

# OOI-CI Prototype Exchange-Space Implementation Report

Tony Garnock-Jones <tonyg@lshift.net>


5th January 2010

## Abstract

Final summary of the prototype phase of the OOI-CI RabbitMQ integration project, listing deliverables, design alternatives, and open issues.

## 1 Introduction and background

This document reports on the outcome of the RabbitMQ part of the prototype OOI-CI AMQP messaging work. It should be read as a companion to the following documents:

- The RabbitMQ OOI Statement of Work;
- The OOI-CI Prototype Exchange-Space API documentation 
- General OOI-CI context documents on exchange-spaces, exchange-points and federation.

There are three main deliverables: a preconfigured RabbitMQ EC2 AMI; support for AMQP 0-9-1 in the RabbitMQ server and selected client libraries; and a prototype implementation of exchange-spaces and exchange-points, along with documentation and examples. Each is briefly described below, along with a summary of its status and pointers to further information.

## 2 RabbitMQ EC2 images

A collection of RabbitMQ AMIs is now available for use. Links and instructions are available at <http://www.rabbitmq.com/ec2.html>. The maintainer of the AMIs is Marek Majkowski (majek@lshift.net).

- Both 32- and 64-bit images are available.
- Images are based on Ubuntu 9.10.
- RabbitMQ v1.7.0 is preinstalled.

## 3 AMQP 0-9-1 interoperability

RabbitMQ's support for the 0-9-1 variant of the AMQP protocol is now largely complete. Three-way interoperability between RabbitMQ, Qpid, and OpenAMQ has been demonstrated through the use of standard off-the-shelf examples shipped with each system, and also through the 0-9-1 test suite developed by Michael Bridgen (mikeb@lshift.net). The status of our interoperability testing is kept up-to-date at <https://dev.rabbitmq.com/wiki/Amqp091Interoperability>.

Support for 0-9-1 is available on the `amqp_0_9_1` branches of the various RabbitMQ mercurial repositories. The following shell commands check out and build an 0-9-1-compatible broker, C client, and Java client:

```
hg clone http://hg.rabbitmq.com/rabbitmq-codegen
hg clone http://hg.rabbitmq.com/rabbitmq-server
hg clone http://hg.rabbitmq.com/rabbitmq-c
hg clone http://hg.rabbitmq.com/rabbitmq-java-client
(cd rabbitmq-codegen; hg up amqp_0_9_1)
(cd rabbitmq-server; hg up amqp_0_9_1)
(cd rabbitmq-c; hg up amqp_0_9_1)
(cd rabbitmq-java-client; hg up amqp_0_9_1)
make -C rabbitmq-server
(cd rabbitmq-c; autoreconf -i; ./configure; make)
(cd rabbitmq-java-client; ant dist)
```

At this point, you can start the server,

```
make -C rabbitmq-server run
```

use the java client library in `build/dist`,

```
(cd rabbitmq-java-client/build/dist; \
sh runjava.sh com.rabbitmq.examples.SendString)
```

and run the C example programs:

```
./rabbitmq-c/examples/amqp_sendstring
```

To run the 0-9-1 test suite,

```
ant test-suite-prepare test-functional
```

The results of the suite will be available in `build/TEST-*.txt`.

## 4 Exchange-space prototype implementation

The prototype takes the form of a new RabbitMQ plugin, available via git on [amoeba.ucsd.edu](http://amoeba.ucsd.edu), and a set of

supporting modifications to the RabbitMQ server. Instructions for building and running the prototype are available in the file `README.md` in the plugin source code checkout.

## 4.1 Design alternatives

Three potential alternative designs were considered:

**Layered NIPCA-style networks.** <sup>1</sup>As part of the background research for this project, the idea of unifying AMQP with NIPCA-style network operations was discussed extensively. This is still an extremely promising line of investigation, and should be revisited in future phases of the project because of its formidable descriptive power.

### Exchange-space extensions to the AMQP vocabulary.

An experimental new class, `ExchangeSpace`, was considered. It would have permitted full control over exchange-space functionality without being restricted by limitations of the existing AMQP vocabulary; however, it would have also required changes to client libraries in order for them to use the new functionality, which would have made experimenting with the new features prohibitively difficult.

**New AMQP exchange types.** This was the option chosen in the end: it helped draw out the similarities and differences between the existing AMQP vocabulary and the vocabulary required for controlling exchange-spaces, and worked seamlessly with existing AMQP client libraries. While it does not result in a perfect fit, leaving some exchange-space features inaccessible to AMQP client libraries (but controllable via out-of-band methods), it does permit flexible and rapid experimentation with the new functionality.

## 4.2 Retrieving and building the code

To retrieve the plugin,

```
git clone \
  git@amoeba.ucsd.edu:exchange_space_prototype.git
```

The changes to the RabbitMQ server are available on the branch named `bug22019`. For more information on retrieving, building, configuring and running the system, see the plugin's `README.md` file.

Example programs using `Pika`<sup>2</sup> are provided in the `examples` directory of the plugin source code checkout, but similar examples can be written for any AMQP

<sup>1</sup>Network IPC Architecture, NIPCA: a powerful network architecture described in John Day's book, *Patterns in Network Architecture*, 2008.

<sup>2</sup>A python AMQP client library. <http://github.com/tonyg/pika>

client library that can support the provision of custom exchange-type names and exchange declaration argument tables; the prototype was specifically designed to be compatible with unmodified existing AMQP client libraries. See the prototype API documentation for more information.

## 4.3 Open issues and next steps

The experimental implementation, with its new exchange-types and otherwise unmodified AMQP vocabulary, is a stop-gap on the road to full integration between AMQP and NIPCA-style concepts. A synthesis of the two systems seems likely to be able to uniformly and flexibly describe and manipulate many common configurations of networks small and large, including several that are not amenable to description in any manner within the bounds of today's network programming APIs.<sup>3</sup> Many open questions in this area remain; among others:

- What is the relationship between associations (roughly analogous to connections) and subscriptions?
- What is the relationship between associations and individual messages?
- What is the relationship between name-binding, directory, and subscription?
- What is the relationship between AMQP's exchanges and queues and NIPCA's names and addresses?
- Where do applications and application instances fit in AMQP? How do they relate to applications and application instances in NIPCA?

In terms of the experimental implementation itself, one interesting issue is that the existing AMQP 0-9-1 vocabulary is not rich enough to describe all exchange-space and exchange-point operations: in particular, it can neither describe options for managing buffers and transfer-of-responsibility within the network nor express the difference between exchange-point deletion and exchange-point unsubscription. If the current implementation is to be taken forward, either an out-of-band control channel will need to be provided for controlling such things, or the core AMQP vocabulary itself will have to be altered, perhaps in a manner guided by future work on the AMQP/NIPCA synthesis.

<sup>3</sup>"Whereof one cannot speak, thereof one must be silent." Firewalls, just to pick one important example, are not represented in the BSD sockets API.

The experimental implementation provides only the barest rudiments of a system for management and monitoring of exchange-spaces and exchange-points. Future work will need to explore the options for improving the management and monitoring situation.

Finally, federation is one of the areas being worked on as part of the development of the AMQP 1.0 specification. The design ideas explored in this project could potentially contribute to the design of standardised federation for AMQP, especially in the area of managing the scoping of addressing and naming.

## **5 Conclusion**

All the project deliverables are complete, with respect to the goals outlined in the Statement of Work and refined by subsequent email and face-to-face communication. OOI's internal prototyping work can proceed, based on the experimental exchange-space implementation provided, and clear lines of potential future investigation exist.

# OOI-CI Prototype Exchange-Space API

Tony Garnock-Jones <tonyg@lshift.net>

28th December 2009

## Abstract

Description of the prototype OOI-CI AMQP API for managing and making use of exchange-spaces and exchange-points alongside existing AMQP 0-9-1 resources, and relation of this API to other ongoing work.

## 1 Introduction

A fully general AMQP model would not only cover AMQP 1.0 and 0-9-1, but would also be an instance of a NIPCA-style DIF architecture, and would permit uniform naming and use of exchanges, resources, gateways and so forth. While we're working on such a generalisation of the AMQP model, an interim API will be useful not only for a first generation of prototype applications but also for clarifying our thinking around enrolment, DIF architecture, and generalised AMQP.

This document discusses the programming and wire-protocol interfaces to the new exchange-space and exchange-point functionality, but does not discuss details of potential implementations of the new interfaces. It also touches on the relationship of this prototyping work to the generalisation of AMQP and to the ongoing NIPCA development work.

## 2 Behavioural Definitions

The new entities have particular behaviours that need to be defined independent of the protocols used to manipulate them.

### 2.1 Exchange Space

Exchange spaces are "virtual" or "cloud" analogues of AMQP 0-9-1's virtual-hosts. Some aspects of "real" AMQP virtual-hosts are not available in exchange-spaces.<sup>1</sup> An exchange space logically

- contains exchange-points, and
- mediates access to the contained exchange-points.

<sup>1</sup>Notably, queues are not available within an exchange-space.

In effect, an exchange-space is an administrative scope for the exchange-point resources it contains. The lifecycle of the exchange-space bounds the lifecycle of the contained exchange-points: when the space is deleted, so are its contained points.

### 2.2 Exchange Point

Exchange points are "virtual" or "cloud" analogues of AMQP 0-9-1's exchanges.<sup>2</sup> An exchange point

- can be of any AMQP 0-9-1 type
- can be published to
- can have queues bound to it

Exchange points, just like regular AMQP exchanges, have lifetimes that are independent of the lifetimes of the client connections that created them. In AMQP, an exchange that is declared by a connected client persists until either it is explicitly deleted, the containing broker is restarted (in the case of non-durable exchanges), or the containing broker's configuration database or the containing broker itself is deleted (in the case of durable exchanges). Exchange-points behave similarly: they persist until explicitly deleted, or until the containing exchange-space is reset or destroyed.

## 3 API overview

The existing AMQP 0-9-1 wire protocol is used unchanged. Exchange-spaces and exchange-points within them are encoded as new exchange types in the system.

### 3.1 Declaring exchange spaces

Exchange spaces are declared<sup>3</sup> by invoking `Exchange.Declare` with an exchange type

<sup>2</sup>In this prototype implementation, exchange-points are also *modelled* as exchanges.

<sup>3</sup>"Declaration" in the AMQP sense of asserting the existence of, creating if necessary.

of `x-exchangespace`.<sup>4</sup> The name given to `Exchange.Declare` will be used as the name of the exchange space. The `arguments` table should contain a string-valued entry named `implementation`, which is used to select an implementation variant.<sup>5</sup> For the prototype, the string should have the value `prototype`. Other entries in the `arguments` table are exchange-space-implementation-specific parameters; for the prototype implementation, no parameters are needed.

### 3.2 Securing exchange spaces

The prototype implementation will have the ability to manage access to created exchange-spaces by using RabbitMQ's existing ACL-like facilities for restricting access to exchanges (as specified in the Statement of Work). While we do have some nascent ideas on controlling authorisation in a distributed setting, development of such ideas is out-of-scope for this phase of development.

The broker holding the exchange declared as type `x-exchangespace` should have users and permissions configured appropriately for restricting access to the declared exchange. Brokers wishing to join an exchange-space must be able to log into the broker on which the `x-exchangespace` exchange is declared, and must be able to both read from and write to the exchange.

### 3.3 Joining and leaving exchange spaces

Joining (i.e., becoming a participating peer in) an exchange space is done implicitly as part of the declaration of a particular exchange point. When an exchange-point is declared, the exchange-space management function is contacted by a behind-the-scenes process. No explicit client action besides declaration of the exchange-point is required.

Leaving an exchange space is done implicitly as part of exchange-point deletion; this is a point of awkwardness in the current design because while there are explicit representations of exchange-spaces and exchange-points, there are no representations of the behind-the-scenes relay processes. Not having a direct handle on relay processes means that all changes to their configuration happen as side-effects of actions on other objects.<sup>6</sup>

<sup>4</sup>Non-standard exchange types are required by the AMQP 0-9-1 specification to have names starting with `x-`.

<sup>5</sup>This is an instance of the *strategy pattern*.

<sup>6</sup>Two possible approaches to avoiding these problems occur to me: either reify the relay processes themselves, or move joining and leaving to being actions that occur on bind and unbind. Neither approach is entirely satisfactory. For the prototype phase of work, we can leave the situation unresolved and concentrate on the main aspects of the system.

### 3.4 Declaring exchange points

Exchange points are declared by invoking `Exchange.Declare` with an exchange type of `x-exchange`. The name given to `Exchange.Declare` will be used as a *local* alias for the exchange point, and does not need to correspond to the name of the exchange point itself. An exchange-point then, from the declaring client's (and the client's local broker's!) point-of-view, has *two* names: the local alias, for use only with the local broker, as well as its "true name", the combination of its exchange-space's name and the exchange-point's unique name within the space.

The `arguments` table given to `Exchange.Declare` must contain the following entries:

- a table-valued entry named `exchangespace`, with contents specific to the implementation variant used but containing at least a string-valued entry named `implementation`, which is used to select an implementation variant. For the prototype, the string should have the value `prototype`.
- a string-valued entry named `exchangename` containing an AMQP exchange name such as might be supplied in the `name` argument to `Exchange.Declare`. This is the name of the exchange-point itself, within the remote exchange-space.
- a string-valued entry named `exchangetype` containing an AMQP exchange type name such as might be supplied in the `exchange type` argument to `Exchange.Declare`.

The `arguments` table may also contain:

- a table-valued entry named `arguments` containing a nested set of exchange declaration arguments such as might be supplied in the `arguments` argument to `Exchange.Declare`. If this entry is omitted, an empty table is assumed.

Taken together, the `exchangename`, `exchangetype` and `arguments` fields in the `arguments` table given to `Exchange.Declare` for the exchange point are sufficient to specify the kind of exchange point being declared.

If the exchange-point declaration fails, including implementation-internal failures such as problems contacting or negotiating with the exchange-space management function, the usual AMQP error reporting procedures are used to notify the declaring client of the problem. All of AMQP's existing restrictions on exchange declarations—that the exchange types must match, and so on—apply to exchange points, as well; that is, if one client declares an exchange-point with type "fanout",

and another declares one with the same name but with type “direct”, the second declaration will fail.

### 3.5 Deleting exchange points and exchange spaces

To delete an exchange point or exchange space, invoke `Exchange.Delete` as usual.

An awkwardness exists here when deleting an exchange point: in AMQP, exchanges can be deleted without first declaring them, but here to delete an exchange point the broker must first know which exchange space to delete the point from, so the point must first have been declared on the local broker.

This is a syntactic problem with the AMQP protocol as it stands. If this prototype design were to be taken to production, the awkwardness would need to be worked around by some out-of-band, non-AMQP resource management tool, perhaps agent- or GUI-based, that could manipulate structures internal to the prototype implementation without being restricted to the existing AMQP operations.

A fully generalised model of AMQP will not have this awkwardness: resource deletion will include enough information to precisely locate the resource to delete without having to construct it beforehand.

### 3.6 Binding to and unbinding from exchange points

A similar awkwardness exists here as described for deleting exchange points: they *must* be declared on the local broker before any kind of binding is attempted. Given this, however, binding to exchange points can be done as usual for AMQP: issue `Queue.Bind` and `Queue.Unbind`, mentioning the local alias for the exchange point that was set up in the `Exchange.Declare`.

### 3.7 Publishing to exchange points

Messages sent to locally-declared `x-exchange` exchanges using `Basic.Publish` will be delivered to other participants in the exchange-space that have subscribed to the exchange-point, with implementation-specific QoS and reliability guarantees. See section 4.6 for details on QoS levels provided by the prototype implementation.

The mandatory and immediate flags to `Basic.Publish` are not supported in the prototype implementation, and must be set `false`.<sup>7</sup> A channel error will be signalled if either is set `true`. Similarly,

---

<sup>7</sup>The reason for this is that it is not clear what they mean in a distributed setting.

Tx-class transactions are not supported in the prototype; if a `Basic.Publish` is issued to an `x-exchange` exchange over a channel on which `Tx.Select` has been called, a channel error will be signalled.

## 4 Prototype implementation

The prototype implementation will provide a single exchange-space implementation variant, `prototype`. Coordination between participants in an exchange-space will be centralised (at the broker on which the exchange-space exchange was declared), analogous to the centralised tracker in early variants of Bittorrent, but the actual exchange of data between exchange-points will be decentralised, managed directly by the exchange-point participants.

This API is flexible enough to accommodate alternative implementations of the exchange-space management function, requiring changes only to the specific arguments passed in the `Exchange.Declare` command used to create an exchange-space. For example, if a decentralised exchange-space network management function were implemented, the parameters to `Exchange.Declare` would include “seed” addresses that would be used to bootstrap the network into existence. An agent-based implementation would provide specifics of the agents that were involved in the network management function.

### 4.1 Exchange-point declaration parameters

When declaring an exchange-point, the table-valued `exchangespace` field of the arguments table of the `Exchange.Declare` command must contain the following fields in addition to the mandatory `implementation` entry:

- `host`, a string: the hostname of the AMQP broker on which the exchange-space is declared.
- `space`, a string: the name given to `Exchange.Declare` when the *exchange-space* was declared.

It may also contain:

- `port`, an integer: the port number of the AMQP broker on which the exchange-space is declared. If omitted, 5672 is used.
- `username`, a string: the user name to use when connecting to the AMQP broker on which the exchange-space is declared. If omitted, `guest` is used.

- `password`, a string: the password to use when connecting to the AMQP broker on which the exchange-space is declared. Optional unless `username` is provided; if omitted, `guest` is used.
- `virtualhost`, a string: the virtual host to use when connecting to the AMQP broker on which the exchange-space is declared. If omitted, `/` is used.

The synchronisation technique and protocol used to implement exchange-points for the prototype has yet to be determined, though it is likely to be similar to the experiments embodied in [http://github.com/tonyg/pika/blob/master/examples/demo\\_relay.py](http://github.com/tonyg/pika/blob/master/examples/demo_relay.py). (It's also worth noting that the particular technique chosen is not visible at the API.)

## 4.2 Potential implementation alternatives

This API definition makes use of the strategy pattern: alternative implementations following the defined interface can be selected by varying the `implementation` arguments given to `Exchange.Declare` commands.

In order to get a feel for the space of possible implementations, it's useful to consider an analogy between this system and the Bittorrent synchronisation system. In Bittorrent, each active torrent consists of exactly one tracker, zero or more peers, and exactly one finitely-bounded set of information to be synchronised (the dataset being distributed). One could say that the tracker is logically identified with the torrent.

In the exchange-point/exchange-space system, each exchange space consists of exactly one network management function, zero or more peers, and zero or more unbounded streams of information to be synchronised (the exchange points themselves). One could say that the network management function is logically identified with the exchange-space.

Multiple implementation variants for Bittorrent are possible: the tracker can be centralised or decentralised; peers can employ different transports and protocols for performing the synchronisation task ahead of them. Similarly, multiple implementation variants for exchange-spaces are possible: the network management function can be centralised or decentralised, and peers can coordinate in many different ways to synchronise their views of the information flowing through the exchange points contained within the exchange space. A decentralised network management function would be analogous to the Distributed Hash Table used to manage Bittorrent trackers in a decentralised way.

## 4.3 Network substrate independence

One of the purposes of this prototype is to integrate with the DIF implementation being developed in parallel with it. As it stands, our implementation of this API uses AMQP over TCP to communicate from peer to peer and from peer to exchange-space network manager. When the DIF implementation becomes available, all that will be required will be formalisation of the means of specifying the addresses involved: again, the strategy pattern instance inherent in the use of the `arguments` table in `Exchange.Declare` commands gives us a convenient location to place the required syntax.

## 4.4 Broker independence and AMQP 1.0

Both the design presented here and the concrete implementation being developed are, to some extent, broker-independent. While we do take advantage of some special features of RabbitMQ for ease and speed of development, it would be possible to operate in an entirely implementation-neutral fashion (at a cost in robustness, complexity and efficiency) by changing the design, reifying implementation structures such as exchange-spaces, exchange-points, and exchange-point relays, turning them into ordinary AMQP applications.

Mapping the design onto AMQP 1.0 has yet to be discussed: it is unclear whether it will involve the kind of alteration discussed in the previous paragraph, or whether a more natural embedding will be possible.

## 4.5 Relationship to generalised AMQP work

The generalisation of AMQP mentioned in the introduction is still being developed, in parallel with this prototyping work and the aforementioned work on a generalised network transport interface. This prototyping work will help clarify some issues in the generalisation work, particularly those surrounding enrolment and the roles of queues and of relays.

## 4.6 Confirmation of receipt

The prototype implementation will deliver on a best-effort basis. Failure of a network element<sup>8</sup> on the path between two peer brokers could lead to message loss. Therefore, application-level acknowledgements, timeouts, and retransmissions will be required in situations where application semantics require confirmation of delivery.

<sup>8</sup>For example: netsplits, routing misconfiguration, and other connectivity failures; firewall reboots; broker restarts; unexpected TCP resets; and so on.

For example, let us imagine a system where an exchange-point is used to relay job requests to some remote service. The service instance receiving the requests should send an acknowledgement of receipt to the originator of each request. If the jobs complete quickly enough, the acknowledgment and the reply (if there is one) can be bundled together in a single message; if the jobs are long-running, a separate acknowledgement message should be sent when a job is received from the broker, with a reply message following later once the job completes. Furthermore, on receipt of a duplicate request, a duplicate acknowledgement<sup>9</sup> should be sent (without actually running the job again). If the originator of a request does not receive an acknowledgement in a reasonable amount of time, it should retransmit the request. This way, robust confirmation of delivery can be achieved, and a wide variety of responsibility-transfer schemes can be straightforwardly implemented without depending on the fine implementation detail of the exchange-point system.

## 5 Example

In this example,

- an exchange-space, `market`, is created by user Margaret on her AMQP server at `amqp.market.example`;
- user Alice, who runs an AMQP server at `amqp.acme.example`, runs a program that wishes to subscribe to and publish to the exchange-points `bids` and `mailbox` within `market`;
- user Bob, who runs an AMQP server at `amqp.brandcorp.example`, runs a program that also wishes to subscribe to and publish to the exchange-points `bids` and `mailbox` within `market`;
- the exchange-point `bids` is a topic exchange;
- the exchange-point `mailbox` is a direct exchange;<sup>10</sup>
- we assume that all machines have prearranged direct TCP-level connectivity available.

The pseudocode snippets given below are non-normative: they are simply illustrative examples of possible commands that could be issued in a scenario such as the one we are examining in this section.

<sup>9</sup>And a duplicate reply, if a reply is available at the time.

<sup>10</sup>One interesting benefit of direct exchanges is that it is not possible to make a wildcard binding to the exchange: if two parties are communicating via a direct exchange with a secret routing key, then all other things being secured there is no means by which an eavesdropper can overhear their conversation other than guessing the routing key being used. (Interestingly, it doesn't look like this kind of 'security' is possible in AMQP 1.0.)

## The exchange space is created

Margaret issues the following command to her server at `amqp.market.example`:

```
Exchange.Declare(exchange = 'market',
                  arguments =
                    {'implementation':
                     'prototype'},
                  type = 'x-exchangespace')
```

This command could fail for all the usual AMQP reasons: duplicate name with differing type, permission denied, unsupported exchange type, and so on.

## Alice tries to join the space, but fails

Alice issues the following command to her server at `amqp.acme.example`:

```
Exchange.Declare(exchange = 'marketbids',
                  type = 'x-exchange-point',
                  arguments =
                    {'exchangespace':
                     {'implementation': 'prototype',
                      'host': 'amqp.market.example',
                      'space': 'market',
                      'username': 'alice',
                      'password': '41iC3'},
                     'exchangenname': 'bids',
                     'exchangetype': 'topic'})
```

Alice's server instantiates an exchange-space relay, behind the scenes, and causes it to try to contact the exchange-space denoted by the `exchangespace` argument to the exchange declaration. This contact fails, ultimately causing the command that Alice issued to fail, because Margaret's server has not been configured to permit a user named `alice` to log in and access the exchange-space embodied by the exchange called `market`.

The command could also have failed for all the other usual AMQP exchange-declaration-failure reasons, including in particular either a *local* exchange already existing with different exchange-type,<sup>11</sup> or the remote *exchange-point* already existing with different exchange-type.<sup>12</sup>

<sup>11</sup>That is, a local exchange called `marketbids` with exchange-type something other than `x-exchange-point`.

<sup>12</sup>That is, a remote exchange-point called `bids` within the `market` space with exchange-type something other than `topic`.



## Margaret configures permissions for Alice and Bob

Alice calls Margaret, who creates a RabbitMQ user<sup>13</sup> on `amqp.market.example` called `alice`, sets the password appropriately, and permits that user to write to and read from the exchange called `market`. Thinking ahead, Margaret does the same for `bob`, to prevent trouble later.

### Alice joins the space

Alice reissues her `Exchange.Declare` on `amqp.acme.example`. This time, the behind-the-scenes relay successfully contacts `amqp.market.example`, joins the exchange space, and sets up the internal plumbing required to make the exchange called `marketbids` on `amqp.acme.example` an alias for the exchange-point `bids` in the exchange-space `market`. Alice issues a similar command, setting up a local alias `marketmail` for exchange-point `mailbox` in exchange-space `market`:

```
Exchange.Declare(exchange = 'marketmail',
  type = 'x-exchange-point',
  arguments =
    {'exchangespace':
      {'implementation': 'prototype',
       'host': 'amqp.market.example',
       'space': 'market',
       'username': 'alice',
       'password': '41iC3'},
     'exchangenname': 'mailbox',
     'exchangetype': 'direct'})
```

Alice creates queues and binds them to `marketbids` and `marketmail` as usual.

### Bob joins the space

Bob issues commands to `amqp.brandcorp.example`, very similar to those Alice issued, with differences that let the exchange-point aliases being constructed fit smoothly in with Bob's existing application architecture.

```
Exchange.Declare(exchange = 'b',
  type = 'x-exchange-point',
  arguments =
    {'exchangespace':
      {'implementation': 'prototype',
       'host': 'amqp.market.example',
       'space': 'market',
       'username': 'bob',
       'password': '808'},
     'exchangenname': 'bids',
     'exchangetype': 'topic'})
```

<sup>13</sup>In this prototype, anyway. AMQP has no general notion of "user", instead delegating such issues to SASL, so here I'm referring instead to the authentication system currently implemented in RabbitMQ.

```
Exchange.Declare(exchange = 'm',
  type = 'x-exchange-point',
  arguments =
    {'exchangespace':
      {'implementation': 'prototype',
       'host': 'amqp.market.example',
       'space': 'market',
       'username': 'bob',
       'password': '808'},
     'exchangenname': 'mailbox',
     'exchangetype': 'direct'})
```

## The network management function helps coordinate Alice's and Bob's local exchange-point aliases

The behind-the-scenes relays in each of Alice's and Bob's local AMQP servers have been coordinating with each other as a result of Alice's and Bob's exchange-space-manipulating commands (the `Exchange.Declares` and the `Queue.Binds`).

### Alice and Bob communicate

Alice and Bob `Basic.Publish` messages to `marketbids/marketmail` and `b/m`, their respective local aliases for the exchange-points `bids` and `mailbox` in the market exchange-space. The behind-the-scenes relays arrange for messages to be delivered to appropriately bound queues.