [
                                    'x-dead-letter-exchange' => 'demo.error', // must␣
↪be defined as exchange before
                                ],
                            ],
                        ],
                    ],
                    'connection' => 'default-amqp-connection',
                ],
                'demo-rpc-client' => [
                    'name' => '',
                    'exchanges' => [
                        'demo-rpc-client' => [],
                    ],
                    'connection' => 'default-amqp-connection',
                ],
                'demo-rpc-server' => [
                    'name' => 'demo-rpc-server',
                    'exchanges' => [
                        'demo-rpc-server' => [],
                    ],
                    'connection' => 'default-amqp-connection',
                ],
                'demo-rpc-server2' => [
                    'name' => 'demo-rpc-server2',
                    'exchanges' => [
                        'demo-rpc-server2' => [],
                    ],
                    'connection' => 'default-amqp-connection',
                ],
                'info-queue' => [
                    'name' => 'info-queue',
                    'exchanges' => [
                        'topic-exchange' => [
                            [
                                'routing_keys' => [
                                    '#.err',
                                ],
                            ]
```

(continues on next page)

(continued from previous page)

```
                    ],
                ],
                'connection' => 'default-amqp-connection',
            ],
        ],
        'connection' => [
            'default-amqp-connection' => [
                'host' => 'localhost',
                'port' => 5672,
                'login' => 'guest',
                'password' => 'guest',
                'vhost' => '/',
                'persistent' => true,
                'read_timeout' => 3, //sec, float allowed
                'write_timeout' => 1, //sec, float allowed
            ],
        ],
        'producer' => [
            'demo-producer' => [
                'type' => 'plain',
                'exchange' => 'demo',
                'qos' => [
                    'prefetch_size' => 0,
                    'prefetch_count' => 10,
                ],
                'auto_setup_fabric' => true,
            ],
            'topic-producer' => [
                'exchange' => 'topic-exchange',
                'auto_setup_fabric' => true,
            ],
        ],
        'callback_consumer' => [
            'demo-consumer' => [
                'queue' => 'foo',
                'callback' => 'echo',
                'idle_timeout' => 10,
                'delivery_callback' => 'my_callback',
            ],
            'topic-consumer-error' => [
                'queue' => 'info-queue',
                'qos' => [
                    'prefetch_count' => 100,
                ],
                'auto_setup_fabric' => true,
                'callback' => 'echo',
                'logger' => 'consumer-logger',
            ],
        ],
        'json_rpc_server' => [
            'demo-rpc-server' => [
                'callback' => 'poweroftwo',
                'queue' => 'demo-rpc-server',
                'auto_setup_fabric' => true,
            ],
            'demo-rpc-server2' => [
                'callback' => 'randomint',
```

(continues on next page)

(continued from previous page)

```
                    'queue' => 'demo-rpc-server2',
                    'auto_setup_fabric' => true,
                ],
            ],
            'json_rpc_client' => [
                'demo-rpc-client' => [
                    'queue' => 'demo-rpc-client',
                    'auto_setup_fabric' => true,
                ],
            ],
        ],
    ],
];
```

2.2.5 What to Read Next

The documentation is organized as *a number of guides*, covering various topics.

We recommend that you read the following guides first, if possible, in this order:

- *Connecting to RabbitMQ with HumusAmqp*
- *Exchanges and Publishing*
- *HumusAmqp Producer's*
- *Queues and Consumers*
- *Bindings*
- *Consumers*
- *CLI*
- *Durability and Related Matters*
- *JSON RPC Server & Client*
- *RabbitMQ Extensions to AMQP 0.9.1*
- *Error Handling and Recovery*
- *Troubleshooting*
- *Deployment*

2.2.6 Tell Us What You Think!

Please take a moment to tell us what you think about this guide: Send an e-mail, say hello in the HumusAmqp gitter chat. or raise an issue on Github.

Let us know what was unclear or what has not been covered. Maybe you do not like the guide style or grammar or discover spelling mistakes. Reader feedback is key to making the documentation better.

2.3 Connecting to RabbitMQ with HumusAmqp

2.3.1 Connection configuration

Map options that HumusAmqp will recognize are

- `:host` - amqp.host The host to connect too. Note: Max 1024 characters.
- `:port` - amqp.port Port on the host.
- `:vhost` - amqp.vhost The virtual host on the host. Note: Max 128 characters.
- `:login` - amqp.login The login name to use. Note: Max 128 characters.
- `:password` - amqp.password Password. Note: Max 128 characters.
- `:persistent` - Establish a persistent connection with the AMQP broker, if set to true.
- `:connect_timeout` - Connection timeout. Note: 0 or greater seconds. May be fractional.
- `:read_timeout` - Timeout in for income activity. Note: 0 or greater seconds. May be fractional.
- `:write_timeout` - Timeout in for outcome activity. Note: 0 or greater seconds. May be fractional.
- `:heartbeat` - The heartbeat to use. Should be approx half the read_timeout.
- `:cacert` - CA Cert for SSL Connections
- `:cert` - Cert for SSL Connections
- `:key` - Key for SSL Connections
- `:verify` - Verify SSL Certs (true or false)

2.3.2 Default parameters

Default connection parameters are

```
[
  'host'         => "localhost",
  'port'         => 5672,
  'vhost'        => "/",
  'login'        => "guest",
  'password'     => "guest",
  'persistent'   => false,
  'connect_timeout' => 1.0,
  'read_timeout'  => 1.0,
  'write_timeout' => 1.0,
  'heartbeat' => 0,
]
```

2.3.3 Creating a connection

```
$options = new Humus\Amqp\ConnectionOptions();
$options->setLogin('username');
$options->setPassword('password');

$connection = new Humus\Amqp\Driver\AmqpExtension\Connection($options);
$connection->connect();
```

2.3.4 Opening a Channel

Some applications need multiple connections to RabbitMQ. However, it is undesirable to keep many TCP connections open at the same time because doing so consumes system resources and makes it more difficult to configure firewalls. AMQP 0-9-1 connections are multiplexed with channels that can be thought of as "lightweight connections that share a single TCP connection".

To open a channel:

```
<?php

$connection = new Humus\Amqp\Driver\PhpAmqpLib\SocketConnection(new␣
↪Humus\Amqp\ConnectionOptions());
$channel    = $connection->newChannel();
```

Channels are typically long lived: you open one or more of them and use them for a period of time, as opposed to opening a new channel for each published message, for example.

2.3.5 Using configuration and factory

You can simply configure as many connection as needed and simply give them a name.

```
<?php

return [
    'dependencies' => [
        'factories' => [
            Driver::class => Humus\Amqp\Container\DriverFactory::class,
            'default-amqp-connection' =>␣
↪[Humus\Amqp\Container\ConnectionFactory::class, 'default'],
        ],
    ],
    'humus' => [
        'amqp' => [
            'driver' => 'php-amqplib',
            'connection' => [
                'default' => [
                    'type' => 'socket',
                    'host' => 'localhost',
                    'port' => 5672,
                    'login' => 'guest',
                    'password' => 'guest',
                    'vhost' => '/',
                    'persistent' => true,
                    'read_timeout' => 3, //sec, float allowed
                    'write_timeout' => 1, //sec, float allowed
                ],
            ],
        ]
    ]
);
```

Note: When using php-amqplib as driver, you have five different connection classes available: *LazyConnection*, *LazySocketConnection*, *SocketConnection*, *SslConnection* and *StreamConnection*. Which connection type should be used, can be configured with the *type* option and the settings: *lazy*, *lazy_socket*, *socket*, *ssl* or *stream*.

2.3.6 Getting a connection

```
<?php

$defaultConnection = $container->get('default-amqp-connection');
```

2.3.7 Troubleshooting

If you have read this guide and still have issues with connecting, check our *Troubleshooting guide* and feel free to raise an issue at Github.

2.3.8 What to Read Next

The documentation is organized as *a number of guides*, covering various topics.

We recommend that you read the following guides first, if possible, in this order:

- *Exchanges and Publishing*
- *HumusAmqp Producer's*
- *Queues and Consumers*
- *Bindings*
- *Consumers*
- *CLI*
- *Durability and Related Matters*
- *JSON RPC Server & Client*
- *RabbitMQ Extensions to AMQP 0.9.1*
- *Error Handling and Recovery*
- *Troubleshooting*
- *Deployment*

2.3.9 Tell Us What You Think!

Please take a moment to tell us what you think about this guide: Send an e-mail, say hello in the HumusAmqp gitter chat. or raise an issue on Github.

Let us know what was unclear or what has not been covered. Maybe you do not like the guide style or grammar or discover spelling mistakes. Reader feedback is key to making the documentation better.

2.4 Exchanges

2.4.1 Exchanges in AMQP 0.9.1 — Overview

What are AMQP exchanges?

An *exchange* accepts messages from a producer application and routes them to message queues. They can be thought of as the "mailboxes" of the AMQP world. Unlike some other messaging middleware products and protocols, in AMQP, messages are *not* published directly to queues. Messages are published to exchanges that route them to queue(s) using pre-arranged criteria called *bindings*.

There are multiple exchange types in the AMQP 0.9.1 specification, each with its own routing semantics. Custom exchange types can be created to deal with sophisticated routing scenarios (e.g. routing based on geolocation data or edge cases) or just for convenience.

Concept of Bindings

A *binding* is an association between a queue and an exchange. A queue must be bound to at least one exchange in order to receive messages from publishers. Learn more about bindings in the *Bindings Guide*.

Exchange attributes

Exchanges have several attributes associated with them:

- Name
- Type (direct, fanout, topic, headers or some custom type)
- Durability
- Whether the exchange is auto-deleted when no longer used
- Other metadata (sometimes known as *X-arguments*)

2.4.2 Exchange types

There are four built-in exchange types in AMQP v0.9.1:

- Direct
- Fanout
- Topic
- Headers

As stated previously, each exchange type has its own routing semantics and new exchange types can be added by extending brokers with plugins. Custom exchange types begin with "x-", much like custom HTTP headers, e.g. x-consistent-hash exchange or x-random exchange.

2.4.3 Message attributes

Before we start looking at various exchange types and their routing semantics, we need to introduce message attributes. Every AMQP message has a number of *attributes*. Some attributes are important and used very often, others are rarely used. AMQP message attributes are metadata and are similar in purpose to HTTP request and response headers.

Every AMQP 0.9.1 message has an attribute called *routing key*. The routing key is an "address" that the exchange may use to decide how to route the message. This is similar to, but more generic than, a URL in HTTP. Most exchange types use the routing key to implement routing logic, but some ignore it and use other criteria (e.g. message content).

2.4.4 Declaring an exchange

```
$options = new Humus\Amqp\ConnectionOptions();
$options->setLogin('username');
$options->setPassword('password');

$connection = new Humus\Amqp\Driver\AmqpExtension\Connection($options);
$connection->connect();

$channel = $connection->newChannel();

$exchange = $channel->newExchange();
$exchange->setName('my-exchange');
$exchange->setType('direct');
$exchange->declareExchange();
```

2.4.5 Using configuration and factory

```
<?php

return [
    'dependencies' => [
        'factories' => [
            Driver::class => Humus\Amqp\Container\DriverFactory::class,
            'default-amqp-connection' =>␣
↪[Humus\Amqp\Container\ConnectionFactory::class, 'default'],
        ],
    ],
    'humus' => [
        'amqp' => [
            'driver' => 'php-amqplib',
            'connection' => [
                'default' => [
                    'type' => 'socket',
                    'host' => 'localhost',
                    'port' => 5672,
                    'login' => 'guest',
                    'password' => 'guest',
                    'vhost' => '/',
                    'persistent' => false,
                    'read_timeout' => 3, //sec, float allowed
                    'write_timeout' => 1, //sec, float allowed
                ],
            ],
            'exchange' => [
                'my-exchange' => [
                    'name' => 'my-exchange',
                    'type' => 'direct',
                    'connection' => 'default-amqp-connection',
                    'auto_setup_fabric' => true,
```

(continues on next page)

(continued from previous page)

```
                    ],
                ]
            ]
        ]
);
```

When *auto_setup_fabric* is set to true, the exchange factory will automatically declare the configured exchanged. For durable exchanges and queues it's recommended for production to disable it, and declare all needed exchanges and queues upfront.

2.4.6 Fanout exchanges

How fanout exchanges route messages

A fanout exchange routes messages to all of the queues that are bound to it and the routing key is ignored. If N queues are bound to a fanout exchange, when a new message is published to that exchange a *copy of the message* is delivered to all N queues. Fanout exchanges are ideal for the broadcast routing of messages.

Graphically this can be represented as:

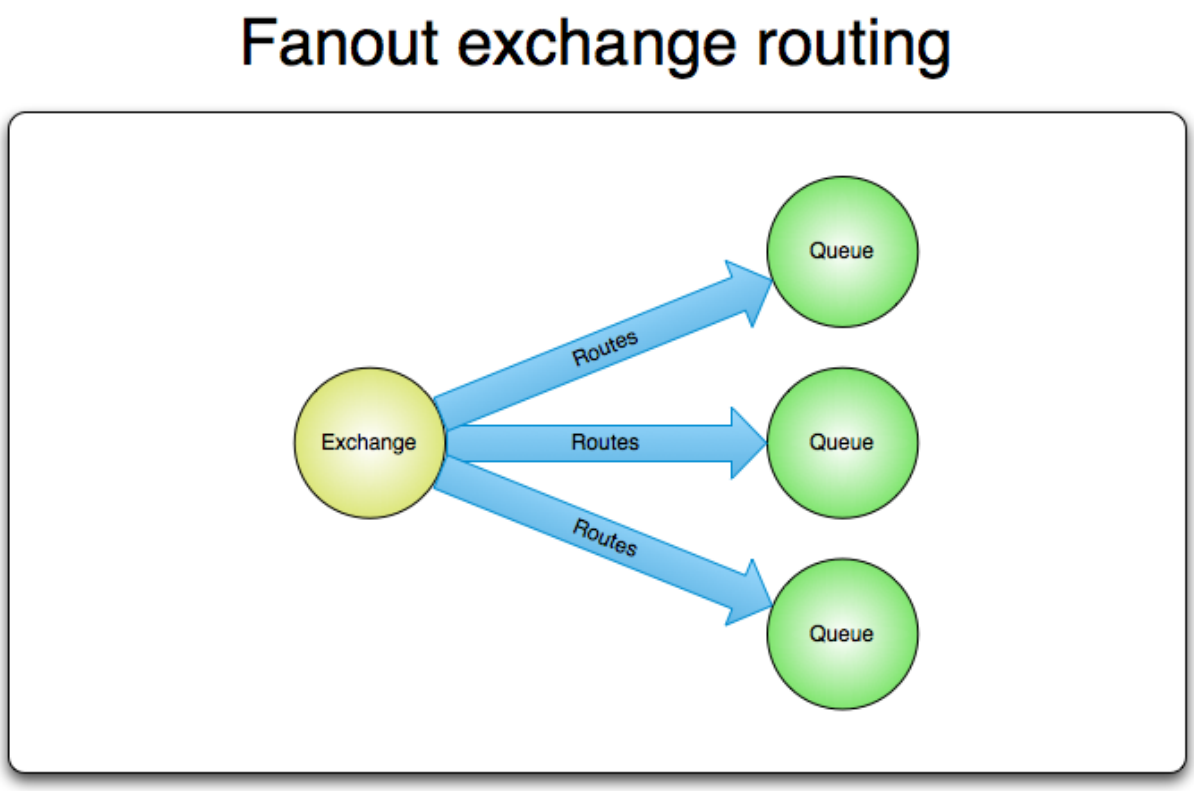

Fig. 1: fanout exchange routing

Fanout use cases

Because a fanout exchange delivers a copy of a message to every queue bound to it, its use cases are quite similar:

- Massively multiplayer online (MMO) games can use it for leaderboard updates or other global events
- Sport news sites can use fanout exchanges for distributing score updates to mobile clients in near real-time
- Distributed systems can broadcast various state and configuration updates
- Group chats can distribute messages between participants using a fanout exchange (although AMQP does not have a built-in concept of presence, so XMPP may be a better choice)

Pre-declared fanout exchanges

AMQP 0.9.1 brokers must implement a fanout exchange type and pre-declare one instance with the name of "`amq.fanout`".

Applications can rely on that exchange always being available to them. Each vhost has a separate instance of that exchange, it is *not shared across vhosts* for obvious reasons.

2.4.7 Direct exchanges

How direct exchanges route messages

A direct exchange delivers messages to queues based on a *message routing key*, an attribute that every AMQP v0.9.1 message contains.

Here is how it works:

- A queue binds to the exchange with a routing key K
- When a new message with routing key R arrives at the direct exchange, the exchange routes it to the queue if K = R

A direct exchange is ideal for the unicast routing of messages (although they can be used for multicast routing as well).

Here is a graphical representation:

Direct routing example

Since direct exchanges use the *message routing key* for routing, message producers need to specify it:

The routing key will then be compared for equality with routing keys on bindings, and consumers that subscribed with the same routing key each get a copy of the message.

Direct Exchanges and Load Balancing of Messages

Direct exchanges are often used to distribute tasks between multiple workers (instances of the same application) in a round robin manner. When doing so, it is important to understand that, in AMQP 0.9.1, *messages are load balanced between consumers and not between queues*.

The *Queues and Consumers* guide provides more information on this subject.

Pre-declared direct exchanges

AMQP 0.9.1 brokers must implement a direct exchange type and pre-declare two instances:

- `amq.direct`
- *“”* exchange known as *default exchange* (unnamed, referred to as an empty string)

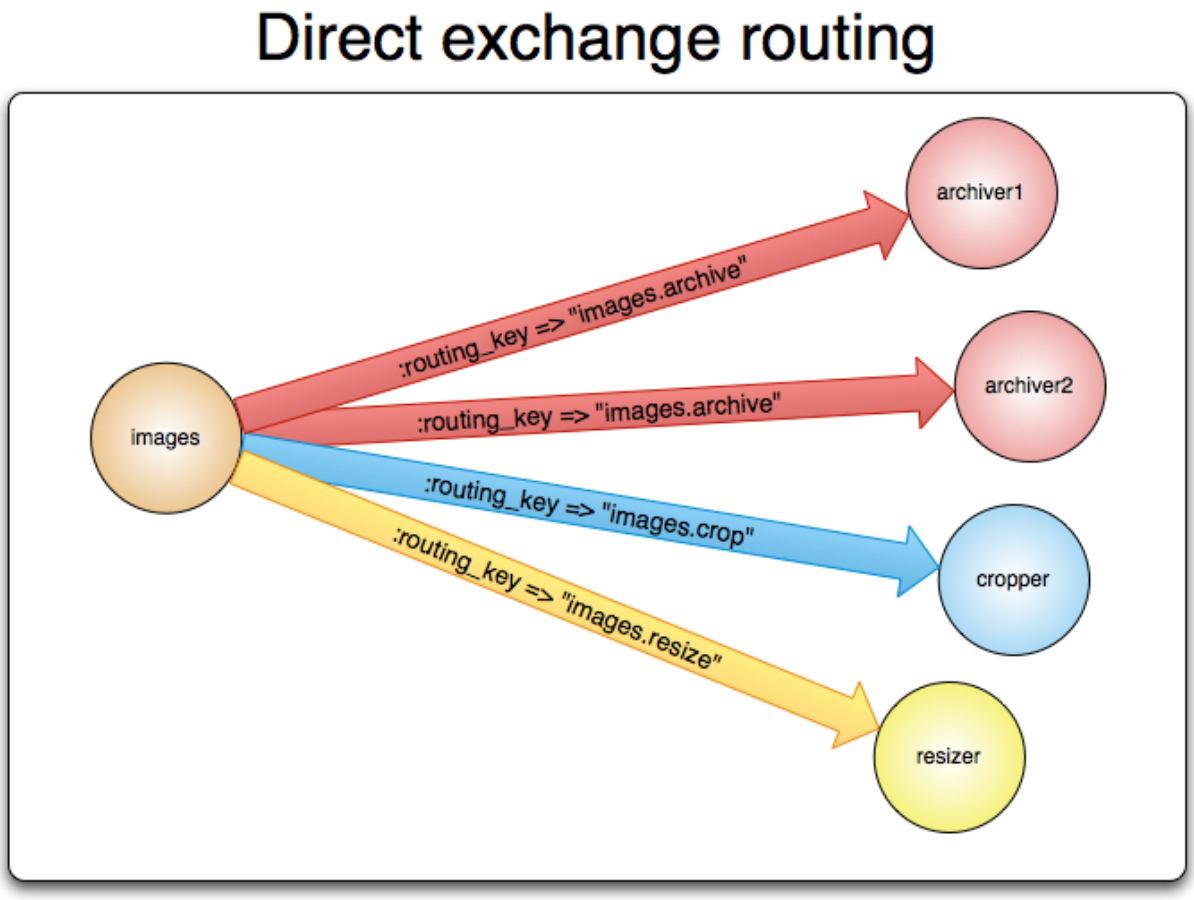

Fig. 2: direct exchange routing

Applications can rely on those exchanges always being available to them. Each vhost has separate instances of those exchanges, they are *not shared across vhosts* for obvious reasons.

Default exchange

The default exchange is a direct exchange with no name pre-declared by the broker. It has one special property that makes it very useful for simple applications, namely that *every queue is automatically bound to it with a routing key which is the same as the queue name*.

For example, when you declare a queue with the name of "search.indexing.online", RabbitMQ will bind it to the default exchange using "search.indexing.online" as the routing key. Therefore a message published to the default exchange with routing key = "search.indexing.online" will be routed to the queue "search.indexing.online". In other words, the default exchange makes it *seem like it is possible to deliver messages directly to queues*, even though that is not technically what is happening.

Direct Exchange Use Cases

Direct exchanges can be used in a wide variety of cases:

- Direct (near real-time) messages to individual players in an MMO game
- Delivering notifications to specific geographic locations (for example, points of sale)
- Distributing tasks between multiple instances of the same application all having the same function, for example, image processors
- Passing data between workflow steps, each having an identifier (also consider using headers exchange)
- Delivering notifications to individual software services in the network

2.4.8 Topic Exchanges

How Topic Exchanges Route Messages

Topic exchanges route messages to one or many queues based on matching between a message routing key and the pattern that was used to bind a queue to an exchange. The topic exchange type is often used to implement various publish/subscribe pattern variations.

Topic exchanges are commonly used for the multicast routing of messages.

Topic exchanges can be used for broadcast routing, but fanout exchanges are usually more efficient for this use case.

Topic Exchange Routing Example

Two classic examples of topic-based routing are stock price updates and location-specific data (for instance, weather broadcasts). Consumers indicate which topics they are interested in (think of it like subscribing to a feed for an individual tag of your favourite blog as opposed to the full feed).

A routing pattern consists of several words separated by dots, in a similar way to URI path segments being joined by slash. A few of examples:

- asia.southeast.thailand.bangkok
- sports.basketball
- usa.nasdaq.aapl

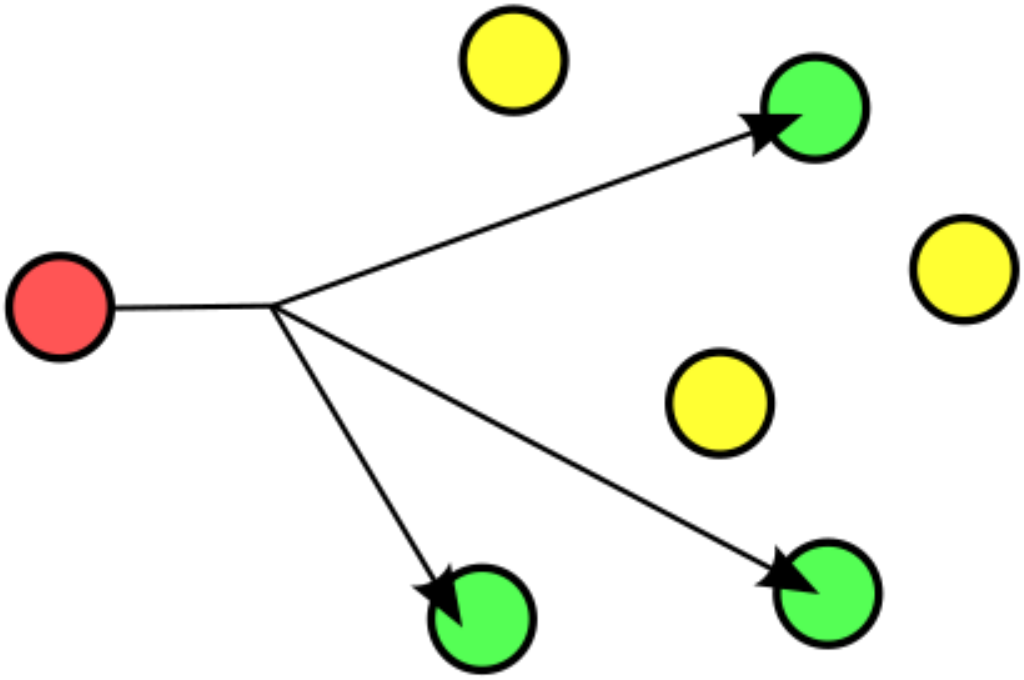

- tasks.search.indexing.accounts

The following routing keys match the "americas.south.#" pattern:

- americas.south
- americas.south.*brazil*
- americas.south.*brazil.saopaolo*
- americas.south.*chile.santiago*

In other words, the "#" part of the pattern matches 0 or more words.

For the pattern "americas.south.*", some matching routing keys are:

- americas.south.*brazil*
- americas.south.*chile*
- americas.south.*peru*

but not

- americas.south
- americas.south.chile.santiago

As you can see, the "*" part of the pattern matches 1 word only.

Topic Exchange Use Cases

Topic exchanges have a very broad set of use cases. Whenever a problem involves multiple consumers/applications that selectively choose which type of messages they want to receive, the use of topic exchanges should be considered. To name a few examples:

- Distributing data relevant to specific geographic location, for example, points of sale
- Background task processing done by multiple workers, each capable of handling specific set of tasks
- Stocks price updates (and updates on other kinds of financial data)
- News updates that involve categorization or tagging (for example, only for a particular sport or team)
- Orchestration of services of different kinds in the cloud
- Distributed architecture/OS-specific software builds or packaging where each builder can handle only one architecture or OS

2.4.9 Publishing messages

```php
<?php

$exchange->publish(
    'some message',
    'routing_key',
    Constants::AMQP_NOPARAM,
    [
        'arguments' => [
            'arg1' => 'value'
        ],
    ]
);
```

Data serialization

You are encouraged to take care of data serialization before publishing (i.e. by using JSON, Thrift, Protocol Buffers or some other serialization library). Note that because AMQP is a binary protocol, text formats like JSON largely lose their advantage of being easy to inspect as data travels across the network, so if bandwidth efficiency is important, consider using MessagePack or Protocol Buffers.

A few popular options for data serialization are:

- JSON
- BSON
- Message Pack
- XML
- Protocol Buffers

Message attributes

RabbitMQ messages have various metadata attributes that can be set when a message is published. Some of the attributes are well-known and mentioned in the AMQP 0.9.1 specification, others are specific to a particular application. Well-known attributes are listed here as options that HumusAmqp takes:

- :persistent
- :delivery_mode
- :mandatory
- :timestamp
- :expiration
- :type
- :reply_to
- :content_type
- :content_encoding
- :correlation_id
- :priority
- :cluster_id
- :user_id
- :app_id
- :message_id

All other attributes can be added to a *headers table*.

An example:

```
<?php

$exchange->publish(
    '{"foo": "bar"}',
    'routing_key',
    Constants::AMQP_NOPARAM,
    [
        'app_id' => 'amqp.example',
        'type' => 'kinda.checkin',
        'headers' => [
            'latitude' => 59.35,
            'longituide' => 18.0666667
        ],
        'timestamp' => time(),
        'correlation_id' => 'r-1',
        'content_type' => 'application/json',
        'delivery_mode' => 2,
        'content_encoding' => 'UTF-8',
    ]
);
```

:routing_key

Used for routing messages depending on the exchange type and configuration.

:persistent

When set to true, RabbitMQ will persist message to disk.

:mandatory

This flag tells the server how to react if the message cannot be routed to a queue. If this flag is set to true, the server will return an unroutable message to the producer with a `basic.return` AMQP method. If this flag is set to false, the server silently drops the message.

:content_type

MIME content type of message payload. Has the same purpose/semantics as HTTP Content-Type header.

:content_encoding

MIME content encoding of message payload. Has the same purpose/semantics as HTTP Content-Encoding header.

:priority

Message priority, from 0 to 9.

:message_id

Message identifier as a string. If applications need to identify messages, it is recommended that they use this attribute instead of putting it into the message payload.

:reply_to

Commonly used to name a reply queue (or any other identifier that helps a consumer application to direct its response). Applications are encouraged to use this attribute instead of putting this information into the message payload.

:correlation_id

ID of the message that this message is a reply to. Applications are encouraged to use this attribute instead of putting this information into the message payload.

:type

Message type as a string. Recommended to be used by applications instead of including this information into the message payload.

:user_id

Sender's identifier. Note that RabbitMQ will check that the value of this attribute is the same as username AMQP connection was authenticated with, it SHOULD NOT be used to transfer, for example, other application user ids or be used as a basis for some kind of Single Sign-On solution.

:app_id

Application identifier string, for example, "eventoverse" or "webcrawler"

:timestamp

Timestamp of the moment when message was sent, in seconds since the Epoch

:expiration

Message expiration specification as a string

:arguments

A map of any additional attributes that the application needs. Nested hashes are supported. Keys must be strings.

It is recommended that application authors use well-known message attributes when applicable instead of relying on custom headers or placing information in the message body. For example, if your application messages have priority, publishing timestamp, type and content type, you should use the respective AMQP message attributes instead of reinventing the wheel.

Validated User ID

In some scenarios it is useful for consumers to be able to know the identity of the user who published a message. RabbitMQ implements a feature known as validated User ID. If this property is set by a publisher, its value must be the same as the name of the user used to open the connection. If the user-id property is not set, the publisher's identity is not validated and remains private.

Publishing Callbacks and Reliable Delivery in Distributed Environments

A commonly asked question about RabbitMQ clients is "how to execute a piece of code after a message is received".

Message publishing with HumusAmqp happens in several steps:

- AMQPExchange::publish takes a payload and various metadata attributes
- Resulting payload is staged for writing
- On the next event loop tick, data is transferred to the OS kernel using one of the underlying NIO APIs
- OS kernel buffers data before sending it
- Network driver may also employ buffering

As you can see, "when data is sent" is a complicated issue and while methods to flush buffers exist, flushing buffers does not guarantee that the data was received by the broker because it might have crashed while data was travelling down the wire.

The only way to reliably know whether data was received by the broker or a peer application is to use message acknowledgements. This is how TCP works and this approach is proven to work at the enormous scale of the modern Internet. AMQP 0.9.1 fully embraces this fact and HumusAmqp follows.

In cases when you cannot afford to lose a single message, AMQP 0.9.1 applications can use one (or a combination of) the following protocol features:

- Publisher confirms (a RabbitMQ-specific extension to AMQP 0.9.1)
- Publishing messages as mandatory
- Transactions (these introduce noticeable overhead and have a relatively narrow set of use cases)

A more detailed overview of the pros and cons of each option can be found in a blog post that introduces Publisher Confirms extension by the RabbitMQ team. The next sections of this guide will describe how the features above can be used with HumusAmqp.

Publishing messages as mandatory

When publishing messages, it is possible to use the `:mandatory` option to publish a message as "mandatory". When a mandatory message cannot be *routed* to any queue (for example, there are no bindings or none of the bindings match), the message is returned to the producer.

```
<?php

$connection = new \Humus\Amqp\Driver\AmqpExtension\Connection();
$connection->connect();

$channel = $connection->newChannel();

$exchange = $channel->newExchange();
$exchange->setName('my-exchange');
```

(continues on next page)

(continued from previous page)

```
$exchange->publish(
    '{"foo": "bar"}',
    'routing_key',
    Constants::AMQP_MANDATORY
);

$this->channel->setReturnCallback(
    function (
        int $replyCode,
        string $replyText,
        string $exchange,
        string $routingKey,
        Envelope $envelope,
        string $body
    ) use (&$result) {
        $result[] = 'Message returned';
        $result[] = func_get_args();
        return false;
    }
);

try {
    $this->channel->waitForBasicReturn();
} catch (\Exception $e) {
    //$result[] = get_class($e) . ': ' . $e->getMessage(); //@todo: make php amqplib
↪throw these exceptions
}

var_dump($result);
```

Returned messages

When a message is returned, the application that produced it can handle that message in different ways:

- Store it for later redelivery in a persistent store
- Publish it to a different destination
- Log the event and discard the message

A returned message handler has access to AMQP method (`basic.return`) information, message metadata and payload (as a byte array). The metadata and message body are returned without modifications so that the application can store the message for later redelivery.

Publishing Persistent Messages

Messages potentially spend some time in the queues to which they were routed before they are consumed. During this period of time, the broker may crash or experience a restart. To survive it, messages must be persisted to disk. This has a negative effect on performance, especially with network attached storage like NAS devices and Amazon EBS. AMQP 0.9.1 lets applications trade off performance for durability, or vice versa, on a message-by-message basis.

To publish a persistent message, use the `:persistent` option:

Note that in order to survive a broker crash, the messages MUST be persistent and the queue that they were routed to MUST be durable.

Durability and Message Persistence provides more information on the subject.

Publishing In Multi-threaded Environments

When using HumusAmqp in multi-threaded environments, the rule of thumb is: avoid sharing channels across threads.

In other words, publishers in your application that publish from separate threads should use their own channels. The same is a good idea for consumers.

2.4.10 Headers exchanges

Now that message attributes and publishing have been introduced, it is time to take a look at one more core exchange type in AMQP 0.9.1. It is called the *headers exchange type* and is quite powerful.

How headers exchanges route messages

An Example Problem Definition

The best way to explain headers-based routing is with an example. Imagine a distributed continuous integration system that distributes builds across multiple machines with different hardware architectures (x86, IA-64, AMD64, ARM family and so on) and operating systems. It strives to provide a way for a community to contribute machines to run tests on and a nice build matrix like the one WebKit uses. One key problem such systems face is build distribution. It would be nice if a messaging broker could figure out which machine has which OS, architecture or combination of the two and route build request messages accordingly.

A headers exchange is designed to help in situations like this by routing on multiple attributes that are more easily expressed as message metadata attributes (headers) rather than a routing key string.

Routing on Multiple Message Attributes

Headers exchanges route messages based on message header matching. Headers exchanges ignore the routing key attribute. Instead, the attributes used for routing are taken from the "headers" attribute. When a queue is bound to a headers exchange, the `:arguments` attribute is used to define matching rules:

```
<?php

return [
    'dependencies' => [
        'factories' => [
            Driver::class => Humus\Amqp\Container\DriverFactory::class,
            'default-amqp-connection' =>␣
↪[Humus\Amqp\Container\ConnectionFactory::class, 'default'],
        ],
    ],
    'humus' => [
        'amqp' => [
            'driver' => 'php-amqplib',
            'connection' => [
                'default' => [
                    'type' => 'socket',
                    'host' => 'localhost',
                    'port' => 5672,
```

(continues on next page)

(continued from previous page)

```
                    'login' => 'guest',
                    'password' => 'guest',
                    'vhost' => '/',
                    'persistent' => false,
                    'read_timeout' => 3, //sec, float allowed
                    'write_timeout' => 1, //sec, float allowed
                ],
            ],
            'exchange' => [
                'header-exchange' => [
                    'name' => 'header-exchange',
                    'type' => 'headers',
                    'connection' => 'default-amqp-connection',
                ],
            ],
            'queue' => [
                'myqueue-1' => [
                    'name' => 'myqueue',
                    'exchanges' => [
                        'header-exchange' => [
                            [
                                'arguments' => [
                                    'os' => 'linux',
                                    'x-match' => 'all'
                                ],
                            ]
                        ],
                    ],
                    'connection' => 'default-amqp-connection',
                ],
                'myqueue-2' => [
                    'name' => 'myqueue',
                    'exchanges' => [
                        'header-exchange' => [
                            [
                                'arguments' => [
                                    'os' => 'osx',
                                    'x-match' => 'any'
                                ],
                            ]
                        ]
                    ],
                    'connection' => 'default-amqp-connection',
                ],
            ],
        ],
    ],
];
```

When matching on one header, a message is considered matching if the value of the header equals the value specified upon binding. An example that demonstrates headers routing:

```
<?php

$exchange->publish(
    '8 cores/Linux',
```

(continues on next page)

(continued from previous page)

```
    '',
    Constants::AMQP_NOPARAM,
    [
        'headers' => [
            'os' => 'linux',
            'cores' => 8
        ],
    ]
);

$exchange->publish(
    '4 cores/OS X',
    '',
    Constants::AMQP_NOPARAM,
    [
        'headers' => [
            'os' => 'osx',
            'cores' => 4
        ],
    ]
);
```

When executed, it outputs

```
myqueue-2 received 8 cores/Linux
```

The myqueue-1 has not matched, because of x-match: all

Matching All vs Matching One

It is possible to bind a queue to a headers exchange using more than one header for matching. In this case, the broker needs one more piece of information from the application developer, namely, should it consider messages with any of the headers matching, or all of them? This is what the "x-match" binding argument is for.

When the `"x-match"` argument is set to `"any"`, just one matching header value is sufficient. So in the example above, any message with a "cores" header value equal to 8 will be considered matching.

Headers Exchange Routing

When there is just one queue bound to a headers exchange, messages are routed to it if any or all of the message headers match those specified upon binding. Whether it is "any header" or "all of them" depends on the `"x-match"` header value. In the case of multiple queues, a headers exchange will deliver a copy of a message to each queue, just like direct exchanges do. Distribution rules between consumers on a particular queue are the same as for a direct exchange.

Headers Exchange Use Cases

Headers exchanges can be looked upon as "direct exchanges on steroids" and because they route based on header values, they can be used as direct exchanges where the routing key does not have to be a string; it could be an integer or a hash (dictionary) for example.

Some specific use cases:

- Transfer of work between stages in a multi-step workflow (routing slip pattern)

- Distributed build/continuous integration systems can distribute builds based on multiple parameters (OS, CPU architecture, availability of a particular package).

Pre-declared Headers Exchanges

RabbitMQ implements a headers exchange type and pre-declares one instance with the name of `"amq.match"`. RabbitMQ also pre-declares one instance with the name of `"amq.headers"`. Applications can rely on those exchanges always being available to them. Each vhost has a separate instance of those exchanges and they are *not shared across vhosts* for obvious reasons.

2.4.11 Custom Exchange Types

consistent-hash

The consistent hashing AMQP exchange type is a custom exchange type developed as a RabbitMQ plugin. It uses consistent hashing to route messages to queues. This helps distribute messages between queues more or less evenly.

A quote from the project README:

> In various scenarios, you may wish to ensure that messages sent to an exchange are consistently and equally distributed across a number of different queues based on the routing key of the message. You could arrange for this to occur yourself by using a direct or topic exchange, binding queues to that exchange and then publishing messages to that exchange that match the various binding keys.
>
> However, arranging things this way can be problematic:
>
> It is difficult to ensure that all queues bound to the exchange will receive a (roughly) equal number of messages without baking in to the publishers quite a lot of knowledge about the number of queues and their bindings.
>
> If the number of queues changes, it is not easy to ensure that the new topology still distributes messages between the different queues evenly.
>
> Consistent Hashing is a hashing technique whereby each bucket appears at multiple points throughout the hash space, and the bucket selected is the nearest higher (or lower, it doesn't matter, provided it's consistent) bucket to the computed hash (and the hash space wraps around). The effect of this is that when a new bucket is added or an existing bucket removed, only a very few hashes change which bucket they are routed to.
>
> In the case of Consistent Hashing as an exchange type, the hash is calculated from the hash of the routing key of each message received. Thus messages that have the same routing key will have the same hash computed, and thus will be routed to the same queue, assuming no bindings have changed.

x-random

The x-random AMQP exchange type is a custom exchange type developed as a RabbitMQ plugin by Jon Brisbin. A quote from the project README:

> It is basically a direct exchange, with the exception that, instead of each consumer bound to that exchange with the same routing key getting a copy of the message, the exchange type randomly selects a queue to route to.

This plugin is licensed under Mozilla Public License 1.1, same as RabbitMQ.

2.4.12 Using the Publisher Confirms Extension

Please refer to *RabbitMQ Extensions guide*

Message Acknowledgements and Their Relationship to Transactions and Publisher Confirms

Consumer applications (applications that receive and process messages) may occasionally fail to process individual messages, or might just crash. Additionally, network issues might be experienced. This raises a question - "when should the RabbitMQ remove messages from queues?" This topic is covered in depth in the *Queues guide*, including prefetching and examples.

In this guide, we will only mention how message acknowledgements are related to AMQP transactions and the Publisher Confirms extension. Let us consider a publisher application (P) that communications with a consumer (C) using AMQP 0.9.1. Their communication can be graphically represented like this:

We have two network segments, S1 and S2. Each of them may fail. A publisher (P) is concerned with making sure that messages cross S1, while the broker (B) and consumer (C) are concerned with ensuring that messages cross S2 and are only removed from the queue when they are processed successfully.

Message acknowledgements cover reliable delivery over S2 as well as successful processing. For S1, P has to use transactions (a heavyweight solution) or the more lightweight Publisher Confirms, a RabbitMQ-specific extension.

2.4.13 Binding Queues to Exchanges

Queues are bound to exchanges. This topic is described in detail in the *Queues and Consumers guide*.

2.4.14 Unbinding Queues from Exchanges

Queues are unbound from exchanges using. This topic is described in detail in the *Queues and Consumers guide*.

2.4.15 Deleting Exchanges

Explicitly deleting an Exchange

Exchanges are deleted using the `Humus\Amqp\Exchange#delete`:

```
<?php

$exchange->delete();
```

Auto-delete exchanges via configuration

Exchanges can be *auto-deleted*. To declare an exchange as auto-deleted, use the `:auto_delete` option on declaration:

```
<?php

return [
    'humus' => [
        'amqp' => [
            'exchange' => [
```

(continues on next page)

(continued from previous page)

```
                'header-exchange' => [
                    'name' => 'header-exchange',
                    'type' => 'headers',
                    'auto_delete' => true
                ],
            ],
        ],
    ],
];
```

An auto-deleted exchange is removed when the last queue bound to it is unbound.

2.4.16 Exchange durability vs Message durability

See *Durability guide*

2.4.17 Wrapping Up

Publishers publish messages to exchanges. Messages are then routed to queues according to rules called bindings that applications define. There are 4 built-in exchange types in RabbitMQ and it is possible to create custom types.

Messages have a set of standard properties (e.g. type, content type) and can carry an arbitrary map of headers.

2.4.18 What to Read Next

The documentation is organized as *a number of guides*, covering various topics.

We recommend that you read the following guides first, if possible, in this order:

- *HumusAmqp Producer's*
- *Queues and Consumers*
- *Bindings*
- *Consumers*
- *CLI*
- *Durability and Related Matters*
- *JSON RPC Server & Client*
- *RabbitMQ Extensions to AMQP 0.9.1*
- *Error Handling and Recovery*
- *Troubleshooting*
- *Deployment*

2.4.19 Tell Us What You Think!

Please take a moment to tell us what you think about this guide: Send an e-mail, say hello in the HumusAmqp gitter chat. or raise an issue on Github.

Let us know what was unclear or what has not been covered. Maybe you do not like the guide style or grammar or discover spelling mistakes. Reader feedback is key to making the documentation better.

2.5 HumusAmqp Producers

2.5.1 What are producers?

The concept of producers has nothing to do with the AMQP protocol itself. It's a feature added in HumusAmqp to allow sending messages with default attributes in a very comfortable way.

2.5.2 Concept of producers

The producer wraps the exchange and corresponding channel and allows to set default attributes and parameters for all messages sent.

Producer types

There are two built-in exchange types in HumusAmqp

- JSON
- Plain

The JSON Producer takes a message as first argument and json encodes it for you. Additional the following attributes are added by default:

```
[
    'content_type' => 'application/json',
    'content_encoding' => 'UTF-8',
    'delivery_mode' => 2,
]
```

This means all messages are persisted by the broker (delivery mode) and the content-type and content-encoding is already set for you.

The plain producer takes a message as first arguments and does nothing with it (so it should be a string or at least it should be possible to cast it to string, like integer and float). Additional the following attributes are added by default:

```
[
    'content_type' => 'text/plain',
    'content_encoding' => 'UTF-8',
    'delivery_mode' => 2,
]
```

Creating a producer

```
<?php

$connection = new \Humus\Amqp\Driver\AmqpExtension\Connection();
$connection->connect();

$channel = $connection->newChannel();
```

(continues on next page)

(continued from previous page)

```
$exchange = $channel->newExchange();
$exchange->setName('my-exchange');

$producer = new \Humus\Amqp\JsonProducer($exchange);
```

Using custom default headers

```
<?php

$producer = new \Humus\Amqp\JsonProducer($exchange, [
    'content_type' => 'application/json',
    'content_encoding' => 'UTF-8',
    'delivery_mode' => 2 // persistent,
    'app_id' => 'DemoApplication',
    'expiration' => 10000,
]);
```

This will additionally add the app_id and message expiration attributes.

Using configuration and factory

```
<?php

return [
    'dependencies' => [
        'factories' => [
            Driver::class => \Humus\Amqp\Container\DriverFactory::class,
            'default-amqp-connection' =>␣
↪[\Humus\Amqp\Container\ConnectionFactory::class, 'default'],
            'my-producer' => [\Humus\Amqp\Container\ProducerFactory::class, 'my-
↪producer'],
        ],
    ],
    'humus' => [
        'amqp' => [
            'driver' => 'php-amqplib',
            'connection' => [
                'default' => [
                    'type' => 'socket',
                    'host' => 'localhost',
                    'port' => 5672,
                    'login' => 'guest',
                    'password' => 'guest',
                    'vhost' => '/',
                    'persistent' => false,
                    'read_timeout' => 3, //sec, float allowed
                    'write_timeout' => 1, //sec, float allowed
                ],
            ],
            'exchange' => [
                'my-exchange' => [
                    'name' => 'my-exchange',
```

(continues on next page)

(continued from previous page)

```
                    'type' => 'direct',
                    'connection' => 'default-amqp-connection',
                    'auto_setup_fabric' => true,
                ],
            ],
            'producer' => [
                'my-producer' => [
                    'type' => 'json',
                    'exchange' => 'my-exchange',
                ],
            ],
        ],
    ],
];

$producer = $container->get('my-producer');
```

Publishing messages

```
<?php

$exchange->publish(
    'some message',
    'routing_key',
    Constants::AMQP_NOPARAM,
    [
        'arguments' => [
            'arg1' => 'value'
        ],
    ]
);
```

2.5.3 Publishing messages

```
<?php

$producer->publish('my message', 'routing_key');
```

2.5.4 Publishing messages transactional

```
<?php

$producer->startTransaction();

$producer->publish('my message', 'routing_key');

$producer->commitTransaction();
```

2.5.5 Publishing messages with confirm select

```php
<?php

$producer->confirmSelect();

$producer->setConfirmCallback(
    function (
        int $deliveryTag,
        bool $multiple = false
    ) use (&$cnt, &$result) {
        $result[] = 'Message acked';
        $result[] = func_get_args();
        return --$cnt > 0;
    },
    function (
        int $deliveryTag,
        bool $multiple,
        bool $requeue
    ) use (&$result) {
        $result[] = 'Message nacked';
        $result[] = func_get_args();
        return false;
    }
);

$producer->publish('my message', 'routing_key');

$producer->waitForConfirm();

var_dump($result);
```

2.5.6 Publishing messages as mandatory

```php
<?php

$producer->setReturnCallback(
    function (
        int $replyCode,
        string $replyText,
        string $exchange,
        string $routingKey,
        Envelope $envelope,
        string $body
    ) {
        throw new \RuntimeException('Message returned: ' . $replyText);
    }
);

$producer->publish(
    'my message',
    'routing_key',
    Constants::AMQP_MANDATORY
);

$producer->waitForBasicReturn();
```

Wrapping Up

Using a producer simplifies the client code when working with exchanges a lot by adding your needed default message attributes. Use them whenever possible instead of handling with the exchange directly.

What to Read Next

The documentation is organized as *a number of guides*, covering various topics.

We recommend that you read the following guides first, if possible, in this order:

- *Queues and Consumers*
- *Bindings*
- *Consumers*
- *CLI*
- *Durability and Related Matters*
- *JSON RPC Server & Client*
- *RabbitMQ Extensions to AMQP 0.9.1*
- *Error Handling and Recovery*
- *Troubleshooting*
- *Deployment*

Tell Us What You Think!

Please take a moment to tell us what you think about this guide: Send an e-mail, say hello in the HumusAmqp gitter chat. or raise an issue on Github.

Let us know what was unclear or what has not been covered. Maybe you do not like the guide style or grammar or discover spelling mistakes. Reader feedback is key to making the documentation better.

2.6 Queues

2.6.1 Queues in AMQP 0.9.1: Overview

What are AMQP Queues?

Queues store and forward messages to consumers. They are similar to mailboxes in SMTP. Messages flow from producing applications to *Exchanges* that route them to queues and finally, queues deliver the messages to consumer applications (or consumer applications fetch messages as needed).

Note: Note that unlike some other messaging protocols/systems, messages are not delivered directly to queues. They are delivered to exchanges that route messages to queues using rules known as *bindings*.

AMQP 0.9.1 is a programmable protocol, so queues and bindings alike are declared by applications.

Concept of Bindings

A *binding* is an association between a queue and an exchange. Queues must be bound to at least one exchange in order to receive messages from publishers. Learn more about bindings in the *Bindings guide*.

Queue Attributes

Queues have several attributes associated with them:

- Name
- Exclusivity
- Durability
- Whether the queue is auto-deleted when no longer used
- Other metadata (sometimes called *X-arguments*)

These attributes define how queues can be used, their life-cycle, and other aspects of queue behavior.

2.6.2 Queue Names and Declaring Queues

Every AMQP queue has a name that identifies it. Queue names often contain several segments separated by a dot ".", in a similar fashion to URI path segments being separated by a slash "/", although almost any string can represent a segment (with some limitations - see below).

Before a queue can be used, it has to be *declared*. Declaring a queue will cause it to be created if it does not already exist. The declaration will have no effect if the queue does already exist and its attributes are the *same as those in the declaration*. When the existing queue attributes are not the same as those in the declaration a channel-level exception is raised. This case is explained later in this guide.

Explicitly Named Queues

Applications may pick queue names or ask the broker to generate a name for them. The configure a queue with an explicit name:

```
<?php

return [
    'dependencies' => [
        'factories' => [
            Driver::class => Humus\Amqp\Container\DriverFactory::class,
            'default-amqp-connection' =>␣
↪[Humus\Amqp\Container\ConnectionFactory::class, 'default'],
        ],
    ],
    'humus' => [
        'amqp' => [
            'driver' => 'php-amqplib',
            'connection' => [
                'default' => [
                    'type' => 'socket',
                    'host' => 'localhost',
                    'port' => 5672,
                    'login' => 'guest',
```

(continues on next page)

(continued from previous page)

```
                    'password' => 'guest',
                    'vhost' => '/',
                    'persistent' => true,
                    'read_timeout' => 3, //sec, float allowed
                    'write_timeout' => 1, //sec, float allowed
                ],
            ],
            'exchange' => [
                'demo-exchange' => [
                    'name' => 'demo-exchange',
                    'type' => 'direct'
                ],
            ],
            'queue' => [
                'my-queue' => [
                    'name' => 'my-queue',
                    'exchanges' => [
                        'demo-exchange' => [],
                    ],
                ],
            ],
        ],
    ],
];
```

or

```
<?php

$queue->setName('my-queue');
```

Server-named queues

To ask an AMQP broker to generate a unique queue name for you, pass an *empty string* as the queue name argument. A generated queue name (like *amq.gen-JZ46KgZEOZWg-pAScMhhig*) will be assigned to the queue instance that the method returns:

```
<?php

return [
    'dependencies' => [
        'factories' => [
            Driver::class => Humus\Amqp\Container\DriverFactory::class,
            'default-amqp-connection' =>␣
↪[Humus\Amqp\Container\ConnectionFactory::class, 'default'],
        ],
    ],
    'humus' => [
        'amqp' => [
            'driver' => 'php-amqplib',
            'connection' => [
                'default' => [
                    'type' => 'socket',
                    'host' => 'localhost',
                    'port' => 5672,
```

(continues on next page)

(continued from previous page)

```
                    'login' => 'guest',
                    'password' => 'guest',
                    'vhost' => '/',
                    'persistent' => true,
                    'read_timeout' => 3, //sec, float allowed
                    'write_timeout' => 1, //sec, float allowed
                ],
            ],
            'exchange' => [
                'demo-exchange' => [
                    'name' => 'demo-exchange',
                    'type' => 'direct'
                ],
            ],
            'queue' => [
                'my-queue' => [
                    'name' => '',
                    'exchanges' => [
                        'demo-exchange' => [],
                    ],
                ],
            ],
        ],
    ],
];
```

Note: While it is common to declare server-named queues as `:exclusive`, it is not necessary.

Reserved Queue Name Prefix

Queue names starting with "amq." are reserved for server-named queues and queues for internal use by the broker. Attempts to declare a queue with a name that violates this rule will result in an exception with reply code `403` and an exception message similar to this:

```
ACCESS_REFUSED - exchange name 'amq.queue' contains reserved prefix 'amq.*'
```

This error results in the channel that was used for the declaration being forcibly closed by RabbitMQ. If the program subsequently tries to communicate with RabbitMQ using the same channel without re-opening it then the AMQP Extension will throw an `\Humus\Amqp\Exception\ChannelException' with message 'Could not create exchange. No channel available.`

Queue Re-Declaration With Different Attributes

When queue declaration attributes are different from those that the queue already has, a channel-level exception with code `406` will be raised. The reply text will be similar to this:

```
PRECONDITION_FAILED - cannot redeclare exchange 'foo' in vhost '/' with different␣
↪type, durable, internal or autodelete value
```

This error results in the channel that was used for the declaration being forcibly closed by RabbitMQ. If the program subsequently tries to communicate with RabbitMQ using the same channel without re-opening it then HumusAmqp will throw an `Humus\Amqp\Exception\ChannelException' with message 'Could not create`

`exchange. No channel available`. In order to continue communications in the same program after such an error, a different channel would have to be used.

2.6.3 Queue Life-cycle Patterns

According to the AMQP 0.9.1 specification, there are two common message queue life-cycle patterns:

- Durable queues that are shared by many consumers and have an independent existence: i.e. they will continue to exist and collect messages whether or not there are consumers to receive them.
- Temporary queues that are private to one consumer and are tied to that consumer. When the consumer disconnects, the message queue is deleted.

There are some variations of these, such as shared message queues that are deleted when the last of many consumers disconnects.

Let us examine the example of a well-known service like an event collector (event logger). A logger is usually up and running regardless of the existence of services that want to log anything at a particular point in time. Other applications know which queues to use in order to communicate with the logger and can rely on those queues being available and able to survive broker restarts. In this case, explicitly named durable queues are optimal and the coupling that is created between applications is not an issue.

Another example of a well-known long-lived service is a distributed metadata/directory/locking server like Apache Zookeeper, Google's Chubby or DNS. Services like this benefit from using well-known, not server-generated, queue names and so do any other applications that use them.

A different sort of scenario is in "a cloud setting" when some kind of worker/instance might start and stop at any time so that other applications cannot rely on it being available. In this case, it is possible to use well-known queue names, but a much better solution is to use server-generated, short-lived queues that are bound to topic or fanout exchanges in order to receive relevant messages.

Imagine a service that processes an endless stream of events — Twitter is one example. When traffic increases, development operations may start additional application instances in the cloud to handle the load. Those new instances want to subscribe to receive messages to process, but the rest of the system does not know anything about them and cannot rely on them being online or try to address them directly. The new instances process events from a shared stream and are the same as their peers. In a case like this, there is no reason for message consumers not to use queue names generated by the broker.

In general, use of explicitly named or server-named queues depends on the messaging pattern that your application needs. Enterprise Integration Patterns discusses many messaging patterns in depth and the RabbitMQ FAQ also has a section on use cases.

2.6.4 Declaring a Durable Shared Queue

To declare a durable shared queue, you pass a queue name that is a non-blank string and use the `:durable` option:

```
<?php

return [
    'dependencies' => [
        'factories' => [
            Driver::class => Humus\Amqp\Container\DriverFactory::class,
            'default-amqp-connection' =>␣
↪[Humus\Amqp\Container\ConnectionFactory::class, 'default'],
        ],
    ],
```

(continues on next page)

(continued from previous page)

```
    'humus' => [
        'amqp' => [
            'driver' => 'php-amqplib',
            'connection' => [
                'default' => [
                    'type' => 'socket',
                    'host' => 'localhost',
                    'port' => 5672,
                    'login' => 'guest',
                    'password' => 'guest',
                    'vhost' => '/',
                    'persistent' => true,
                    'read_timeout' => 3, //sec, float allowed
                    'write_timeout' => 1, //sec, float allowed
                ],
            ],
            'exchange' => [
                'demo-exchange' => [
                    'name' => 'demo-exchange',
                    'type' => 'direct'
                ],
            ],
            'queue' => [
                'my-queue' => [
                    'name' => 'demo-queue',
                    'exchanges' => [
                        'demo-exchange'
                    ],
                    'durable' => true
                ],
            ],
        ],
    ],
];
```

Note: Using the HumusAmqpContainerQueueFactory queues are durable by default, when no value is set.

2.6.5 Declaring a Temporary Exclusive Queue

To declare a server-named, exclusive, auto-deleted queue, pass “” (an empty string) as the queue name and use the `:exclusive` option:

```
<?php

return [
    'dependencies' => [
        'factories' => [
            Driver::class => Humus\Amqp\Container\DriverFactory::class,
            'default-amqp-connection' =>␣
↪[Humus\Amqp\Container\ConnectionFactory::class, 'default'],
        ],
    ],
    'humus' => [
```

(continues on next page)

(continued from previous page)

```
        'amqp' => [
            'driver' => 'php-amqplib',
            'connection' => [
                'default' => [
                    'type' => 'socket',
                    'host' => 'localhost',
                    'port' => 5672,
                    'login' => 'guest',
                    'password' => 'guest',
                    'vhost' => '/',
                    'persistent' => true,
                    'read_timeout' => 3, //sec, float allowed
                    'write_timeout' => 1, //sec, float allowed
                ],
            ],
            'exchange' => [
                'demo-exchange' => [
                    'name' => 'demo-exchange',
                    'type' => 'direct'
                ],
            ],
            'queue' => [
                'my-queue' => [
                    'name' => '',
                    'exchange' => 'demo-exchange',
                    'exclusive' => true
                ],
            ],
        ]
    ]
];
```

Exclusive queues may only be accessed by the current connection and are deleted when that connection closes. The declaration of an exclusive queue by other connections is not allowed and will result in a channel-level exception with the code `405 (RESOURCE_LOCKED)`

Exclusive queues will be deleted when the connection they were declared on is closed.

2.6.6 Checking if a Queue Exists

Sometimes it's convenient to check if a queue exists. To do so, at the protocol level you use `queue.declareQueue` with `passive` set to `true`. In response RabbitMQ responds with a channel exception if the queue does not exist. This will lead to an `\Humus\Amqp\QueueException` with message `NOT_FOUND - no queue 'test-queue' in vhost '/'`

2.6.7 Binding Queues with Routing Keys

In order to receive messages, a queue needs to be bound to at least one exchange. Most of the time binding is explcit (done by applications). **Please note:** All queues are automatically bound to the default unnamed RabbitMQ direct exchange with a routing key that is the same as the queue name, see: *Exchanges and Publishing* guide for more details.

```
<?php
```

(continues on next page)

(continued from previous page)

```
return [
    'dependencies' => [
        'factories' => [
            Driver::class => Humus\Amqp\Container\DriverFactory::class,
            'default-amqp-connection' =>␣
↪[Humus\Amqp\Container\ConnectionFactory::class, 'default'],
        ],
    ],
    'humus' => [
        'amqp' => [
            'driver' => 'php-amqplib',
            'connection' => [
                'default' => [
                    'type' => 'socket',
                    'host' => 'localhost',
                    'port' => 5672,
                    'login' => 'guest',
                    'password' => 'guest',
                    'vhost' => '/',
                    'persistent' => false,
                    'read_timeout' => 3, //sec, float allowed
                    'write_timeout' => 1, //sec, float allowed
                ],
            ],
            'exchange' => [
                'demo-exchange' => [
                    'name' => 'demo-exchange',
                    'type' => 'direct'
                    'connection' => 'default-amqp-connection',
                ],
            ],
            'queue' => [
                'my-queue' => [
                    'name' => 'demo-queue',
                    'exchanges' => [
                        'demo-exchange' => [
                            [
                                'routing_keys => [
                                    'v1.0.*',
                                    'v1.1.0',
                                    'v2.0.0'
                                ],
                            ],
                        ],
                    ],
                    'connection' => 'default-amqp-connection',
                ],
            ],
        ],
    ],
];
```

2.6.8 Unbinding Queues From Exchanges

To unbind a queue from an exchange use the `AMQPQueue#unbind` function:

```php
<?php

$queue->unbind('exchange-name');
```

Note: Trying to unbind a queue from an exchange that the queue was never bound to will result in a channel-level exception.

2.6.9 Purging queues

It is possible to purge a queue (remove all of the messages from it) using the `\Humus\Amqp\Queue#purge` method:

```php
<?php

$queue->purge();
```

Note: When a server named queue is declared, it is empty, so for server-named queues, there is no need to purge them before they are used.

2.6.10 Deleting Queues

Queues can be deleted either indirectly or directly. To delete a queue indirectly you can include either of the following two arguments in the queue declaration:

- `:exclusive => true`
- `:auto_delete => true`

If the *exclusive* flag is set to true then the queue will be deleted when the connection that was used to declare it is closed.

If the *auto_delete* flag is set to true then the queue will be deleted when there are no more consumers subscribed to it. The queue will remain in existence until at least one consumer accesses it.

To delete a queue directly, use the `AMQPQueue#delete` method:

```php
<?php

$queue->delete();
```

When a queue is deleted, all of the messages in it are deleted as well.

2.6.11 Queue Durability vs Message Durability

See *Durability and Related Matters*

2.6.12 RabbitMQ Extensions Related to Queues

See *RabbitMQ Extensions to AMQP 0.9.1*

2.6.13 Wrapping Up

In RabbitMQ, queues can be client-named or server-named. For messages to be routed to queues, queues need to be bound to exchanges.

2.6.14 What to Read Next

The documentation is organized as *a number of guides*, covering various topics.

We recommend that you read the following guides first, if possible, in this order:

- *Queues and Consumers*
- *Bindings*
- *Consumers*
- *CLI*
- *Durability and Related Matters*
- *JSON RPC Server & Client*
- *RabbitMQ Extensions to AMQP 0.9.1*
- *Error Handling and Recovery*
- *Troubleshooting*
- *Deployment*

2.6.15 Tell Us What You Think!

Please take a moment to tell us what you think about this guide: Send an e-mail, say hello in the HumusAmqp gitter chat. or raise an issue on Github.

Let us know what was unclear or what has not been covered. Maybe you do not like the guide style or grammar or discover spelling mistakes. Reader feedback is key to making the documentation better.

2.7 Bindings

2.7.1 What Are AMQP 0.9.1 Bindings

Bindings are rules that exchanges use (among other things) to route messages to queues. To instruct an exchange E to route messages to a queue Q, Q has to *be bound* to E. Bindings may have an optional *routing key* attribute used by some exchange types. The purpose of the routing key is to selectively match only specific (matching) messages published to an exchange to the bound queue. In other words, the routing key acts like a filter.

To draw an analogy:

- Queue is like your destination in New York city
- Exchange is like JFK airport
- Bindings are routes from JFK to your destination. There may be no way, or more than one way, to reach it

Some exchange types use routing keys while some others do not (routing messages unconditionally or based on message metadata). If an AMQP message cannot be routed to any queue (for example, because there are no bindings for the exchange it was published to), it is either dropped or returned to the publisher, depending on the message attributes that the publisher has set.

If an application wants to connect a queue to an exchange, it needs to *bind* them. The opposite operation is called *unbinding*.

2.7.2 Binding Queues to Exchanges

In order to receive messages, a queue needs to be bound to at least one exchange. Most of the time binding is explicit (done by applications).

2.7.3 Example using config and factory

```
return [
    'dependencies' => [
        'factories' => [
            Driver::class => Humus\Amqp\Container\DriverFactory::class,
            'default-amqp-connection' =>␣
↪[Humus\Amqp\Container\ConnectionFactory::class, 'default'],
        ],
    ],
    'humus' => [
        'amqp' => [
            'driver' => 'php-amqplib',
            'connection' => [
                'default' => [
                    'type' => 'socket',
                    'host' => 'localhost',
                    'port' => 5672,
                    'login' => 'guest',
                    'password' => 'guest',
                    'vhost' => '/',
                    'persistent' => true,
                    'read_timeout' => 3,
                    'write_timeout' => 1,
                ],
            ],
            'exchanges' => [
                'demo-exchange' => [
                    'name' => 'demo-exchange',
                    'type' => 'direct',
                    'connection' => 'default-amqp-connection',
                ]
            ],
            'queues' => [
                'my-queue' => [
                    'name' => 'my-queue',
                    'exchanges' => [
                        'demo-exchange' => [],
                    ],
                    'connection' => 'default-amqp-connection',
                ]
            ]
```

(continues on next page)

(continued from previous page)

```
        ]
    ]
];
```

2.7.4 Unbinding queues from exchanges

```
<?php

$queue->unbind('exchange-name');
```

2.7.5 Exchange-to-Exchange Bindings

Exchange-to-Exchange bindings is a RabbitMQ extension to AMQP 0.9.1. It is covered in the *RabbitMQ Extensions to AMQP 0.9.1*

2.7.6 Bindings, Routing and Returned Messages

How RabbitMQ Routes Messages

After a message reaches RabbitMQ and before it reaches a consumer, several things happen:

- RabbitMQ needs to find one or more queues that the message needs to be routed to, depending on type of exchange
- RabbitMQ puts a copy of the message into each of those queues or decides to return the message to the publisher
- RabbitMQ pushes message to consumers on those queues or waits for applications to fetch them on demand

A more in-depth description is this:

- RabbitMQ needs to consult bindings list for the exchange the message was published to in order to find one or more queues that the message needs to be routed to (step 1)
- If there are no suitable queues found during step 1 and the message was published as mandatory, it is returned to the publisher (step 1b)
- If there are suitable queues, a *copy* of the message is placed into each one (step 2)
- If the message was published as mandatory, but there are no active consumers for it, it is returned to the publisher (step 2b)
- If there are active consumers on those queues and the basic.qos setting permits, message is pushed to those consumers (step 3)

The important thing to take away from this is that messages may or may not be routed and it is important for applications to handle unroutable messages.

Handling of Unroutable Messages

Unroutable messages are either dropped or returned to producers. RabbitMQ extensions can provide additional ways of handling unroutable messages: for example, RabbitMQ's Alternate Exchanges extension makes it possible to route unroutable messages to another exchange. HumusAmqp support for it is documented in the *RabbitMQ Extensions to AMQP 0.9.1*.

Exchanges and Publishing documentation guide provides more information on the subject, including full code examples.

2.7.7 What to Read Next

The documentation is organized as *a number of guides*, covering various topics.

We recommend that you read the following guides first, if possible, in this order:

- *Consumers*
- *CLI*
- *Durability and Related Matters*
- *JSON RPC Server & Client*
- *RabbitMQ Extensions to AMQP 0.9.1*
- *Error Handling and Recovery*
- *Troubleshooting*
- *Deployment*

2.7.8 Tell Us What You Think!

Please take a moment to tell us what you think about this guide: Send an e-mail, say hello in the HumusAmqp gitter chat. or raise an issue on Github.

Let us know what was unclear or what has not been covered. Maybe you do not like the guide style or grammar or discover spelling mistakes. Reader feedback is key to making the documentation better.

2.8 Consumers

HumusAmqp provides a default consumer implementation that suites most use-cases. If you have a special use-case, you can extend the abstract class or implement the consumer interface yourself.

2.8.1 Consumer Callbacks

In order do reduce extending consumer classes and avoid factory duplication, the consumer expects a delivery callback. This callback get executed every time a new message is delivered to the consumer. The consumer expects the callback to take 2 arguments (the envelope and the queue) and returns a delivery result. A very simple callback would look like this:

```
<?php

$callback = function(\Humus\Amqp\Envelope $envelope, \Humus\Amqp\Queue $queue) {
    echo $envelope->getBody();
    return \Humus\Amqp\DeliveryResult::MSG_ACK();
}
```

The delivery result will signal the consumer whether it should ack, nack, reject, reject and requeue or defer the message until the block size is reached or an timeout occurs. So you can handle blocks of messages (reducing network overhead).

2.8.2 Handling Messages in Batches

If you have collected messages (returned DeliveryResult::MSG_DEFER() in the delivery callback) and the block size or timeout is reached, the flush callback will get executed. If you did not specify a flush callback, it will FlushDeferredResult::MSG_ACK() leading to all messages collected being acknowledged at once. You have the possibility to add a custom flush callback where you have to take care whether to ack, nack or reject all messages.

2.8.3 Message Acknowledgements & Rejecting

Consumer applications — applications that receive and process messages - may occasionally fail to process individual messages, or will just crash. There is also the possibility of network issues causing problems. This raises a question — "When should the AMQP broker remove messages from queues?"

The AMQP 0.9.1 specification proposes two choices:

- After broker sends a message to an application (using either basic.deliver or basic.get-ok methods).
- After the application sends back an acknowledgement (using basic.ack AMQP method).

The former choice is called the *automatic acknowledgement model*, while the latter is called the *explicit acknowledgement model*. With the explicit model, the application chooses when it is time to send an acknowledgement. It can be right after receiving a message, or after persisting it to a data store before processing, or after fully processing the message (for example, successfully fetching a Web page, processing and storing it into some persistent data store).

Note: Acknowledgements are channel-specific. Applications MUST NOT receive messages on one channel and acknowledge them on another.

2.8.4 Logging

The consumer expects you to inject a logger instance (PsrLogLoggerInterface).

2.8.5 Error-Handling

By default, all errors are logged on the configured logger. If you want to, you can specify your own error callback that will get executed instead.

2.8.6 QoS — Prefetching messages

For cases when multiple consumers share a queue, it is useful to be able to specify how many messages each consumer can be sent at once before sending the next acknowledgement. This can be used as a simple load balancing technique to improve throughput if messages tend to be published in batches. For example, if a producing application sends messages every minute because of the nature of the work it is doing.

Imagine a website that takes data from social media sources like Twitter or Facebook during the Champions League (european soccer) final (or the Superbowl), and then calculates how many tweets mentioned a particular team during the last minute. The site could be structured as 3 applications:

- A crawler that uses streaming APIs to fetch tweets/statuses, normalizes them and sends them in JSON for processing by other applications ("app A").
- A calculator that detects what team is mentioned in a message, updates statistics and pushes an update to the Web UI once a minute ("app B").

- A Web UI that fans visit to see the stats ("app C").

In this imaginary example, the "tweets per second" rate will vary, but to improve the throughput of the system and to decrease the maximum number of messages that the AMQP broker has to hold in memory at once, applications can be designed in such a way that application "app B", the "calculator", receives 5000 messages and then acknowledges them all at once. The broker will not send message 5001 unless it receives an acknowledgement.

In AMQP 0.9.1 parlance this is known as *QoS* or *message prefetching*. Prefetching is configured on a per-channel basis.

The default implementation of the HumusAmqp's consumer works with prefetch count, so if you set the prefetch count to 50, a block size of 50 messages will be used.

2.8.7 Timeouts

The idle timeout takes effect, when there are no more messages coming in and you expect a block size > 1 (prefetch count). The wait timeout applies every time the consumer tries to fetch a message from the queue but doesn't receive any.

2.8.8 Set up the consumer

```php
<?php

$logger = new \Psr\Log\NullLogger();

$connection = new \Humus\Amqp\Driver\AmqpExtension\Connection();
$connection->connect();

$channel = $connection->newChannel();
// handle 20 messages or wait for timeout until flush deferred callback is executed
$channel->setPrefetchCount(20);

$queue = $channel->newQueue();
$queue->setName('test-queue');

$consumer = new \Humus\Amqp\CallbackConsumer(
    $queue,
    $logger,
    12.5, // idle timeout, float in seconds
    function (\Humus\Amqp\Envelope $envelope, \Humus\Amqp\Queue $queue) {
        echo $envelope->getBody();
        return \Humus\Amqp\DeliveryResult::MSG_DEFER();
    },
    function (\Humus\Amqp\Queue $queue) {
        return \Humus\Amqp\FlushDeferredResult::MSG_ACK();
    },
    null, // no custom error callback
    'demo-consumer-tag'
);

$consumer->consume(2000); // consume 2000 messages
```

2.8.9 Set up the consumer using config and factory

```
<?php

// declare callbacks as invokable classes first

namespace My
{
    class EchoCallback
    {
        public function __invoke(\Humus\Amqp\Envelope $envelope, \Humus\Amqp\Queue
↪$queue)
        {
            echo $envelope->getBody();
            return \Humus\Amqp\DeliveryResult::MSG_DEFER();
        }
    }

    class FlushDeferredCallback
    {
        public function (\Humus\Amqp\Queue $queue)
        {
            return \Humus\Amqp\FlushDeferredResult::MSG_ACK();
        }
    }
}

return [
    'dependencies' => [
        'factories' => [
            Driver::class => Humus\Amqp\Container\DriverFactory::class,
            'default-amqp-connection' =>␣
↪[Humus\Amqp\Container\ConnectionFactory::class, 'default'],
            \My\EchoCallback::class =>␣
↪\Zend\ServiceManager\Factory\InvokableFactory::class,
            \My\FlushDeferredCallback::class =>␣
↪\Zend\ServiceManager\Factory\InvokableFactory::class,
            \Psr\Log\NullLogger =>␣
↪\Zend\ServiceManager\Factory\InvokableFactory::class,
        ],
    ],
    'humus' => [
        'amqp' => [
            'driver' => 'amqp-extension',
            'connection' => [
                'default' => [
                    'host' => 'localhost',
                    'port' => 5672,
                    'login' => 'guest',
                    'password' => 'guest',
                    'vhost' => '/',
                    'persistent' => false,
                    'read_timeout' => 3,
                    'write_timeout' => 1,
                ],
            ],
            'queue' => [
```

(continues on next page)

(continued from previous page)

```
                'my-queue' => [
                'name' => 'demo-queue',
                'connection' => 'default-amqp-connection',
                'exchanges' => [
                    'demo-exchange' => [
                        [
                            'routing_keys => [
                                'v1.0.*',
                                'v1.1.0',
                                'v2.0.0'
                            ],
                        ],
                    ],
                ],
            ],
            'callback_consumer' => [
                'demo-consumer' => [
                    'queue' => 'demo-queue',
                    'delivery_callback' => \My\EchoCallback::class,
                    'flush_callback' => \My\FlushDeferredCallback::class,
                    'logger' => \Psr\Log\NullLogger::class,
                    'idle_timeout' => 12.5,
                    'consumer_tag' => 'demo-consumer-tag',
                    'qos' => [
                        'prefetch_count' => 50,
                    ],
                ],
            ],
        ],
    ],
];

$consumer = $container->get('demo-consumer');
$consumer->consume(2000);
```

2.8.10 Using Multiple Consumers Per Queue

It is possible to have multiple non-exclusive consumers on queues. In that case, messages will be distributed between them according to prefetch levels of their channels (more on this later in this guide). If prefetch values are equal for all consumers, each consumer will get about the same number of messages.

2.8.11 Starting a consumer from CLI

This requires setting up the consumer via config and container factory.

```
$ ./vendor/bin/humus-amqp consumer -n demo-consumer -a 2000
```

See: *Running from CLI* for more informations.

2.8.12 Killing a Consumer gracefully

You can send a SIGUSER1 signal to gracefully shutdown the consumer.

```
kill -10 23453
```

Where 23453 is the process id of the consumer process.

2.8.13 What to Read Next

The documentation is organized as *a number of guides*, covering various topics.

We recommend that you read the following guides first, if possible, in this order:

- *CLI*
- *Durability and Related Matters*
- *JSON RPC Server & Client*
- *RabbitMQ Extensions to AMQP 0.9.1*
- *Error Handling and Recovery*
- *Troubleshooting*
- *Deployment*

2.8.14 Tell Us What You Think!

Please take a moment to tell us what you think about this guide: Send an e-mail, say hello in the HumusAmqp gitter chat. or raise an issue on Github.

Let us know what was unclear or what has not been covered. Maybe you do not like the guide style or grammar or discover spelling mistakes. Reader feedback is key to making the documentation better.

2.9 Running from CLI

In order to run cli commands, you need to setup your connection, exchange and queue configuration first.

You can run cli commands like this:

```
$ ./vendor/bin/humus-amqp
```

2.9.1 Setting up the CLI

When you run the humus-amqp cli you may encounter this error on first run:

```
You are missing a "humus-amqp-config.php" or "config/humus-amqp-config.php" file in␣
↪your
project, which is required to get the Humus Amqp Console working. You can use the
following sample as a template:

<?php

use Humus\Amqp\Console\ConsoleRunner;

// replace with file to your own project bootstrap
```

(continues on next page)

(continued from previous page)

```
require_once 'bootstrap.php';

// replace with mechanism to retrieve the container in your app
$container = GetContainer();

return ConsoleRunner::createHelperSet($container);
```

To solve this, simply add a file *config/humus-amqp-config.php* file and use the template to setup the cli config.

2.9.2 Setup-Fabric

To setup all exchanges and queues configured:

```
$ ./vendor/bin/humus-amqp setup-fabric
```

This will create all exchanges and queues.

2.9.3 Running consumers

To start a consumer:

```
$ ./vendor/bin/humus-amqp consumer -n myconsumer -a 100
```

This will start the myconsumer and consume 100 messages until if stops or times out.

2.9.4 Running JSON-RPC servers

To start a JSON-RPC server

```
$ ./vendor/bin/humus-amqp json_rpc_server -n myserver -a 100
```

This will start the myserver and consume 100 messages until if stops or times out.

2.9.5 List amqp types

Show availables connections, exchanges, queues, callback_consumers, producers, json_rpc_clients and json_rpc_servers

```
$ ./vendor/bin/humus-amqp show -t exchanges
```

This will list all known exchanges.

2.9.6 Purge queues

To purge a queue:

```
$ ./vendor/bin/humus-amqp purge-queue -p myqueue
```

This will remove all messages from the given queue.

2.9.7 Publishing from CLI

To publish a message to an exchane via CLI:

```
$ ./vendor/bin/humus-amqp publish-message -p myproducer -m "my text" -c -r "my.
↪routing.key"
```

This will send a message with body "my text" and routing key "my.routing.key" via the "myproducer"-producer using confirm select mode.

2.9.8 Troubleshooting

If you have read this guide and still have issues with connecting, check our *Troubleshooting guide* and feel free to raise an issue at Github.

2.9.9 What to Read Next

The documentation is organized as *a number of guides*, covering various topics.

We recommend that you read the following guides first, if possible, in this order:

- *Durability and Related Matters*
- *RabbitMQ Extensions to AMQP 0.9.1*
- *Error Handling and Recovery*
- *Troubleshooting*
- *Deployment*

2.9.10 Tell Us What You Think!

Please take a moment to tell us what you think about this guide: Send an e-mail, say hello in the HumusAmqp gitter chat. or raise an issue on Github.

Let us know what was unclear or what has not been covered. Maybe you do not like the guide style or grammar or discover spelling mistakes. Reader feedback is key to making the documentation better.

2.10 Durability

2.10.1 Entity durability and message persistence

Exchange Durability

AMQP separates the concept of entity durability (queues, exchanges) from message persistence. Exchanges can be durable or transient. Durable exchanges survive broker restart, transient exchanges do not (they have to be redeclared when the broker comes back online), however, not all scenarios and use cases mandate exchanges to be durable.

To create a durable exchange, declare it with the `:durable => true` argument.

Queue Durability

Queues can be durable or transient. Durable queues survive broker restart, transient queues do not (they have to be redeclared when the broker comes back online), however, not all scenarios and use cases mandate queues to be durable.

To create a durable queue, declare it with the `:durable => true` argument.

Durability of a queue does not make *messages* that are routed to that queue durable. If a broker is taken down and then brought back up, durable queues will be re-declared during broker startup, however, only *persistent* messages will be recovered.

Binding Durability

Bindings of durable queues to durable exchanges are automatically durable and are restored after a broker restart. The AMQP 0.9.1 specification states that the binding of durable queues to transient exchanges must be allowed. In this case, since the exchange would not survive a broker restart, neither would any bindings to such and exchange.

Message Persistence

Messages may be published as persistent and this, in conjunction with queue durability, is what makes an AMQP broker persist them to disk. If the server is restarted, the system ensures that received persistent messages in durable queues are not lost. Simply publishing a message to a durable exchange or the fact that a queue to which a message is routed is durable does not make that message persistent. Message persistence depends on the persistence mode of the message itself.

Note: Publishing persistent messages affects performance (just like with data stores, durability comes at a certain cost to performance).

Clustering and High Availability

To achieve the degree of durability that critical applications need, it is necessary but not enough to use durable queues, exchanges and persistent messages. You need to use a cluster of brokers because otherwise, a single hardware problem may bring a broker down completely.

RabbitMQ offers a number of high availability features for both scenarios with more (LAN) and less (WAN) reliable network connections.

See the RabbitMQ clustering and high availability guides for in-depth discussion of this topic.

Highly Available (Mirrored) Queues

Whilst the use of clustering provides for greater durability of critical systems, in order to achieve the highest level of resilience for queues and messages, high availability configuration should be used. This is because although exchanges and bindings survive the loss of individual nodes by using clustering, messages do not. Without mirroring, queue contents reside on exactly one node, thus the loss of a node will cause message loss.

See the RabbitMQ high availability guide for more information about mirrored queues.

2.10.2 What to Read Next

The documentation is organized as *a number of guides*, covering various topics.

We recommend that you read the following guides first, if possible, in this order:

- *RabbitMQ Extensions to AMQP 0.9.1*
- *Error Handling and Recovery*
- *Troubleshooting*
- *Deployment*

2.10.3 Tell Us What You Think!

Please take a moment to tell us what you think about this guide: Send an e-mail, say hello in the HumusAmqp gitter chat. or raise an issue on Github.

Let us know what was unclear or what has not been covered. Maybe you do not like the guide style or grammar or discover spelling mistakes. Reader feedback is key to making the documentation better.

2.11 JSON RPC Server & Client

2.11.1 Setup the JSON RPC Server

Assuming you want to use the provided interop-factories, let's start with a sample configuration:

A sample configuration might look like this, more details and explanation will be in the coming chapters.

```
return [
    'dependencies' => [
        'factories' => [
            Driver::class => Container\DriverFactory::class,
            'default-amqp-connection' => [Container\ConnectionFactory::class, 'default
↪'],
            'demo-rpc-server' => [Container\JsonRpcServerFactory::class, 'demo-rpc-
↪server'],
            'timestwo' => function () {
                return function (\Humus\Amqp\JsonRpc\Request $request):␣
↪\Humus\Amqp\JsonRpc\Response {
                    return \Humus\Amqp\JsonRpc\JsonRpcResponse::withResult($request->
↪id(), $request->params() * 2);
                };
            },
        ],
    ],
    'humus' => [
        'amqp' => [
            'driver' => 'amqp-extension',
            'exchange' => [
                'demo-rpc-server' => [
                    'name' => 'demo-rpc-server',
                    'type' => 'direct',
                    'connection' => 'default-amqp-connection',
                ],
```

(continues on next page)

(continued from previous page)

```
            ],
            'queue' => [
                'demo-rpc-server' => [
                    'name' => 'demo-rpc-server',
                    'exchanges' => [
                        'demo-rpc-server' => [],
                    ],
                    'connection' => 'default-amqp-connection',
                ],
            ],
            'connection' => [
                'default-amqp-connection' => [
                    'host' => 'localhost',
                    'port' => 5672,
                    'login' => 'guest',
                    'password' => 'guest',
                    'vhost' => '/',
                    'persistent' => true,
                    'read_timeout' => 3, //sec, float allowed
                    'write_timeout' => 1, //sec, float allowed
                ],
            ],
            'json_rpc_server' => [
                'demo-rpc-server' => [
                    'delivery_callback' => 'timestwo',
                    'idle_timeout' => 10,
                    'queue' => 'demo-rpc-server',
                    'auto_setup_fabric' => true,
                ],
            ],
        ],
    ],
];
```

So what's important here? The JSON RPC Server needs an exchange and a queue. All messages routed to the exchange, will get routed to the server queue. In this example we use a direct exchange and a single queue for the server. This is pretty much the simplest setup we can have.

The second this is, the server needs a callback, we use *timestwo* here. As you can see in the dependencies setup, the callback is simply turning a Request into a Response like this:

```
function (\Humus\Amqp\JsonRpc\Request $request): \Humus\Amqp\JsonRpc\Response {
    return \Humus\Amqp\JsonRpc\JsonRpcResponse::withResult($request->id(), $request->
↪params() * 2);
}
```

For the callback function consider this:

- Return an instance of *HumusAmqpJsonRpcResponse*
- Just return some value, the server will automatically wrap the result with an instance of *HumusAmqpJsonRpcJsonRpcResponse*
- All thrown exceptions will return an error response to the client

The *HumusAmqpJsonRpcRequest* also has a method named *method()* - this allows you to a single callback to return different results, based on the method. For example:

```
function (\Humus\Amqp\JsonRpc\Request $request): \Humus\Amqp\JsonRpc\Response {
    switch ($request->method()) {
        case 'times2':
            return \Humus\Amqp\JsonRpc\JsonRpcResponse::withResult($request->id(),
↪$request->params() * 2);
        case 'times3:
            return \Humus\Amqp\JsonRpc\JsonRpcResponse::withResult($request->id(),
↪$request->params() * 3);
        case 'plus5':
            return \Humus\Amqp\JsonRpc\JsonRpcResponse::withResult($request->id(),
↪$request->params() + 5);
        default:
            return \Humus\Amqp\JsonRpc\JsonRpcResponse::withError($request->id(), new␣
↪\Humus\Amqp\JsonRpc\JsonRpcError(32601));
    }
}
```

2.11.2 Running JSON-RPC servers

To start a JSON-RPC server

```
$ ./vendor/bin/humus-amqp json_rpc_server -n demo-rpc-server -a 100
```

This will start the *demo-rpc-server* and consume 100 messages until it stops or times out.

2.11.3 Setup the JSON RPC Client

Again, let's start with a sample configuration first (and skip the server config part, to make it easier to read):

```
return [
    'dependencies' => [
        'factories' => [
            Driver::class => Container\DriverFactory::class,
            'default-amqp-connection' => [Container\ConnectionFactory::class, 'default
↪'],
            'demo-rpc-client' => [Container\JsonRpcClientFactory::class, 'demo-rpc-
↪client'],
        ],
    ],
    'humus' => [
        'amqp' => [
            'driver' => 'amqp-extension',
            'exchange' => [
                'demo-rpc-client' => [
                    'name' => 'demo-rpc-client',
                    'type' => 'direct',
                    'connection' => 'default-amqp-connection',
                ],
            ],
            'queue' => [
                'demo-rpc-client' => [
                    'name' => '',
                    'exchanges' => [
                        'demo-rpc-client' => [],
```

(continues on next page)

(continued from previous page)

```
                    ],
                    'connection' => 'default-amqp-connection',
                ],
            ],
            'connection' => [
                'default-amqp-connection' => [
                    'host' => 'localhost',
                    'port' => 5672,
                    'login' => 'guest',
                    'password' => 'guest',
                    'vhost' => '/',
                    'persistent' => true,
                    'read_timeout' => 3, //sec, float allowed
                    'write_timeout' => 1, //sec, float allowed
                ],
            ],
            'json_rpc_client' => [
                'demo-rpc-client' => [
                    'queue' => 'demo-rpc-client',
                    'auto_setup_fabric' => true,
                    'exchanges' => [
                        'demo-rpc-server'
                    ],
                ],
            ],
        ],
    ],
];
```

So what's important here: The RPC client also needs an exchange and a queue. But the important thing to note is, that the queue has no name, an empty string is given as queue name. This will automatically create a queue with a unique name that will get destroyed, when the client is no longer in use. Also the client needs an array of exchanges, where the client can send messages to. In this example we use a single exchange *demo-rpc-server*.

2.11.4 Using the JSON RPC client

As an exercise, let's send two requests to our JSON RPC server and see what results we get:

```
$request1 = new \Humus\Amqp\JsonRpc\JsonRpcRequest('demo-rpc-server', 'timestwo', 1,
↪'request-1');
$request2 = new \Humus\Amqp\JsonRpc\JsonRpcRequest('demo-rpc-server', 'timestwo', 2,
↪'request-2');

$client->addRequest($request1);
$client->addRequest($request2);

$responses = $client->getResponseCollection();

$response1 = $responses->getResponse('request-1');
$response2 = $responses->getResponse('request-2');

var_dump($response1->isError()); // false
var_dump($response2->isError()); // false

var_dump($response1->result()); // 2
```

(continues on next page)

(continued from previous page)

```
var_dump($response2->result()); // 4
```

2.11.5 Troubleshooting

If you have read this guide and still have issues with connecting, check our *Troubleshooting guide* and feel free to raise an issue at Github.

2.11.6 What to Read Next

The documentation is organized as *a number of guides*, covering various topics.

We recommend that you read the following guides first, if possible, in this order:

- *RabbitMQ Extensions to AMQP 0.9.1*
- *Error Handling and Recovery*
- *Troubleshooting*
- *Deployment*

2.11.7 Tell Us What You Think!

Please take a moment to tell us what you think about this guide: Send an e-mail, say hello in the HumusAmqp gitter chat. or raise an issue on Github.

Let us know what was unclear or what has not been covered. Maybe you do not like the guide style or grammar or discover spelling mistakes. Reader feedback is key to making the documentation better.

2.12 RabbitMQ Extensions

HumusAmqp supports all RabbitMQ extensions to AMQP 0.9.1 that the PHP AMQP Extension supports, too:

- Negative acknowledgements (basic.nack)
- Exchange-to-Exchange Bindings
- Alternate Exchanges
- Per-queue Message Time-to-Live
- Per-message Time-to-Live
- Queue Leases
- Sender-selected Distribution
- Dead Letter Exchanges
- Publisher confirms
- Validated user_id

This guide briefly describes how to use these extensions with HumusAmqp.

2.12.1 Enabling RabbitMQ Extensions

You don't need to require any additional files to make HumusAmqp support RabbitMQ extensions. The support is built into the core.

2.12.2 Per-queue Message Time-to-Live

Per-queue Message Time-to-Live (TTL) is a RabbitMQ extension to AMQP 0.9.1 that allows developers to control how long a message published to a queue can live before it is discarded. A message that has been in the queue for longer than the configured TTL is said to be dead. Dead messages will not be delivered to consumers and cannot be fetched.

```
<?php

return [
    'dependencies' => [
        'factories' => [
            Driver::class => Humus\Amqp\Container\DriverFactory::class,
            'default-amqp-connection' =>␣
↪[Humus\Amqp\Container\ConnectionFactory::class, 'default'],
        ],
    ],
    'humus' => [
        'amqp' => [
            'driver' => 'amqp-extensions',
            'connection' => [
                'default' => [
                    'host' => 'localhost',
                    'port' => 5672,
                    'login' => 'guest',
                    'password' => 'guest',
                ],
            ],
            'exchange' => [
                'demo-exchange' => [
                    'name' => 'demo-exchange',
                    'type' => 'direct',
                    'connection' => 'default-amqp-connection',
                ],
            ],
            'queue' => [
                'my-queue' => [
                    'name' => 'my-queue',
                    'exchanges' => [
                        'demo-exchange' => [
                            [
                                'arguments' => [
                                    'x-message-ttl' => 1000
                                ],
                            ],
                        ],
                    ],
                    'connection' => 'default-amqp-connection',
                ],
            ],
        ],
```

(continues on next page)

(continued from previous page)

```
    ],
];
```

When a published message is routed to multiple queues, each of the queues gets a *copy of the message*. If the message subsequently dies in one of the queues, it has no effect on copies of the message in other queues.

Learn More

See also rabbitmq.com section on Per-queue Message TTL

2.12.3 basic.nack

The AMQP 0.9.1 specification defines the basic.reject method that allows clients to reject individual, delivered messages, instructing the broker to either discard them or requeue them. Unfortunately, basic.reject provides no support for negatively acknowledging messages in bulk.

To solve this, RabbitMQ supports the basic.nack method that provides all of the functionality of basic.reject whilst also allowing for bulk processing of messages.

How To Use It With HumusAmqp

The HumusAmqp makes already use of the nack method to reject a block of messages. You don't need to take care of that, unless you want to write a consumer yourself.

Learn More

See also rabbitmq.com section on basic.nack

2.12.4 Alternate Exchanges

The Alternate Exchanges RabbitMQ extension to AMQP 0.9.1 allows developers to define "fallback" exchanges where unroutable messages will be sent.

```
<?php

return [
    'dependencies' => [
        'factories' => [
            Driver::class => Humus\Amqp\Container\DriverFactory::class,
            'default-amqp-connection' =>␣
↪[Humus\Amqp\Container\ConnectionFactory::class, 'default'],
        ],
    ],
    'humus' => [
        'amqp' => [
            'driver' => 'amqp-extensions',
            'connection' => [
                'default' => [
                    'host' => 'localhost',
                    'port' => 5672,
```

(continues on next page)

(continued from previous page)

```
                    'login' => 'guest',
                    'password' => 'guest',
                ],
            ],
            'exchanges' => [
                'demo-exchange' => [
                    'name' => 'demo-exchange',
                    'type' => 'direct',
                    'arguments' => [
                        'alternate_exchange' => 'alternate-exchange-name'
                    ],
                    'connection' => 'default-amqp-connection',
                ],
            ],
        ],
    ],
];
```

Learn More

See also rabbitmq.com section on Alternate Exchanges

2.12.5 Exchange-To-Exchange Bindings

RabbitMQ supports exchange-to-exchange bindings to allow even richer routing topologies as well as a backbone for some other features (e.g. tracing).

```
<?php

return [
    'dependencies' => [
        'factories' => [
            Driver::class => Humus\Amqp\Container\DriverFactory::class,
            'default-amqp-connection' =>␣
↪[Humus\Amqp\Container\ConnectionFactory::class, 'default'],
        ],
    ],
    'humus' => [
        'amqp' => [
            'driver' => 'amqp-extensions',
            'connection' => [
                'default' => [
                    'host' => 'localhost',
                    'port' => 5672,
                    'login' => 'guest',
                    'password' => 'guest',
                ],
            ],
            'exchange' => [
                'exchange1' => [
                    'name' => 'exchange1',
                    'type' => 'direct',
                    'connection' => 'default-amqp-connection',
                ],
```

(continues on next page)

(continued from previous page)

```
                'exchange2' => [
                    'name' => 'exchange2',
                    'type' => 'direct',
                    'connection' => 'default-amqp-connection',
                ],
                'demo-exchange' => [
                    'name' => 'demo-exchange',
                    'type' => 'direct',
                    'exchange_bindings' => [
                        'exchange1' => [
                            [
                                'routing_keys' => [
                                    'routingKey.1',
                                    'routingKey.2'
                                ],
                            ],
                        ),
                        'exchange2' => [
                            [
                                'routing_keys' => [
                                    'routingKey.3'
                                ],
                            ],
                        ],
                    ],
                    'connection' => 'default-amqp-connection',
                ],
            ],
        ],
    ],
];
```

Learn More

See also rabbitmq.com section on Exchange-to-Exchange Bindings

2.12.6 Queue Leases

Queue Leases is a RabbitMQ feature that lets you set for how long a queue is allowed to be *unused*. After that moment, it will be deleted. *Unused* here means that the queue

- has no consumers
- is not redeclared
- no message fetches happened

```
<?php

return [
    'dependencies' => [
        'factories' => [
            Driver::class => Humus\Amqp\Container\DriverFactory::class,
            'default-amqp-connection' =>␣
↪[Humus\Amqp\Container\ConnectionFactory::class, 'default'],
```

(continues on next page)

(continued from previous page)

```
        ],
    ],
    'humus' => [
        'amqp' => [
            'driver' => 'amqp-extensions',
            'connection' => [
                'default' => [
                    'host' => 'localhost',
                    'port' => 5672,
                    'login' => 'guest',
                    'password' => 'guest',
                ],
            ],
            'exchange' => [
                'demo-exchange' => [
                    'name' => 'demo-exchange',
                    'type' => 'direct',
                    'arguments' => [
                        'x-expires' => 10000
                    ],
                    'connection' => 'default-amqp-connection',
                ],
            ],
        ],
    ],
];
```

Learn More

See also rabbitmq.com section on Queue Leases

2.12.7 Per-Message Time-to-Live

A TTL can be specified on a per-message basis, by setting the `:expiration` property when publishing.

```
<?php

$attribs = new MessageAttributes()
$attribs->setExpiration(5000);

$producer->publish('some message', '', $attribs);
```

Learn More

See also rabbitmq.com section on Per-message TTL

2.12.8 Sender-Selected Distribution

Generally, the RabbitMQ model assumes that the broker will do the routing work. At times, however, it is useful for routing to happen in the publisher application. Sender-Selected Routing is a RabbitMQ feature that lets clients have extra control over routing.

The values associated with the "CC" and "BCC" header keys will be added to the routing key if they are present. If neither of those headers is present, this extension has no effect.

```
<?php

$producer->publish('some message', '', Constants::AMQP_NOPARAM, [
    'headers' => [
        'CC' => [
            'two',
            'three'
        ],
    ],
]);
```

Learn More

See also rabbitmq.com section on Sender-Selected Distribution

2.12.9 Dead Letter Exchange (DLX)

The x-dead-letter-exchange argument to queue.declare controls the exchange to which messages from that queue are 'dead-lettered'. A message is dead-lettered when any of the following events occur:

The message is rejected (basic.reject or basic.nack) with requeue=false; or the TTL for the message expires.

How To Use It With HumusAmqp

Dead-letter Exchange is a feature that is used by specifying additional queue arguments:

- "x-dead-letter-exchange" specifies the exchange that dead lettered messages should be published to by RabbitMQ
- "x-dead-letter-routing-key" specifies the routing key that should be used (has to be a constant value)

```
<?php

return [
    'dependencies' => [
        'factories' => [
            Driver::class => Humus\Amqp\Container\DriverFactory::class,
            'default-amqp-connection' =>␣
↪[Humus\Amqp\Container\ConnectionFactory::class, 'default'],
        ],
    ],
    'humus' => [
        'amqp' => [
            'driver' => 'amqp-extensions',
            'connection' => [
                'default' => [
                    'host' => 'localhost',
                    'port' => 5672,
                    'login' => 'guest',
                    'password' => 'guest',
                ],
            ],
```

(continues on next page)

(continued from previous page)

```
            'queue' => [
                'foo' => [
                    'name' => 'foo',
                    'exchanges' => [
                        'demo' => [],
                    ],
                    'arguments' => [
                        'x-dead-letter-exchange' => 'demo.error'
                    ],
                    'connection' => 'default-amqp-connection',
                ],
            ],
        ],
    ],
];
```

Learn More

See also rabbitmq.com section on Dead Letter Exchange

2.12.10 Wrapping Up

RabbitMQ provides a number of useful extensions to the AMQP 0.9.1 specification.

HumusAmqp releases have RabbitMQ extensions support built into the core. Some features are based on optional arguments for queues, exchanges or messages, and some are HumusAmqp public API features. Any future argument-based extensions are likely to be useful with HumusAmqp immediately, without any library modifications.

2.12.11 What to Read Next

The documentation is organized as *a number of guides*, covering various topics.

We recommend that you read the following guides first, if possible, in this order:

- *RabbitMQ Extensions to AMQP 0.9.1*
- *Error Handling and Recovery*
- *Troubleshooting*
- *Deployment*

2.12.12 Tell Us What You Think!

Please take a moment to tell us what you think about this guide: Send an e-mail, say hello in the HumusAmqp gitter chat. or raise an issue on Github.

Let us know what was unclear or what has not been covered. Maybe you do not like the guide style or grammar or discover spelling mistakes. Reader feedback is key to making the documentation better.

2.13 Error Handling

2.13.1 Client Exceptions

Here is the break-down of exceptions that can be raised by HumusAmqp

```
Humus\Amqp\Exception\BadMethodCallException
Humus\Amqp\Exception\ChannelException
Humus\Amqp\Exception\ConnectionException
Humus\Amqp\Exception\InvalidArgumentException
Humus\Amqp\Exception\InvalidJsonRpcRequest
Humus\Amqp\Exception\InvalidJsonRpcVersion
Humus\Amqp\Exception\JsonParseError
Humus\Amqp\Exception\QueueException
Humus\Amqp\Exception\RuntimeException
```

2.13.2 Initial RabbitMQ Connection Failures

When applications connect to the broker, they need to handle connection failures. Networks are not 100% reliable, even with modern system configuration tools like Chef or Puppet misconfigurations happen and the broker might also be down. Error detection should happen as early as possible. To handle TCP connection failure, catch the `Humus\Amqp\Exception\ConnectionException` exception.

2.13.3 Authentication Failures

Another reason why a connection may fail is authentication failure. Handling authentication failure is very similar to handling initial TCP connection failure.

When you try to access RabbitMQ with invalid credentials, you'll get an `\Humus\Amqp\Exception\ConnectionException` with message `Library error: a socket error occurred - Potential login failure.`.

In case you are wondering why the exception name has "potential" in it: AMQP 0.9.1 spec requires broker implementations to simply close TCP connection without sending any more data when an exception (such as authentication failure) occurs before AMQP connection is open. In practice, however, when broker closes TCP connection between successful TCP connection and before AMQP connection is open, it means that authentication has failed.

2.13.4 Channel-level Exceptions

Channel-level exceptions are more common than connection-level ones and often indicate issues applications can recover from (such as consuming from or trying to delete a queue that does not exist).

Common channel-level exceptions and what they mean

A few channel-level exceptions are common and deserve more attention.

406 Precondition Failed

Description

The client requested a method that was not allowed because some precondition failed.

What might cause it

AMQP entity (a queue or exchange) was re-declared with attributes different from original declaration. Maybe two applications or pieces of code declare the same entity with different attributes. Note that different RabbitMQ client libraries historically use slightly different defaults for entities and this may cause attribute mismatches.

PRECONDITION_FAILED - parameters for queue 'examples.channel_exception' in vhost '/' not equivalent

PRECONDITION_FAILED - channel is not transactional

405 Resource Locked

Description

The client attempted to work with a server entity to which it has no access because another client is working with it.

What might cause it

Multiple applications (or different pieces of code/threads/processes/routines within a single application) might try to declare queues with the same name as exclusive.

Multiple consumer across multiple or single app might be registered as exclusive for the same queue.

Example RabbitMQ error message

RESOURCE_LOCKED - cannot obtain exclusive access to locked queue 'examples.queue' in vhost '/'

404 Not Found

Description

The client attempted to use (publish to, delete, etc) an entity (exchange, queue) that does not exist.

What might cause it

Application miscalculates queue or exchange name or tries to use an entity that was deleted earlier

Example RabbitMQ error message

NOT_FOUND - no queue 'queue_that_should_not_exist0.6798199937619038' in vhost '/'

403 Access Refused

Description

The client attempted to work with a server entity to which it has no access due to security settings.

What might cause it

Application tries to access a queue or exchange it has no permissions for (or right kind of permissions, for example, write permissions)

Example RabbitMQ error message

ACCESS_REFUSED - access to queue 'examples.channel_exception' in vhost '_testbed' refused for user '_reader'

2.13.5 What to Read Next

The documentation is organized as *a number of guides*, covering various topics.

We recommend that you read the following guides first, if possible, in this order:

- *Troubleshooting*
- *Deployment*

2.13.6 Tell Us What You Think!

Please take a moment to tell us what you think about this guide: Send an e-mail, say hello in the HumusAmqp gitter chat. or raise an issue on Github.

Let us know what was unclear or what has not been covered. Maybe you do not like the guide style or grammar or discover spelling mistakes. Reader feedback is key to making the documentation better.

2.14 Troubleshooting

2.14.1 First steps

Whenever something doesn't work, check the following things before asking on the mailing list:

- RabbitMQ log.
- List of users in a particular vhost you are trying to connect.
- Network connectivity, firewall settings, DNS host resolution.

2.14.2 Inspecting RabbitMQ log file

In this section we will cover typical problems that can be tracked down by reading RabbitMQ log.

RabbitMQ logs abrupt TCP connection failures, timeouts, protocol version mismatches and so on. If you are running RabbitMQ, log file location depends on the operating systems and installation method. See RabbitMQ installation guide for more information.

OS X with Homebrew

On Mac OS X, RabbitMQ installed via Homebrew logs to $HOMEBREW_HOME/var/log/rabbitmq/rabbit@$HOSTNAME.log. For example, if you have Homebrew installed at `/usr/local` and your hostname is giove, the log will be at `/usr/local/var/log/rabbitmq/rabbit@giove.log`.

Authentication Failures

Here is what authentication failure looks like in a RabbitMQ log:

```
=ERROR REPORT==== 12-Jul-2013::16:49:03 ===
closing AMQP connection <0.31567.1> (127.0.0.1:50458 -> 127.0.0.1:5672):
{handshake_error,starting,0,
                 {amqp_error,access_refused,
```

(continues on next page)

(continued from previous page)

```
                                "PLAIN login refused: user 'pipeline_agent' - invalid␣
↪credentials",
                                'connection.start_ok'}}
```

This means that the connection attempt with the username pipeline_agent failed because the credentials were invalid. If you are seeing this message, make sure username, password and vhost are correct.

The following entry:

```
=ERROR REPORT==== 17-May-2011::17:26:28 ===
exception on TCP connection <0.4201.62> from 10.8.0.30:57990
{bad_header,<<65,77,81,80,0,0,9,1>>}
```

means that an old RabbitMQ version (pre-2.0) is used. Those versions are not supported by HumusAmqp. It is recommended to use the latest stable release.

2.14.3 Handling Channel-level Exceptions

A broad range of problems result in AMQP channel exceptions: an indication by the broker that there was an issue that the application needs to be aware of. Channel-level exceptions are typically not fatal and can be recovered from. Some examples are:

- Exchange is re-declared with attributes different from the original declaration. For example, a non-durable exchange is being re-declared as durable.
- Queue is re-declared with attributes different from the original declaration. For example, an auto-deletable queue is being re-declared as non-auto-deletable.
- Queue is bound to an exchange that does not exist.

and so on. These will result in a reasonably descriptive exception `Humus\Amqp\Exception\ChannelException`. Handling and logging them will likely reveal an issue when it arises.

2.14.4 Network connection issues

Testing Network Connection with RabbitMQ using Telnet

One simple way to check network connection between a particular network node and a RabbitMQ node is to use telnet:

```
telnet [host or ip] 5672
```

then enter any random string of text and hit Enter. RabbitMQ should immediately close down the connection. Here is an example session:

```
telnet localhost 5672
Connected to localhost.
Escape character is '^]'.
adjasd
AMQP    Connection closed by foreign host.
```

If Telnet exits after printing instead

```
telnet: connect to address [host or ip]: Connection refused
telnet: Unable to connect to remote host
```

then the connection between the machine that you are running Telnet tests on and RabbitMQ fails. This can be due to many different reasons, but it is a good idea to check these two things first:

- Firewall configuration for port 5672 or 5671 (if TLS/SSL is used)
- DNS resolution (if hostname is used)

Connecting to localhost on VPN

Using VPN almost certainly changes your DNS server configuration which may affect connections to `localhost` as well as to remote hosts. If you keep getting `Got an exception when receiving data: IO timeout when reading 7 bytes (Timeout::Error)` errors and you're on VPN try switching VPN off.

2.14.5 RabbitMQ Startup Issues

Missing erlang-os-mon on Debian and Ubuntu

The following error on RabbitMQ startup on Debian or Ubuntu

```
ERROR: failed to load application os_mon: {"no such file or directory","os_mon.app"}
```

suggests that the erlang-os-mon package is not installed.

2.14.6 asn1 Issue with Erlang R16B01

```
BOOT FAILED
===========

Error description:
    {error,{cannot_start_application,public_key,{not_started,asn1}}}
```

is an issue in RabbitMQ 3.1 on Erlang R16B01+. It is resolved in RabbitMQ 3.1.2 and later versions.

2.14.7 What to Read Next

The documentation is organized as *a number of guides*, covering various topics.

We recommend that you read the following guides first, if possible, in this order:

- *Deployment*

2.14.8 Tell Us What You Think!

Please take a moment to tell us what you think about this guide: Send an e-mail, say hello in the HumusAmqp gitter chat. or raise an issue on Github.

Let us know what was unclear or what has not been covered. Maybe you do not like the guide style or grammar or discover spelling mistakes. Reader feedback is key to making the documentation better.

2.15 Deployment Strategies

2.15.1 Shut down the system, update and restart

While the easiest way to deploy a RabbitMQ environment is to just destroy the old one and create a new one, that's not always possible. Sometimes the consumers are not part of your application or it's complete unacceptable to put the payment system offline and not to process messages even for half an hour, things get more complicated.

This guide tries to show some different deployment strategies. Note, that there is no general way to do this, it will always depend on the use-case you have.

2.15.2 Create a new node and switch configuration

Let's say you have a running RabbitMQ configuration on a given node (or vhost). An easy way to update configuration would be to just deploy a new node (or vhost) and switch you application configuration to use that new node (vhost) instead of the old one.

2.15.3 Message-Versioning

Message-Versioning with Routing Keys

You could use the routing keys for versioning, like this:

```
<?php

$producer->publish('some message', 'v1.0.0');
```

Then you just bind the queue to the routing key. When you update your system, newer messages (e.g. 2.0.0) are not delivered to queue that is not able to process them. On the other hand you can write consumers, that are backwards-compatible, so the queues used will bind to a variety of routing keys (v1.0.0 & v2.0.0 or v1.0.*).

Note: Use semantic versioning.

Message-Versioning with extra Attributes

Use some custom message headers to specify version constraints:

```
<?php

$producer->publish('some message', '', [
    'headers' => [
        'x-version' => 'v1.0.0',
    ],
]);
```

This way a message with a version not handled by your consumer, will also be at least delivered to him. The consumer will then need to check for the x-version header and process if possible or throw an exception. Compared to the versioning strategy with routing keys, you'll always know, that somebody is sending still messages in old version, as you see the exceptions in the log file. If you just deploy a consumer that is only able to acknowlegde messages of e.g. v2.0.0, than it's hard to recognize that there are still messages in version 1.0.0 going through the system. On the other hand, there are messages send over network, that nobody cares, so routing keys are often preferable.

2.15.4 Updating Exchanges

Most times it's not neccessary to update an exchange configuration. However if required, you have to remove the old exchange before you can redeclare the new one with the same name. Keep in mind, that when you put an exchange down, no messages can be delivered any more. It's also possible to use an exchange name like: "update-stuff-v1-0-0", so you don't have to deleted the old exchange when deploying the new one.

2.15.5 Updating Queues

Updating queues can be sometimes a little more tricky then updating an exchange. If you bind the new queue before the old one is destroyed, you'll get duplicated messages in your system. But if this is done in concert with a new exchange you can prevent duplicate messages.

2.15.6 Getting queues empty first

Sometimes you might want to upgrade the consumers because of a new message format and you don't want to maintain backwards compatibility. If you don't want to lose any messages, stop all producers first, until the queues are empty, then you can to the switch without losing messages.

2.15.7 Tell Us What You Think!

Please take a moment to tell us what you think about this guide: Send an e-mail, say hello in the HumusAmqp gitter chat. or raise an issue on Github.

Let us know what was unclear or what has not been covered. Maybe you do not like the guide style or grammar or discover spelling mistakes. Reader feedback is key to making the documentation better.

2.16 License Information

CHAPTER 3

Indices and tables

- search