

1
2
3
4
5
6
7
8
9
10
11
12
13
14

AMQP

Advanced Message Queuing Protocol

Protocol Specification

Version 0.8 June 2006 [amqp0-8-june19]
A General-Purpose Middleware Standard

Technical Contacts:

Carl Trieloff	Red Hat
Ciaran McHale	IONA Technology
Gordon Sim	Red Hat
Harold Piskiel	Envoy Technologies
John O'Hara	JPMorgan Chase
Jason Brome	Envoy Technologies
Kim van der Riet	Red Hat
Mark Atwell	JPMorgan Chase
Martin Lucina	iMatix Corporation
Pieter Hintjens	iMatix Corporation
Robert Greig	JPMorgan Chase
Sam Joyce	IONA Technology
Sanjay Shrivastava	Envoy Technologies

1 Copyright Notice

2 © Copyright JPMorgan Chase & Co., Cisco Systems, Inc., Envoy, iMatix Corporation, IONA Technologies,
3 Red Hat, Inc., TWIST Process Innovations, and 29West 2006. All rights reserved.

4 License

5 JPMorgan Chase & Co., Cisco Systems, Inc., Envoy, iMatix Corporation, IONA Technologies, Red Hat,
6 Inc., TWIST Process Innovations, and 29West (collectively, the "Authors") each hereby grants to you a
7 worldwide, perpetual, royalty-free, nontransferable, nonexclusive license to (i) copy, display, and
8 implement the Advanced Messaging Queue Protocol ("AMQP") Specification and (ii) the Licensed Claims
9 that are held by the Authors, all for the purpose of implementing the Advanced Messaging Queue Protocol
10 Specification. Your license and any rights under this Agreement will terminate immediately without notice
11 from any Author if you bring any claim, suit, demand, or action related to the Advanced Messaging Queue
12 Protocol Specification against any Author. Upon termination, you shall destroy all copies of the Advanced
13 Messaging Queue Protocol Specification in your possession or control.

14 As used hereunder, "Licensed Claims" means those claims of a patent or patent application, throughout
15 the world, excluding design patents and design registrations, owned or controlled, or that can be
16 sublicensed without fee and in compliance with the requirements of this Agreement, by an Author or its
17 affiliates now or at any future time and which would necessarily be infringed by implementation of the
18 Advanced Messaging Queue Protocol Specification. A claim is necessarily infringed hereunder only when
19 it is not possible to avoid infringing it because there is no plausible non-infringing alternative for
20 implementing the required portions of the Advanced Messaging Queue Protocol Specification.
21 Notwithstanding the foregoing, Licensed Claims shall not include any claims other than as set forth above
22 even if contained in the same patent as Licensed Claims; or that read solely on any implementations of any
23 portion of the Advanced Messaging Queue Protocol Specification that are not required by the Advanced
24 Messaging Queue Protocol Specification, or that, if licensed, would require a payment of royalties by the
25 licensor to unaffiliated third parties. Moreover, Licensed Claims shall not include (i) any enabling
26 technologies that may be necessary to make or use any Licensed Product but are not themselves expressly
27 set forth in the Advanced Messaging Queue Protocol Specification (e.g., semiconductor manufacturing
28 technology, compiler technology, object oriented technology, networking technology, operating system
29 technology, and the like); or (ii) the implementation of other published standards developed elsewhere and
30 merely referred to in the body of the Advanced Messaging Queue Protocol Specification, or (iii) any
31 Licensed Product and any combinations thereof the purpose or function of which is not required for
32 compliance with the Advanced Messaging Queue Protocol Specification. For purposes of this definition,
33 the Advanced Messaging Queue Protocol Specification shall be deemed to include both architectural and
34 interconnection requirements essential for interoperability and may also include supporting source code
35 artifacts where such architectural, interconnection requirements and source code artifacts are expressly
36 identified as being required or documentation to achieve compliance with the Advanced Messaging Queue
37 Protocol Specification.

38 As used hereunder, "Licensed Products" means only those specific portions of products (hardware,
39 software or combinations thereof) that implement and are compliant with all relevant portions of the
40 Advanced Messaging Queue Protocol Specification.

41 The following disclaimers, which you hereby also acknowledge as to any use you may make of the
42 Advanced Messaging Queue Protocol Specification:

1 THE ADVANCED MESSAGING QUEUE PROTOCOL SPECIFICATION IS PROVIDED "AS IS," AND THE
2 AUTHORS MAKE NO REPRESENTATIONS OR WARRANTIES, EXPRESS OR IMPLIED, INCLUDING,
3 BUT NOT LIMITED TO, WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR
4 PURPOSE, NON-INFRINGEMENT, OR TITLE; THAT THE CONTENTS OF THE ADVANCED
5 MESSAGING QUEUE PROTOCOL SPECIFICATION ARE SUITABLE FOR ANY PURPOSE; NOR THAT
6 THE IMPLEMENTATION OF THE ADVANCED MESSAGING QUEUE PROTOCOL SPECIFICATION
7 WILL NOT INFRINGE ANY THIRD PARTY PATENTS, COPYRIGHTS, TRADEMARKS OR OTHER
8 RIGHTS.

9 THE AUTHORS WILL NOT BE LIABLE FOR ANY DIRECT, INDIRECT, SPECIAL, INCIDENTAL OR
10 CONSEQUENTIAL DAMAGES ARISING OUT OF OR RELATING TO ANY USE, IMPLEMENTATION OR
11 DISTRIBUTION OF THE ADVANCED MESSAGING QUEUE PROTOCOL SPECIFICATION.

12 The name and trademarks of the Authors may NOT be used in any manner, including advertising or
13 publicity pertaining to the Advanced Messaging Queue Protocol Specification or its contents without
14 specific, written prior permission. Title to copyright in the Advanced Messaging Queue Protocol
15 Specification will at all times remain with the Authors.

16 No other rights are granted by implication, estoppel or otherwise.

17 Upon termination of your license or rights under this Agreement, you shall destroy all copies of the
18 Advanced Messaging Queue Protocol Specification in your possession or control.

19 **Status of this Document**

20 "JPMorgan", "JPMorgan Chase", "Chase", the JPMorgan Chase logo and the Octagon Symbol are
21 trademarks of JPMorgan Chase & Co.

22 IMATIX and the iMatix logo are trademarks of iMatix Corporation sprl.

23 IONA, IONA Technologies, and the IONA logos are trademarks of IONA Technologies PLC and/or its
24 subsidiaries.

25 LINUX is a trademark of Linus Torvalds. RED HAT and JBOSS are registered trademarks of Red Hat, Inc. in
26 the US and other countries.

27 Java, all Java-based trademarks and OpenOffice.org are trademarks of Sun Microsystems, Inc. in the United
28 States, other countries, or both.

29 Other company, product, or service names may be trademarks or service marks of others.

Table of Contents

1 Overview.....	7
1.1 Goals of This Document.....	7
1.2 Patents.....	7
1.3 Summary.....	7
1.3.1 What is the AMQ Protocol?.....	7
1.3.2 Why AMQ Protocol?.....	7
1.3.3 Scope of AMQ Protocol.....	7
1.3.4 The Advanced Message Queuing Protocol Model (AMQP Model).....	8
1.3.5 The Advanced Message Queuing Protocol (AMQP).....	9
1.3.6 Scales of Deployment.....	10
1.3.7 Functional Scope.....	10
1.4 Organisation of This Document.....	11
1.5 Conventions.....	11
1.5.1 Guidelines for Implementers.....	11
1.5.2 Technical Terminology.....	12
2 General Architecture.....	14
2.1 AMQ Protocol Model Architecture.....	14
2.1.1 Main Entities.....	14
2.1.2 Message Flow.....	17
2.1.3 Exchanges.....	19
2.1.4 Message Queues.....	19
2.1.5 Bindings.....	20
2.2 AMQ Protocol Command Architecture.....	24
2.2.1 Protocol Commands (Classes & Methods).....	24
2.2.2 Mapping AMQP to a middleware API.....	24
2.2.3 No Confirmations.....	25
2.2.4 The Connection Class.....	26
2.2.5 The Channel Class.....	26
2.2.6 The Access Class.....	27
2.2.7 The Exchange Class.....	27

2.2.8 The Queue Class.....	28
2.2.9 The Content Classes.....	28
2.2.10 The Transaction Class.....	29
2.2.11 The Distributed Transaction Class.....	30
2.3 AMQ Protocol Transport Architecture.....	30
2.3.1 General Description.....	30
2.3.2 Data Types.....	31
2.3.3 Protocol Negotiation.....	31
2.3.4 Delimiting Frames.....	31
2.3.5 Frame Details.....	32
2.3.6 Error Handling.....	33
2.3.7 Closing Channels and Connections.....	34
2.4 AMQ Protocol Client Architecture.....	34
3 Functional Specification.....	36
3.1 Server Functional Specification.....	36
3.1.1 Messages and Content.....	36
3.1.2 Virtual Hosts.....	36
3.1.3 Exchanges.....	37
3.1.4 Message Queues.....	39
3.1.5 Bindings.....	39
3.1.6 Consumers.....	39
3.1.7 Quality of Service.....	40
3.1.8 Acknowledgements.....	40
3.1.9 Flow Control.....	40
3.1.10 Naming Conventions.....	40
3.2 AMQP Command Specification (Classes & Methods).....	41
3.2.1 Explanatory Notes.....	41
3.2.2 Class and Method Ids.....	41
4 Technical Specifications.....	45
4.1 IANA Assigned Port Number.....	45
4.2 AMQP Wire-Level Format.....	45

4.2.1 Format Protocol Grammar.....	45
4.2.2 Protocol Header.....	47
4.2.3 General Frame Format.....	48
4.2.4 Method Frames.....	49
4.2.5 AMQP Data Fields.....	50
4.3 Channel Multiplexing.....	53
4.4 Error Handling.....	54
4.4.1 Exceptions.....	54
4.4.2 Reply Code Format.....	54
4.4.3 Channel Exception Reply Codes.....	55
4.4.4 Connection Exception Reply Codes.....	55
4.5 Limitations.....	56
4.6 Security.....	57
4.6.1 Goals and Principles.....	57
4.6.2 Denial of Service Attacks.....	57
5 Conformance Tests.....	58
5.1 Introduction.....	58
5.2 Design.....	58
5.2.1 “Test Sets” group Tests into meaningful capabilities.....	58
5.2.2 Wire-Level Tests.....	58
5.2.3 Functional Tests.....	58
5.3 Test Sets.....	59

1 Overview

1.1 Goals of This Document

This document defines a networking protocol, the Advanced Message Queuing Protocol (AMQP), which enables conforming client applications to communicate with conforming messaging middleware services. To fully achieve this we also define the normative behaviour of the messaging middleware service.

We address a technical audience with some experience in the domain, and we provide sufficient specifications and guidelines that a suitably skilled engineer can construct conforming solutions in any modern programming language or hardware platform.

1.2 Patents

A conscious design objective of AMQP was to base it on concepts taken from existing, unencumbered, widely implemented standards such those published by the Internet Engineering Task Force (IETF) or the World Wide Web Consortium (W3C).

Consequently, we believe it is possible to create AMQP implementations using only well known techniques such as those found in existing Open Source networking and email routing software or which are otherwise well-known to technology experts.

1.3 Summary

1.3.1 What is the AMQ Protocol?

The Advanced Message Queuing Protocol (AMQ Protocol or AMQP) creates full functional interoperability between conforming clients and messaging middleware servers (also called "brokers").

1.3.2 Why AMQ Protocol?

Our goal is to enable the development and industry-wide use of standardised messaging middleware technology that will lower the cost of enterprise and systems integration and provide industrial-grade integration services to a broad audience.

It is our aim that through AMQ Protocol messaging middleware capabilities may ultimately be driven into the network itself, and that through the pervasive availability of messaging middleware new kinds of useful applications may be developed.

1.3.3 Scope of AMQ Protocol

To enable complete interoperability for messaging middleware requires that both the networking protocol and the semantics of the broker services are sufficiently specified.

AMQP, therefore, defines both the network protocol and the broker services through:

- 1 ◆ A **defined set of messaging capabilities** called the "Advanced Message Queuing Protocol Model"
2 (AMQP Model). The AMQP Model consists of a set of components that route and store messages
3 within the broker service, plus a set of rules for wiring these components together.
- 4 ◆ A **network wire-level protocol**, AMQP, that lets client applications talk to the broker and interact with
5 the AMQP Model it implements.

6 One can partially imply the semantics of the server from the AMQP protocol specifications but we believe
7 that an explicit description of these semantics helps the understanding of the protocol.

8 1.3.4 The Advanced Message Queuing Protocol Model (AMQP Model)

9 We define the server's semantics explicitly, since interoperability demands that these be the same in any
10 given server implementation.

11 The AMQP Model thus specifies a modular set of components and standard rules for connecting these.

12 There are three main types of component, which are connected into processing chains in the server to
13 create the desired functionality:

- 14 ◆ The "**exchange**" receives messages from publisher applications and routes these to "message
15 queues", based on arbitrary criteria, usually message properties or content
- 16 ◆ The "**message queue**" stores messages until they can be safely processed by a consuming client
17 application (or multiple applications)
- 18 ◆ The "**binding**" defines the relationship between a message queue and an exchange and provides the
19 message routing criteria

20 Using this model we can emulate the classic middleware concepts of store-and-forward queues and topic
21 subscriptions trivially. We can also express less trivial concepts such as content-based routing,
22 message queue forking, and on-demand message queues.

23 In very gross terms, an AMQP server is analogous to an email server, with each exchange acting as a
24 message transfer agent, and each message queue as a mailbox. The bindings define the routing tables in
25 each transfer agent. Publishers send messages to individual transfer agents, which then route the
26 messages into mailboxes. Consumers take messages from mailboxes.

27 In many pre-AMQP middleware system, by contrast, publishers send messages directly to individual
28 mailboxes (in the case of store-and-forward queues), or to mailing lists (in the case of topic subscriptions).

29 The difference is that when the rules connecting message queues to exchanges are under control of the
30 architect (rather than embedded in code), it becomes possible to do interesting things, such as define a rule
31 that says, "place a copy of all messages containing such-and-such a header into this message queue".

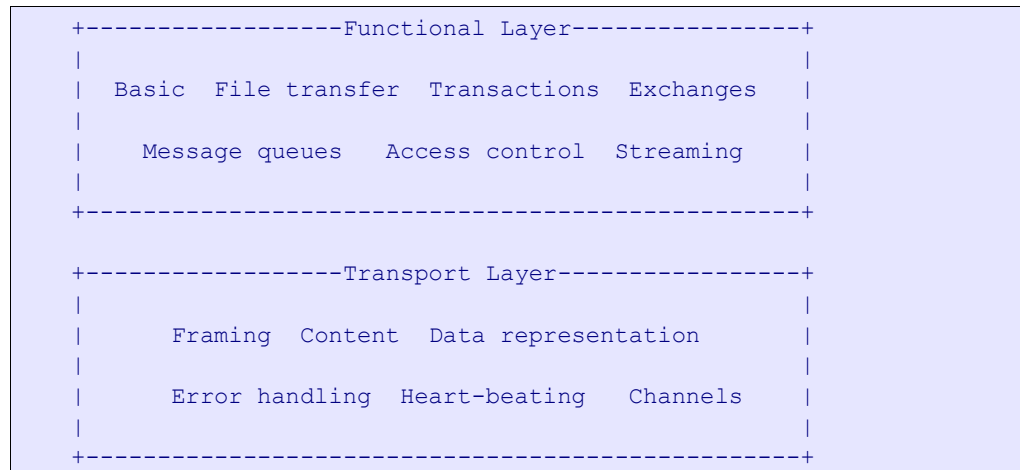
32 The design of the AMQP Model was driven by these main requirements:

- 33 ◆ To support the semantics required by the financial services industry
- 34 ◆ To provide the levels of performance required by the financial services industry
- 35 ◆ To be easily extended for new kinds of message routing and queueing
- 36 ◆ To permit the server's specific semantics to be programmed by the application, via the protocol
- 37 ◆ To be flexible yet simple.

1.3.5 The Advanced Message Queuing Protocol (AMQP)

The AMQP protocol is a binary protocol with modern features: it is multi-channel, negotiated, asynchronous, secure, portable, neutral, and efficient.

AMQP is usefully split into two layers:



The functional layer defines a set of commands (grouped into logical classes of functionality) that do useful work on behalf of the application.

The transport layer that carries these methods from application to server, and back, and which handles channel multiplexing, framing, content encoding, heart-beating, data representation, and error handling.

One could replace the transport layer with arbitrary transports without changing the application-visible functionality of the protocol. One could also use the same transport layer for different high-level protocols.

The design of AMQ Protocol Model was driven by these requirements:

- ◆ To guarantee interoperability between conforming implementations
- ◆ To provide explicit control over the quality of service
- ◆ To support any middleware domain: messaging, file transfer, streaming, RPC, etc
- ◆ To accommodate existing messaging API standards (for example, Sun's JMS)
- ◆ To be consistent and explicit in naming
- ◆ To allow complete configuration of server wiring via the protocol
- ◆ To use a command notation that maps easily into application-level API's
- ◆ To be clear, so each operation does exactly one thing.

The design of AMQP transport layer was driven by these main requirements, in no particular order:

- ◆ To be compact, using a binary encoding that packs and unpacks rapidly
- ◆ To handle messages of any size without significant limit
- ◆ To permit zero-copy data transfer (e.g. remote DMA)

- 1 ◆ To carry multiple channels across a single connection
- 2 ◆ To be long-lived, with no significant in-built limitations
- 3 ◆ To allow asynchronous command pipe-lining
- 4 ◆ To be easily extended to handle new and changed needs
- 5 ◆ To be forward compatible with future versions
- 6 ◆ To be repairable, using a strong assertion model
- 7 ◆ To be neutral with respect to programming languages
- 8 ◆ To fit a code generation process.

9 1.3.6 Scales of Deployment

10 The scope of AMQP covers different levels of scale, roughly as follows:

- 11 ◆ Developer/casual use: 1 server, 1 user, 10 message queues, 1 message per second
- 12 ◆ Production application: 2 servers, 10-100 users, 10-50 message queues, 10 messages per second (36K messages/hour)
- 13 ◆ Departmental mission critical application: 4 servers, 100-500 users, 50-100 message queues, 100 messages per second (360K/hour)
- 14 ◆ Regional mission critical application: 16 servers, 500-2,000 users, 100-500 message queues and topics, 1000 messages per second(3.6M/hour)
- 15 ◆ Global mission critical application: 64 servers, 2K-10K users, 500-1000 message queues and topics, 10,000 messages per second(36M/hour)
- 16 ◆ Market data (trading): 200 servers, 5K users, 10K topics, 100K messages per second (360M/hour)

17 As well as volume, the latency of message transfer can be highly important. For instance, market data becomes worthless very rapidly. Implementations may differentiate themselves by providing differing Quality of Service or Manageability Capabilities whilst remaining fully compliant with this specification.

24 1.3.7 Functional Scope

25 We want to support a variety of messaging architectures:

- 26 ◆ Store-and-forward with many writers and one reader
- 27 ◆ Transaction distribution with many writers and many readers
- 28 ◆ Publish-subscribe with many writers and many readers
- 29 ◆ Content-based routing with many writers and many readers
- 30 ◆ Queued file transfer with many writers and many readers
- 31 ◆ Point-to-point connection between two peers
- 32 ◆ Market data distribution with many sources and many readers.

33 1.4 Organisation of This Document

34 The document is divided into six chapters, most of which are designed to be read independently according to your level of interest:

1. **"Overview"** (this chapter). Read this chapter for an introduction
2. **"General Architecture"**, in which we describe the architecture and overall design of AMQP. This chapter is intended to help systems architects understand how AMQP works
3. **"Functional Specifications"**, in which we define how applications work with AMQP. This chapter consists of a readable discussion, followed by a detailed specification of each protocol command, intended as a reference for implementers. Before reading this chapter you should read the General Architecture
4. **"Technical Specifications"**, in which we define how the AMQP transport layer works. This chapter consists of a short discussion, followed by a detailed specification of the wire-level constructs, intended as a reference for implementers. You can read this chapter by itself if you want to understand how the wire-level protocol works (but not what it is used for)
5. **"Conformance Tests"**, in which we explain the conformance tests, which assert that an AMQP server conforms to the functional and technical specifications defined in this document. You can read this chapter by itself
6. **"Background"**, in which we state and analyse the scope and requirements of the AMQP standard and describe some of the underlying motivations behind the most important features of the protocol. This chapter comes last because it is not part of the knowledge needed to write an AMQP implementation, but it does provide useful background understanding. Note that the specification chapters include statements of key requirements, without analysis.

1.5 Conventions

1.5.1 Guidelines for Implementers

- ◆ We use the terms **MUST**, **MUST NOT**, **SHOULD**, **SHOULD NOT**, and **MAY** as defined by IETF RFC 2119
- ◆ We use the term "the server" when discussing the specific behaviour required of a conforming AMQP server
- ◆ We use the term "the client" when discussing the specific behaviour required of a conforming AMQP client
- ◆ We use the term "the peer" to mean "the server or the client"
- ◆ All numeric values are decimal unless otherwise indicated
- ◆ Protocol constants are shown as upper-case names. AMQP implementations **SHOULD** use these names when defining and using constants in source code and documentation
- ◆ Property names, method arguments, and frame fields are shown as lower-case names. AMQP implementations **SHOULD** use these names consistently in source code and documentation.

1.5.2 Technical Terminology

These terms have special significance within the context of this document:

- ◆ **AMQP Command Architecture:** An encoded wire-level protocol command which executes actions on the state of the AMQP Model Architecture.
- ◆ **AMQP Model Architecture:** A logical framework representing the key entities and semantics which must be made available by an AMQP compliant server implementation, such that the server can be

1 meaningfully manipulated by AMQP Commands sent from a client in order to achieve the semantics
2 defined in this specification.

- 3 ◆ **Connection:** A network connection, e.g. a TCP/IP socket connection
- 4 ◆ **Channel:** A bi-directional stream of communications between two AMQP peers. Channels are
5 multiplexed so that a single network connection can carry multiple channels
- 6 ◆ **Client:** The initiator of an AMQP connection or channel. AMQP is not symmetrical. Clients produce
7 and consume messages while servers queue and route messages
- 8 ◆ **Server:** The process that accepts client connections and implements the AMQP message queueing
9 and routing functions. Also known as "broker"
- 10 ◆ **Peer:** Either party in an AMQP connection. An AMQP connection involves exactly two peers (one is
11 the client, one is the server)
- 12 ◆ **Frame:** A formally-defined package of connection data. Frames are always written and read
13 contiguously - as a single unit - on the connection
- 14 ◆ **Protocol Class:** A collection of AMQP commands (also known as Methods) that deal with a specific
15 type of functionality
- 16 ◆ **Method:** A specific type of AMQP command frame that passes instructions from one peer to the other
- 17 ◆ **Content:** Application data passed from client to server and from server to client. AMQP content can
18 be structured into multiple parts. The term is synonymous with "message"
- 19 ◆ **Content Header:** A specific type of frame that describes a content's properties
- 20 ◆ **Content Body:** A specific type of frame that contains raw application data. Content body frames are
21 entirely opaque - the server does not examine or modify these in any way
- 22 ◆ **Message:** Synonymous with "content"
- 23 ◆ **Exchange:** The entity within the server which receives messages from producer applications and
24 optionally routes these to message queues within the server
- 25 ◆ **Exchange Type:** The algorithm and implementation of a particular model of exchange. In contrast to the
26 "exchange instance", which is the entity that receives and routes messages within the server
- 27 ◆ **Message queue:** A named entity that holds messages and forwards them to consumer applications.
- 28 ◆ **Binding:** An entity that creates a relationship between a message queue and an exchange
- 29 ◆ **Routing key:** A virtual address that an exchange may use to decide how to route a specific message
- 30 ◆ **Durable:** A server resource that survives a server restart
- 31 ◆ **Transient:** A server resource that is wiped or reset after a server restart
- 32 ◆ **Persistent:** A message that the server holds on reliable disk storage and MUST NOT lose after a
33 server restart
- 34 ◆ **Non-persistent:** A message that the server holds in memory and MAY lose after a server restart
- 35 ◆ **Consumer:** A client application that requests messages from a message queue
- 36 ◆ **Producer:** A client application that publishes messages to an exchange
- 37 ◆ **Virtual host:** A collection of exchanges, message queues and associated objects. Virtual hosts are
38 independent server domains that share a common authentication and encryption environment. The
39 client application chooses a virtual host after logging in to the server

- 1 ◆ **Realm:** A set of server resources (exchanges and message queues) covered by a single security policy
2 and access control. Applications ask for access rights for specific realms, rather than for specific
3 resources
- 4 ◆ **Ticket:** A token that a server provides to a client, for access to a specific realm
- 5 ◆ **Streaming:** The process by which the server will send messages to the client at a pre-arranged rate
- 6 ◆ **Staging:** The process by which a peer will transfer a large message to a temporary holding area before
7 formally handing it over to the recipient. This is how AMQP implements re-startable file transfers
- 8 ◆ **Out-of-band transport:** The technique by which data is carried outside the network connection. For
9 example, one might send data across TCP/IP and then switch to using shared memory if one is talking
10 to a peer on the same system
- 11 ◆ **Zero copy:** The technique of transferring data without copying it to or from intermediate buffers. Zero
12 copy requires that the protocol allows the out-of-band transfer of data as opaque blocks, as AMQP
13 does
- 14 ◆ **Assertion:** A condition that must be true for processing to continue
- 15 ◆ **Exception:** A failed assertion, handled by closing either the Channel or the Connection
- 16 These terms have **no special significance** within the context of AMQP:
- 17 ◆ **Topic:** Usually a means of distributing messages; AMQP implements topics using one or more types of
18 exchange
- 19 ◆ **Subscription:** Usually a request to receive data from topics; AMQP implements subscriptions as
20 message queues and bindings
- 21 ◆ **Service:** Usually synonymous with server. The AMQP standard uses "server" to conform with IETF
22 standard nomenclature and to clarify the roles of each party in the protocol (both sides may be AMQP
23 services)
- 24 ◆ **Broker:** synonymous with server. The AMQP standard uses the terms "client" and "server" to
25 conform with IETF standard nomenclature.
- 26 ◆ **Router:** Sometimes used to describe the actions of an exchange. However exchanges can do more than
27 message routing (they can also act as message end-points), and the term "router" has special
28 significance in the network domain, so AMQP avoids using it.


2 General Architecture

2.1 AMQ Protocol Model Architecture

This section explains the server semantics that must be standardised in order to guarantee interoperability between AMQP implementations.

2.1.1 Main Entities

This diagram shows the overall AMQ Protocol Model:



We can summarise what a middleware server is: it is a data server that accepts messages and does two main things with them, it routes them to different consumers depending on arbitrary criteria, and it buffers them in memory or on disk when consumers are not able to accept them fast enough.

In a pre-AMQP server these tasks are done by monolithic engines that implement specific types of routing and buffering. The AMQ Protocol Model takes the approach of smaller, modular pieces that can be combined in more diverse and robust ways. It starts by dividing these tasks into two distinct roles:

- ◆ The exchange, which accepts messages from producers and routes them message queues
- ◆ The message queue, which stores messages and forwards them to consumer applications

There is a clear interface between exchange and message queue, called a "binding", which we will come to later. The usefulness of the AMQ Protocol Model comes from three main features:

1. The ability to create arbitrary exchange and message queue types (some are defined in the standard, but others can be added as server extensions)
2. The ability to wire exchanges and message queues together to create any required message-processing system
3. The ability to control this completely through the protocol

In fact, AMQP provides runtime-programmable semantics.

2.1.1.1 The Message Queue

A message queue stores messages in memory or on disk, and delivers these in sequence to one or more consumer applications. Message queues are message storage and distribution entities. Each message queue is entirely independent and is a reasonably clever object.

A message queue has various properties: private or shared, durable or temporary, client-named or server-named, etc. By selecting the desired properties we can use a message queue to implement conventional middleware entities such as:

- ◆ A standard **store-and-forward queue**, which holds messages and distributes these between consumers on a round-robin basis. Store and forward queues are typically durable and shared between multiple consumers
- ◆ A **temporary reply queue**, which holds messages and forwards these to a single consumer. Reply queues are typically temporary, server-named, and private to one consumer
- ◆ A "**pub-sub**" subscription queue, which holds messages collected from various "subscribed" sources, and forwards these to a single consumer.

Subscription queues are typically temporary, server-named, and private to one consumer.

These categories are not formally defined in AMQP: they are examples of how message queues can be used. It is trivial to create new entities such as durable, shared subscription queues.

2.1.1.2 The Exchange

An exchange accepts messages from a producer application and routes these to message queues according to pre-arranged criteria. These criteria are called "bindings". Exchanges are matching and routing engines. That is, they inspect messages and using their binding tables, decide how to forward these messages to message queues or other exchanges. Exchanges never store messages.

The term "exchange" is used to mean both a class of algorithm, and the instances of such an algorithm. More properly, we speak of the "exchange type" and the "exchange instance".

AMQP defines a number of standard exchange types, which cover the fundamental types of routing needed to do common message delivery. AMQP servers will provide default instances of these exchanges. Applications that use AMQP can additionally create their own exchange instances. Exchange types are named so that applications which create their own exchanges can tell the server what exchange type to use. Exchange instances are also named so that applications can specify how to bind queues and publish messages.

Exchanges can do more than route messages. They can act as intelligent agents that work from within the server, accepting messages and producing messages as needed. The exchange concept is intended to define a model for adding extensibility to AMQP servers in a reasonably standard way, since extensibility has some impact on interoperability.

2.1.1.3 The Routing Key

In the general case an exchange examines a message's properties, its header fields, and its body content, and using this and possibly data from other sources, decides how to route the message.

1 In the majority of simple cases the exchange examines a single key field, which we call the "routing key".
2 The routing key is a virtual address that the exchange may use to decide how to route the message.

3 For **point-to-point routing, the routing key is the name of a message queue.**

4 For **topic pub-sub routing, the routing key is the topic hierarchy value.**

5 In more complex cases the routing key may be combined with routing on message header fields and/or its
6 content.

7 2.1.1.4 Analogy to Email

8 If we make an analogy with an email system we see that the AMQP concepts are not radical:

- 9 ◆ an AMQP message is analogous to an email message
- 10 ◆ a message queue is like a mailbox
- 11 ◆ a consumer is like a mail client that fetches and deletes email
- 12 ◆ a exchange is like a MTA (mail transfer agent) that inspects email and decides, on the basis of routing
13 keys and tables, how to send the email to one or more mailboxes
- 14 ◆ a routing key corresponds to an email To: or Cc: or Bcc: address, without the server information
15 (routing is entirely internal to an AMQP server)
- 16 ◆ each exchange instance is like a separate MTA process, handling some email sub-domain, or particular
17 type of email traffic
- 18 ◆ a binding is like an entry in a MTA routing table.

19 The power of AMQP comes from our ability to create queues (mailboxes), exchanges (MTA processes),
20 and bindings (routing entries), at runtime, and to chain these together in ways that go far beyond a simple
21 mapping from "to" address to mailbox name.

22 We should not take the email-AMQP analogy too far: there are fundamental differences. The challenge in
23 AMQP is to route and store messages within a server, or SMTP¹ parlance calls them "autonomous
24 systems". By contrast, the challenge in email is to route messages between autonomous systems.

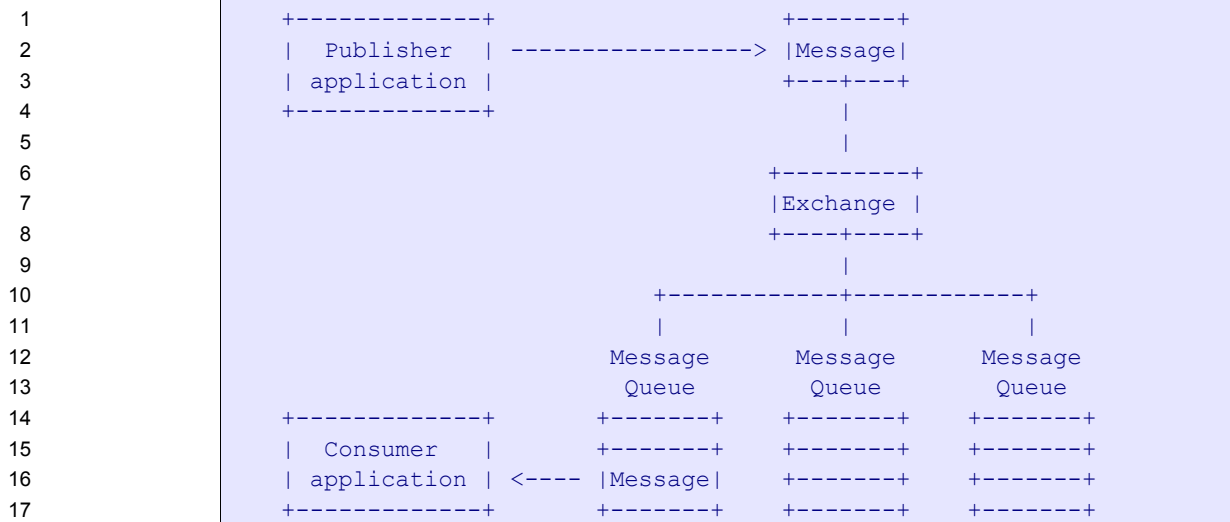
25 Routing within a server and between servers are distinct problems and have distinct solutions, if only for
26 banal reasons such as maintaining transparent performance.

27 To route between AMQP servers owned by different entities, one sets up explicit bridges, where one
28 AMQP server acts and the client of another server for the purpose of transferring messages between those
29 separate entities. This way of working tends to suit the types of businesses where AMQP is expected to
30 be used, because these bridges are likely to be underpinned by business processes, contractual
31 obligations and security concerns. This model also makes AMQP 'spam' more difficult.

32 2.1.2 Message Flow

33 This diagram shows the flow of messages through the AMQP Model server:

¹ SMTP is the Simple Mail Transport Protocol as defined by the IETF.



2.1.2.1 Message Life-cycle

An AMQP message consists of a set of properties plus opaque content.

A new “message” is created by a producer application using an AMQP client API. The producer places “content” in the message and perhaps sets some message “properties”. The producer labels the message with “routing information”, which is superficially similar to an address, but almost any scheme can be created. The producer then sends the message to an “exchange” on the server.

When the message arrives at the server, the exchange (usually) routes the message to a set of message “queues” which also exist on the server. If the message is unroutable, the exchange may drop it silently or return it to the producer. The producer chooses how unroutable messages are treated.

A single message can exist on many message queues. The server can handle this in different ways, by copying the message, by using reference counting, etc. This does not affect interoperability. However, when a message is routed to multiple message queues, it is identical on each message queue. There is no unique identifier that distinguishes the various copies.

When a message arrives in a message queue, the message queue tries immediately to pass it to a consumer application via AMQP. If this is not possible, the message queue stores the message (in memory or on disk as requested by the producer) and waits for a consumer to be ready. If there are no consumers, the message queue may return the message to the producer via AMQP (again, if the producer asked for this).

When the message queue can deliver the message to a consumer, it removes the message from its internal buffers. This can happen immediately, or after the consumer has acknowledged that it has successfully processed the message. The consumer chooses how and when messages are “acknowledged”. The consumer can also reject a message (a negative acknowledgement).

Producer messages and consumer acknowledgements are grouped into “transactions”. When an application plays both roles, which is often, it does a mix of work: sending messages and sending acknowledgements, and then committing or rolling back the transaction.

Message deliveries from the server to the consumer are not transacted; it is sufficient to transact the acknowledgements to these messages

2.1.2.2 What The Producer Sees

By analogy with the email system, we can see that a producer does not send messages directly to a message queue. Allowing this would break the abstraction in the AMQP Model. It would be like allowing email to bypass the MTA's routing tables and arrive directly in a mailbox. This would make it impossible to insert intermediate filtering and processing, spam detection, for instance.

The AMQP Model uses the same principle as an email system: all messages are sent to a single point, the exchange or MTA, which inspects the messages based on rules and information that are hidden from the sender, and routes them to drop-off points that are also hidden from the sender.

2.1.2.3 What The Consumer Sees

Our analogy with email starts to break down when we look at consumers. Email clients are passive - they can read their mailboxes, but they do not have any influence on how these mailboxes are filled. An AMQP consumer can also be passive, just like email clients. That is, we can write an application that expects a particular message queue to be ready and bound, and which will simply process messages off that message queue.

However, we also allow AMQP client applications to:

- ◆ create or destroy message queues
- ◆ define the way these message queues are filled, by making bindings
- ◆ select different exchanges which can completely change the routing semantics

This is like having an email system where one can, via the protocol:

- ◆ create a new mailbox
- ◆ tell the MTA that all messages with a specific header field should be copied into this mailbox
- ◆ completely change how the mail system interprets addresses and other message headers

We see that AMQP is more like a language for wiring pieces together than a system. This is part of our objective, to make the server behaviour programmable via the protocol.

2.1.2.4 Automatic Mode

Most integration architectures do not need this level of sophistication. Like the amateur photographer, a majority of AMQP users need a "point and shoot" mode. AMQP provides this through the use of two simplifying concepts:

- ◆ a **default exchange for message producers**
- ◆ a **default binding for message queues** that selects messages based on a match between routing key and message queue name

In effect, **the default binding lets a producer send messages directly to a message queue**, given suitable authority – it emulates the simplest “send to destination” addressing scheme people have come to expect of traditional middleware.

The default binding does not prevent the message queue from being used in more sophisticated ways. It does, however, let one use AMQP without needing to understand how exchanges and bindings work.

2.1.3 Exchanges

2.1.3.1 Types of Exchange

Each exchange type implements a specific routing algorithm. There are a number of standard exchange types, explained in the "Functional Specifications" chapter, but there are two that are particularly important:

- ◆ the "direct" exchange type, which routes on a routing key
- ◆ the "topic" exchange type, which routes on a routing pattern

Note that:

1. the default exchange is a "direct" exchange
2. the server will create a "direct" and a "topic" exchange at start-up with well-known names and client applications may depend on this

2.1.3.2 Exchange Life-cycle

Each AMQP server pre-creates a number of exchanges (more pedantically, "exchange instances"). These exchanges exist when the server starts and cannot be destroyed.

AMQP applications can also create their own exchanges. AMQP does not use a "create" method as such, it uses a "declare" method which means, "create if not present, otherwise continue". It is plausible that applications will create exchanges for private use and destroy them when their work is finished. AMQP provides a method to destroy exchanges but in general applications do not do this.

In our examples in this chapter we will assume that the exchanges are all created by the server at start-up. We will not show the application declaring its exchanges.

2.1.4 Message Queues

2.1.4.1 Message Queue Properties

When a client application creates a message queue, it can select some important properties:

- ◆ **name** - if left unspecified, the server chooses a name and provides this to the client. Generally, when applications share a message queue they agree on a message queue name beforehand, and when an application needs a message queue for its own purposes, it lets the server provide a name
- ◆ **durable** - if specified, the message queue remains present and active when the server restarts. It may lose non-persistent messages if the server restarts
- ◆ **auto-delete** - if specified, the server will delete the message queue when all clients have finished using it, or shortly thereafter.

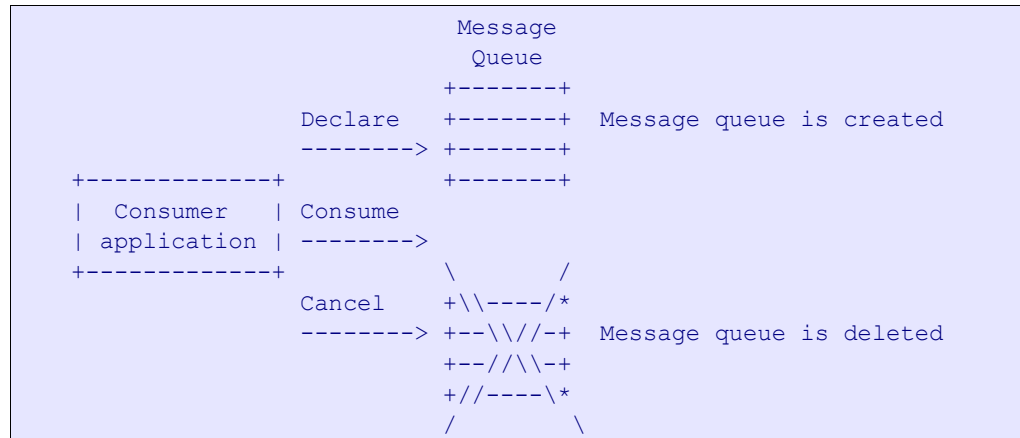
2.1.4.2 Queue Life-cycles

There are two main message queue life-cycles:

- ◆ **Durable message 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 message 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 on these, such as **shared message queues** that are deleted when the last of many consumers disconnects.

This diagram shows the way temporary message queues are created and deleted:



2.1.5 Bindings

A binding is the relationship between an exchange and a message queue that tells the exchange how to route messages. Bindings are constructed from commands from the client application (the one owning and using the message queue) to an exchange. We can express a binding command in pseudo-code as follows:

```
Queue.Bind <queue> TO <exchange> WHERE <condition>
```

Let's look at three typical use cases: shared queues, private reply queues, and pub-sub subscriptions.

2.1.5.1 Constructing a Shared Queue

Shared queues are the classic middleware "point-to-point queue". In AMQP we can use the default exchange and default binding. Let's assume our message queue is called "app.svc01". Here is the pseudo-code for creating the shared queue:

```
Queue.Declare
  queue=app.svc01
  private=FALSE
```

We may have many consumers on this shared queue. To consume from the shared queue, each consumer does this:

```
Basic.Consume
  queue=app.svc01
```

To publish to the shared queue, each producer sends a message to the default exchange:

```
Basic.Publish
  routing_key=app.svc01
```

2.1.5.2 Constructing a Reply Queue

Reply queues are usually temporary, with server-assigned names. They are also usually private, i.e. read by a single consumer. Apart from these particularities, reply queues use the same matching criteria as standard queues, so we can also use default exchange.

Here is the pseudo-code for creating a reply queue, where S: indicates a server reply:

```
Queue.Declare
  queue=<empty>
  private=TRUE
  auto_delete=TRUE
S:Queue.Create-Ok
  queue=tmp.1
```

To publish to the reply queue, a producer sends a message to the default exchange:

```
Basic.Publish
  routing_key=tmp.1
```

One of the standard message properties is Reply-To, which is designed specifically for carrying the name of reply queues.

2.1.5.3 Constructing a Pub-Sub Subscription Queue

In classic middleware the term "subscription" is vague and refers to at least two different concepts: the set of criteria that match messages and the temporary queue that holds matched messages. AMQP separates the work into bindings and message queues. There is no AMQP entity called "subscription".

Let us agree that a pub-sub subscription:

- ◆ holds messages for a single consumer (or in some cases for multiple consumers)
- ◆ collects messages from multiple sources, through a set of bindings that match topics, message fields, or content in different ways.

The key difference between a subscription queue and a named or reply queue is that the subscription queue name is irrelevant for the purposes of routing, and routing is done on abstracted matching criteria rather than a 1-to-1 matching of the routing key field.

Let's take the common pub-sub model of "topic trees" and implement this. We need an exchange type capable of matching on a topic tree. In AMQP this is the "topic" exchange type. The topic exchange matches wild-cards like "STOCK.USD.*" against routing key values like "STOCK.USD.NYSE".

We **cannot** use the default exchange or binding because these do not do topic-style routing. So we have to create a binding explicitly. Here is the pseudo-code for creating and binding the pub-sub subscription queue:

```
1 Queue.Declare
2   queue=<empty>
3   auto_delete=TRUE
4 S:Queue.Declare-Ok
5   queue=tmp.2
6 Queue.Bind
7   queue=tmp.2
8   TO exchange=amq.topic
9   WHERE routing_key=STOCK.USD.*
```

10 To consume from the subscription queue, the consumer does this:

```
11 Basic.Consume
12   queue=tmp.2
```

13 When publishing a message, the producer does something like this:

```
14 Basic.Publish
15   exchange=amq.topic
16   routing_key=STOCK.USD.IBM
```

17 The topic exchange processes the incoming routing key ("STOCK.USD.IBM") with its binding table, and
18 finds one match, for tmp.2. It then routes the message to that subscription queue.

19 2.1.5.4 Chained Bindings

20 The basic structures explained above are enough to implement shared queues and standard pub-sub
21 topics. However, some applications need more than this: they need to be able to combine matching
22 algorithms so that messages are matched several times before they reach a client application.

23 We want to provide the following semantic:

```
24 Queue.Bind <queue> TO <exchange1> WHERE <condition>
25   AND TO <exchange2> WHERE <condition>
```

26 Note that the "OR" semantic is trivial, we simply make two separate bindings for the same message queue.
27 It is the "AND" semantic that is non-trivial.

28 For performance reasons, AMQP does not provide an actual language in which to express such semantics.
29 Rather we will construct the combined semantic from individual methods:

```
30 Queue.Bind <queue> TO <exchange1> WHERE <condition>
31 Queue.Bind <queue> TO <exchange2> VIA <exchange1> WHERE <condition>
```

32 We call this a "chained binding". To see how this would work in practice, let's take two such algorithms as
33 examples: one is "topic", which matches the routing key against a wild-card pattern, and the other is
34 "filter", which detects illegal messages. (Note, "filter" is not real, just an example.) We have two exchange
35 instances, amq.topic and amq.filter.

36 We want to match all messages with routing key like "STOCK.USD.*" and which have a PDF file in their
37 content (one of the abilities of the imaginary filter exchange is to filter according to the content types of
38 messages).

39 We create a temporary message queue and bind it as follows:

```
1 Queue.Declare
2     queue=<empty>
3     auto_delete=TRUE
4 S:Queue.Create-Ok
5     queue=tmp.3
6 Queue.Bind
7     queue=tmp.3
8     TO exchange=amq.topic
9     WHERE routing_key=STOCK.USD.*
10 Queue.Bind
11     queue=tmp.3
12     TO exchange=amq.filter
13     VIA exchange=amq.topic
14     WHERE content-type=application/x-pdf
```

15 To publish a message, we send to the amq.topic as before:

```
16 Basic.Publish
17     exchange=amq.topic
18     routing_key=STOCK.USD.IBM
19     content-type=application/x-pdf
```

20 2.1.5.5 Message Selectors

21 Applications need to be able to select messages from message queues, at the same time as consuming
22 them. While this might be inefficient in a particular AMQP server implementation (it usually means
23 scanning messages sequentially), it is a common requirement because it is conceptually simple and similar
24 in some ways to the SQL SELECT statement. AMQP must therefore support this.

25 We call this a "message selector". We can compare this to a normal binding:

- 26 ◆ A normal binding routes messages into a message queue after which they are dispatched to N
27 consumers on a round-robin basis. For example, we may route print jobs to different message queues
28 based on where the printed documents must go
- 29 ◆ The select-on-consume semantic works on a single message queue for N consumers, but where each
30 consumer may have an independent set of criteria for the messages it wants to process. For example,
31 one of a group of printers servicing a single print queue might ask to receive all oversized documents,
32 by preference.

33 The main difference is that (a) the round-robin nature of message delivery remains in force, so if multiple
34 consumers have the same selector criteria, they will share the messages, and (b) the delivery of messages
35 remains ordered. Using normal bindings, this is not possible.

36 So message selectors apply when the server **searches** for an eligible consumer for the next message.
37 Searches may be inherently slower than normal bindings, because a message will be matched multiple times
38 rather than just once.

39 We can express a message selector command in pseudo-code as follows:

```
Basic.Consume FROM <queue> WHERE <condition>
```

2.2 AMQ Protocol Command Architecture

This section explains how the application talks to the server.

2.2.1 Protocol Commands (Classes & Methods)

Middleware is complex, and our challenge in designing the protocol structure was to tame that complexity. Our approach has been to model a traditional API based on classes which contain methods, and to define methods to do exactly one thing, and do it well. This results in a large command set but one that is relatively easy to understand.

The AMQP commands are grouped into classes. Each class covers a specific functional domain. Some classes are optional - each peer implements the classes it needs to support.

There are two distinct method dialogues:

- ◆ Synchronous request-response, in which one peer sends a request and the other peer sends a reply. Synchronous request and response methods are used for functionality that is not performance critical
- ◆ Asynchronous notification, in which one peer sends a method but expects no reply. Asynchronous methods are used where where performance is critical.

To make method processing simple, we define distinct replies for each synchronous request. That is, no method is used as the reply for two different requests. This means that a peer, sending a synchronous request, can accept and process incoming methods until getting one of the valid synchronous replies. This differentiates AMQP from more traditional RPC protocols.

A method is formally defined as a synchronous request, a synchronous reply (to a specific request), or asynchronous. Lastly, each method is formally defined as being client-side (i.e. server to client), or server-side (client to server).

2.2.2 Mapping AMQP to a middleware API

We have designed AMQP to be mappable to a middleware API. This mapping has some intelligence (not all methods, and not all arguments make sense to an application) but it is also mechanical (given some rules, all methods can be mapped without manual intervention).

The advantages of this are that having learnt the AMQP semantics (the classes that are described in this section), developers will find the same semantics provided in whatever environment they use.

For example, here is a Queue.Declare method example:

```
Queue.Declare
  queue=my.queue
  auto_delete=TRUE
  exclusive=FALSE
```

This can be cast as a wire-level frame:

1	+-----+-----+-----+-----+-----+
2	Queue Declare my.queue 1 0
3	+-----+-----+-----+-----+-----+
4	class method name autodel excl.

5 Or as a higher-level API:

```
6 queue_declare (session, "my.queue", TRUE, FALSE);
```

7 Or as an abstract language:

```
8 <queue_declare name = "my.queue" auto_delete = "1"
9 exclusive = "FALSE" />
```

10 There are two main exceptions to making the entire protocol isomorphic with the client API:

- 11 ♦ Existing API standards, such as JMS, which must be mapped manually onto the AMQP methods.
- 12 ♦ Those AMQP methods concerned with connection and session start-up and shut-down, which are not
- 13 useful to expose in the high-level API.

14 The pseudo-code logic for mapping an asynchronous method is:

```
15 send method to server
```

16 The pseudo-code logic for mapping a synchronous method is:

```
17 send request method to server
18 repeat
19     wait for response from server
20     if response is an asynchronous method
21         process method (usually, delivered or returned content)
22     else
23         assert that method is a valid response for request
24         exit repeat
25     end-if
26 end-repeat
```

27 It is worth commenting that for most applications, middleware can be completely hidden in technical layers,
28 and that the actual API used matters less than the fact that the middleware is robust and capable.

29 2.2.3 No Confirmations

30 A chatty protocol is slow. We use asynchronicity heavily in those cases where performance is an issue.
31 This is generally where we send content from one peer to another. We send off methods as fast as
32 possible, without waiting for confirmations. Where necessary, we implement windowing and throttling at a
33 higher level, e.g. at the consumer level.

34 We can dispense with confirmations because we adopt an assertion model for all actions. Either they
35 succeed, or we have an exception that closes the channel or connection.

36 There are no confirmations in AMQP. Success is silent, and failure is noisy. When applications need
37 explicit tracking of success and failure, they should use transactions.

2.2.4 The Connection Class

AMQP is a connected protocol. The connection is designed to be long-lasting, and can carry multiple channels.

The connection life-cycle is this:

1. The client opens a TCP/IP connection to the server and sends a protocol header. This is the only data the client sends that is not formatted as a method.
2. The server responds with its protocol version and other properties, including a list of the security mechanisms that it supports (the Start method).
3. The client selects a security mechanism (Start-Ok).
4. The server starts the authentication process, which uses the SASL challenge-response model. It sends the client a challenge (Secure).
5. The client sends an authentication response (Secure-Ok). For example using the "plain" mechanism, the response consist of a login name and password.
6. The server repeats the challenge (Secure) or moves to negotiation, sending a set of parameters such as maximum frame size (Tune).
7. The client accepts or lowers these parameters (Tune-Ok).
8. The client formally opens the connection and selects a virtual host (Open).
9. The server confirms that the virtual host is a valid choice (Open-Ok).
10. The client now uses the connection as desired.
11. One peer (client or server) ends the connection (Close).
12. The other peer hand-shakes the connection end (Close-Ok).
13. The server and the client close their socket connection.

2.2.5 The Channel Class

AMQP is a multi-channelled protocol. Channels provide a way to multiplex a heavyweight TCP/IP connection into several light weight connections. This makes the protocol more "firewall friendly" since port usage is predictable. It also means that traffic shaping and other network QoS features can be easily employed.

Channels are independent of each other and can perform different functions simultaneously with other channels, the available bandwidth being shared between the concurrent activities.

It is expected and encouraged that multi-threaded client applications may often use a "channel-per-thread" model as a programming convenience. However, opening several connections to one or more AMQP servers from a single client is also entirely acceptable.

The channel life-cycle is this:

1. The client opens a new channel (Open).
2. The server confirms that the new channel is ready (Open-Ok).
3. The client and server use the channel as desired.
4. One peer (client or server) closes the channel (Close).
5. The other peer hand-shakes the channel close (Close-Ok).

2.2.6 The Access Class

AMQP's access control model is based on "realms". A realm covers some group of server resources (exchanges and message queues) managed under a single security policy and access control. Applications ask for access to specific realms, rather than to specific resources. The server grants access in the form of "tickets", which the client application then uses accordingly. Tickets expire when the channel is closed, or if the server's access controls change.

The tickets granted in AMQP are **not** cryptographically secure, they are a memento that the server MAY use to accelerate access checking. The server **MUST NOT** trust the ticket. The server **MUST** always check a resource is accessible on each action where a ticket is presented. The ticket presented **SHOULD** be used as an opportunity for the system to optimise the access check logic.

Client applications **MUST** treat tickets as opaque data – and **MUST NOT** make assumptions as to ticket uniqueness, generation order, repeatability, etc.

The access ticket life-cycle is:

1. The client requests an access ticket for a realm (Request).
2. The server grants it (Request-Ok).
3. The server can, of course, refuse the request.

2.2.7 The Exchange Class

The exchange class lets an application manage exchanges on the server.

This class lets the application script its own wiring (rather than relying on some configuration interface).

Note: Most applications do not need this level of sophistication, and legacy middleware is unlikely to be able to support this semantic.

The exchange life-cycle is:

1. The client asks the server to make sure the exchange exists (Declare). The client can refine this into, "create the exchange if it does not exist", or "warn me but do not create it, if it does not exist".
2. The client publishes messages to the exchange.
3. The client may choose to delete the exchange (Delete).

2.2.8 The Queue Class

The queue class lets an application manage message queues on the server. This is a basic step in almost all applications that consume messages, at least to verify that an expected message queue is actually present.

The life-cycle for a durable message queue is fairly simple:

1. The client asserts that the message queue exists (Declare, with the "passive" argument).
2. The server confirms that the message queue exists (Declare-Ok).
3. The client reads messages off the message queue.

The life-cycle for a temporary message queue is more interesting:

- 1 1. The client creates the message queue (Declare, often with no message queue name so the server will
- 2 assign a name). The server confirms (Declare-Ok).
- 3 2. The client starts a consumer on the message queue. The precise functionality of a consumer depends
- 4 on the content class.
- 5 3. The client cancels the consumer, either explicitly or by closing the channel and/or connection.
- 6 4. When the last consumer disappears from the message queue, and after a polite time-out, the server
- 7 deletes the message queue.

8 AMQP implements the delivery mechanism for topic subscriptions as message queues. This enables
9 interesting structures where a subscription can be load balanced among a pool of co-operating subscriber
10 applications.

11 The life-cycle for a subscription involves an extra bind stage:

- 12 1. The client creates the message queue (Declare), and the server confirms (Declare-Ok).
- 13 2. The client binds the message queue to a topic exchange (Bind) and the server confirms (Bind-Ok).
- 14 3. The client uses the message queue as in the previous examples.

15 2.2.9 The Content Classes

16 Following the principle of placing functional domains into distinct protocol classes that the server may or
17 may not implement, AMQP also separates content processing into separate classes. The logic is that
18 different types of content have different semantics. For example, basic messages and file transfer are quite
19 different problems. We give each content type a class, and a set of methods that work with it.

20 AMQP currently defines three content classes:

- 21 1. Basic contents, which implement standard messaging semantics.
- 22 2. File contents, which support file-transfer semantics.
- 23 3. Stream contents, which support data streaming semantics.

24 2.2.9.1 The Basic Content Class

25 The Basic content class provides a superset of the message properties and functionality required to enable
26 the implementation of a Java Messaging Service client API which uses AMQP to communicate with any
27 AMQP server on any platform.

28 Most of the messaging capabilities described in this specification are enabled by the Basic content class.

29 The Basic content methods support these main semantics:

- 30 ♦ Sending messages from client to server, which happens asynchronously (Publish)
- 31 ♦ Starting and stopping consumers (Consume, Cancel)
- 32 ♦ Sending messages from server to client, which happens asynchronously (Deliver, Return)
- 33 ♦ Acknowledging messages (Ack, Reject)
- 34 ♦ Taking messages off the message queue synchronously (Get).

2.2.9.2 The File Content Class

The File content class enables AMQP to perform bulk file transfer in addition to messaging.

The File content class has specific support for restarting incomplete file transfers. We do this by sending file messages in two steps:

1. The sender uploads the file to the recipient. We call this "staging". If the upload is interrupted, the sender can recover and send only the missing part of the file.
2. The sender tells the recipient to process the file (e.g. to publish it).

The file content methods support these main semantics:

- ◆ Staging a file, from either peer to the other (Open, Stage)
- ◆ Sending a staged file from client to server, which happens asynchronously (Publish)
- ◆ Starting and stopping consumers (Consume, Cancel)
- ◆ Sending messages from server to client, which happens asynchronously (Deliver, Return)
- ◆ Acknowledging messages (Ack, Reject).

2.2.9.3 The Stream Content Class

The Stream content class is designed for content streaming (voice, video, etc.) It has these main semantics:

- ◆ Sending messages from client to server, which happens asynchronously (Publish)
- ◆ Starting and stopping consumers (Consume, Cancel)
- ◆ Sending messages from server to client, which happens asynchronously (Deliver, Return)

2.2.10 The Transaction Class

AMQP supports three kinds of transactions:

1. Automatic transactions, in which every published message and acknowledgement is processed as a stand-alone transaction.
2. Server local transactions, in which the server will buffer published messages and acknowledgements and commit them on demand from the client.
3. Distributed transactions, in which the server will synchronise its transactions with an external transaction coordinator.

The Transaction class ("tx") gives applications access to the second type, namely server transactions.

The semantics of this class are:

1. The application asks for server transactions in each channel where it wants these transactions (Select).
2. The application does work (Publish, Ack).
3. The application commits or rolls-back the work (Commit, Roll-back).
4. The application does work, ad infinitum.

2.2.11 The Distributed Transaction Class

The distributed transaction class ("dtx") provides simpler semantics because most of the work is done by the server and external transaction coordinator behind the scenes.

The semantics of this class are as follows:

1. The application asks for server transactions in each channel where it wants these transactions (Select).
2. The application does work (Publish, Ack).
3. AMQP arranges to propagate the global transaction ID.
4. Magic happens.

2.3 AMQ Protocol Transport Architecture

This section explains how commands are mapped to the wire-level protocol.

2.3.1 General Description

AMQP is a binary protocol. Information is organised into "frames", of various types. Frames carry protocol methods, structured contents, and other information. All frames have the same general format: frame header, payload, and frame end. The frame payload format depends on the frame type.

We assume a reliable stream-oriented network transport layer (TCP/IP or equivalent).

Within a single socket connection, there can be multiple independent threads of control, called "channels". Each frame is numbered with a channel number. By interleaving their frames, different channels share the connection. For any given channel, frames run in a strict sequence that can be used to drive a protocol parser (typically a state machine).

We construct frames using a small set of data types such as bits, integers, strings, and field tables. Frame fields are packed tightly without making them slow or complex to parse. It is relatively simple to generate framing layer mechanically from the protocol specifications.

The wire-level formatting is designed to be scalable and generic enough to be used for arbitrary high-level protocols (not just AMQP). We assume that AMQP will be extended, improved and otherwise varied over time and the wire-level format will support this.

2.3.2 Data Types

The AMQP data types are:

- ◆ Integers (from 1 to 8 octets), used to represent sizes, quantities, limits, etc. Integers are always unsigned and may be unaligned within the frame
- ◆ Bits, used to represent on/off values. Bits are packed into octets
- ◆ Short strings, used to hold short text properties. Short strings are limited to 255 octets and can be parsed with no risk of buffer overflows
- ◆ Long strings, used to hold chunks of binary data
- ◆ Field tables, which hold name-value pairs. The field values are typed as strings, integers, etc.

2.3.3 Protocol Negotiation

The AMQP client and server negotiate the protocol. This means that when the client connects, the server proposes certain options that the client can accept, or modify. When both peers agree on the outcome, the connection goes ahead. Negotiation is a useful technique because it lets us assert assumptions and preconditions.

In AMQP, we negotiate a number of specific aspects of the protocol:

- ◆ The actual protocol and version. An AMQP server MAY host multiple protocols on the same port
- ◆ Encryption arguments and the authentication of both parties. This is part of the functional layer, explained previously
- ◆ Maximum frame size, number of channels, and other operational limits.

Agreed limits MAY enable both parties to pre-allocate key buffers, avoiding deadlocks. Every incoming frame either obeys the agreed limits, and so is "safe", or exceeds them, in which case the other party IS faulty and MUST be disconnected. This is very much in keeping with the "it either works properly or it doesn't work at all" philosophy of AMQP.

Both peers negotiate the limits to the lowest agreed value as follows:

- ◆ The server MUST tell the client what limits it proposes
- ◆ The client responds and MAY reduce those limits for its connection.


2.3.4 Delimiting Frames

TCP/IP is a stream protocol, i.e. there is no in-built mechanism for delimiting frames. Existing protocols solve this in several different ways:

- ◆ Sending a single frame per connection. This is simple but slow
- ◆ Adding frame delimiters to the stream. This is simple but slow to parse
- ◆ Counting the size of frames and sending the size in front of each frame. This is simple and fast, and our choice.


2.3.5 Frame Details

All frames consist of a header (8 octets), a payload of arbitrary size, and a 'frame-end' octet that detects malformed frames:



To read a frame, we:

1. Read the header and check the frame type and channel.
2. Depending on the frame type, we read the payload and process it.
3. Read the frame end octet.



We place content body in distinct frames (rather than including it in the method) so that AMQP may support "zero copy" techniques in which content is never marshalled or encoded, and can be sent via out-of-band transport such as shared memory or remote DMA.

We place the content properties in their own frame so that recipients can selectively discard contents they do not want to process.

Contents can be structured with sub-contents to any level.

2.3.5.3 Out-of-band Frames

Out-of-band transport can be used in specific high-performance models. Note that this part of the protocol is speculative because we have not built a working out-of-band prototype. This part of the protocol is a place-holder rather than a formal proposal.

The principle of out-of-band transport is that a TCP/IP connection can be used for controlling another, faster but less abstract protocol such as remote-DMA, shared memory, or multicast.

2.3.5.4 Heartbeat Frames

Hearbeating is a technique designed to **undo** one of TCP/IP's features, namely its ability to recover from a broken physical connection by closing only after a quite long time-out. In some scenarios we need to know very rapidly if a peer is disconnected or not responding for other reasons (e.g. it is looping). Since heart-beating can be done at a low level, we implement this as a special type of frame that peers exchange at the transport level, rather than as a class method.

2.3.6 Error Handling

AMQP uses exceptions to handle errors. That is:

- ◆ Any operational error, e.g. message queue not found, insufficient access rights, etc. results in a channel exception.
- ◆ Any structural error, e.g. invalid argument, bad sequence of methods, etc. results in a connection exception.
- ◆ An exception closes the channel or connection, and returns a reply code and reply text to the client application. We use the 3-digit reply code plus textual reply text scheme that is used in HTTP and many other protocols.

2.3.7 Closing Channels and Connections

Closing a channel or connection for any reason - normal or exceptional - must be done carefully. Abrupt closure is not always detected rapidly, and following an exception, we could lose the error reply codes. The correct design is to hand-shake all closure so that we close only after we are sure the other party is aware of the situation.

1 When a peer decides to close a channel or connection, it sends a Close method. The receiving peer
2 responds with Close-Ok, and then both parties can close their channel or connection.

3 2.4 AMQ Protocol Client Architecture

4 It is possible to read and write AMQP frames directly from an application but this would be bad design.
5 Even the simplest AMQP dialogue is rather more complex than, say HTTP, and application developers
6 should not need to understand such things as binary framing formats in order to send a message to a
7 message queue.

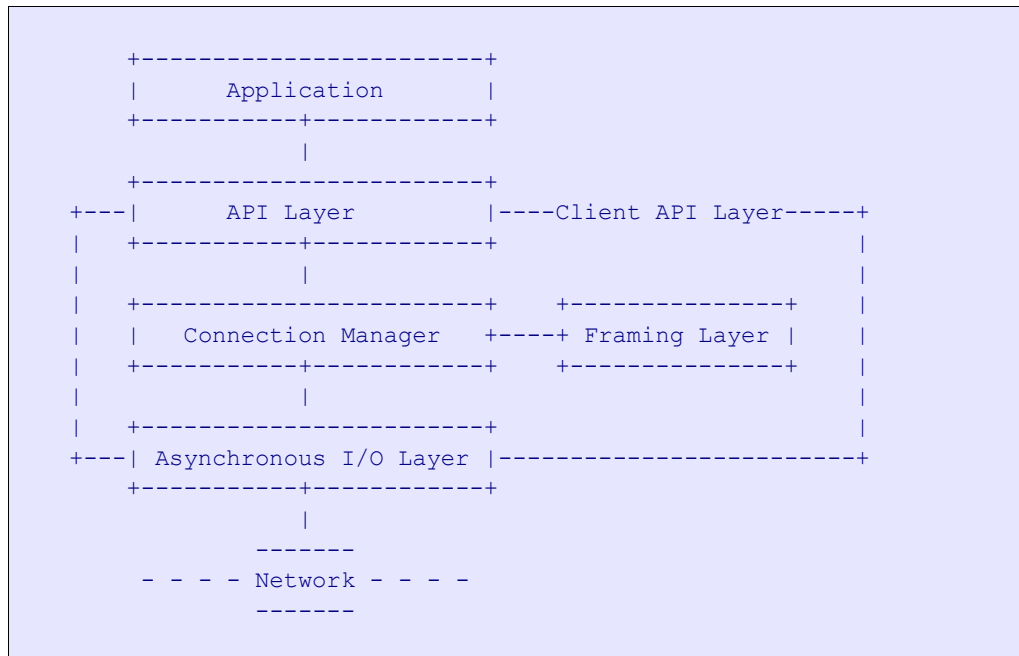
8 The recommended AMQP client architecture consists of several layers of abstraction:

- 9 1. A **framing layer**. This layer takes AMQP protocol methods, in some language-specific format
10 (structures, classes, etc.) and serialises them as wire-level frames. The framing layer can be
11 mechanically generated from the AMQP specifications (which are defined in a protocol modelling
12 language, implemented in XML and specifically designed for AMQP).
- 13 2. A **connection manager layer**. This layer reads and writes AMQP frames and manages the overall
14 connection and session logic. In this layer we can encapsulate the full logic of opening a connection
15 and session, error handling, content transmission and reception, and so on. Large parts of this layer
16 can be produced automatically from the AMQP specifications. For instance, the specifications define
17 which methods carry content, so the logic "send method and then optionally send content" can be
18 produced mechanically.
- 19 3. An **API layer**. This layer exposes a specific API for applications to work with. The API layer may
20 reflect some existing standard, or may expose the high-level AMQP methods, making a mapping as
21 described earlier in this section. The AMQP methods are designed to make this mapping both simple
22 and useful. The API layer may itself be composed of several layers, e.g. a higher-level API constructed
23 on top of the AMQP method API.
- 24 4. A **transaction processing layer**. This layer drives the application by delivering it transactions to
25 process, where the transactions are middleware messages. Using a transaction layer can be very
26 powerful because the middleware becomes entirely hidden, making applications easier to build, test,
27 and maintain.

28 Additionally, there is usually some kind of I/O layer, which can be very simple (synchronous socket reads
29 and writes) or sophisticated (fully asynchronous multi-threaded i/o).

30 This diagram shows the overall recommended architecture (without layer 4, which is a different story):

1
2
3
4
5
6
7
8
9
10
11
12
13
14
15
16
17
18
19
20
21



22
23
24
25

In this document, when we speak of the "client API", we mean all the layers below the application (i/o, framing, connection manager, and API layers). We will usually speak of "the client API" and "the application" as two separate things, where the application uses the client API to talk to the middleware server.

3 Functional Specification

3.1 Server Functional Specification

3.1.1 Messages and Content

A message is the atomic unit of processing of the middleware routing and queuing system. Messages carry a content, which consists of a content header, holding a set of properties, and a content body, holding an opaque block of binary data. Contents can themselves contain child contents, to any level of complexity.

A message can correspond to many different application entities:

- ◆ An application-level message
- ◆ A file to transfer
- ◆ One frame of a data stream
- ◆ etc.

AMQP defines a set of "content classes", each implementing a specific content syntax (the set of content header properties) and semantics (the methods that are available to manipulate messages of that content class).

Messages may be persistent, according to the semantics of each class. A persistent message is held securely on disk and guaranteed to be delivered even if there is a serious network failure, server crash, overflow etc.

Messages may have a priority level, according to the semantics of each class. A high priority message is sent ahead of lower priority messages waiting in the same message queue. When messages must be discarded in order to maintain a specific service quality level the server will first discard low-priority messages.

The server **MUST NOT** modify message content bodies that it receives and passes to consumer applications. The server **MAY** add information to content headers but it **MUST NOT** remove or modify existing information.

3.1.2 Virtual Hosts

A Virtual Host¹ is a data partition within the server, it is an administrative convenience which will prove useful to those wishing to provide AMQP as a service on a shared infrastructure.

A virtual host comprises its own name space, a set of exchanges, message queues, and all associated objects. Each connection **MUST BE** associated with a single virtual host.

The client selects the virtual host in the Connection.Open method, after authentication. This implies that the authentication scheme of the server is shared between all virtual hosts on that server. However, the

¹ The term Virtual Host is taken from the use popularised by the Apache HTTP server. Apache Virtual Hosts enable Internet Service providers to provide bulk hosting from one shared server infrastructure. We hope that the inclusion of this capability within AMQP opens up similar opportunities to larger organisations.

1 authorization scheme used MAY be unique to each virtual host. This is intended to be useful for shared
2 hosting infrastructures. Administrators who need different authentication schemes for each virtual host
3 should use separate servers.

4 All channels within the connection work with the same virtual host. There is no way to communicate with
5 a different virtual host on the same connection, nor is there any way to switch to a different virtual host
6 without tearing down the connection and beginning afresh.

7 The protocol offers no mechanisms for creating or configuring virtual hosts - this is done in an undefined
8 manner within the server and is entirely implementation-dependent.

9 3.1.3 Exchanges

10 An exchange is a message routing agent within a virtual host. An exchange instance (which we commonly
11 call "an exchange") accepts messages and routing information - principally a routing key - and either
12 passes the messages to message queues, or to internal services. Exchanges are named on a per-virtual
13 host basis.

14 Applications can freely create, share, use, and destroy exchange instances, within the limits of their
15 authority.

16 Exchanges may be durable, temporary, or auto-deleted. Durable exchanges last until they are deleted.
17 Temporary exchanges last until the server shuts-down. Auto-deleted exchanges last until they are no
18 longer used.

19 The server provides a specific set of exchange types. Each exchange type implements a specific matching
20 and algorithm, as defined in the next section. AMQP mandates a small number of exchange types, and
21 recommends some more. Further, each server implementation may add its own exchange types.

22 An exchange can route a single message to many message queues in parallel. This creates multiple
23 instances of the message that are consumed independently.

24 3.1.3.1 The Direct Exchange Type

25 The direct exchange type works as follows:

- 26 1. A message queue binds to the exchange using a routing key, K.
- 27 2. A publisher sends the exchange a message with the routing key R.
- 28 3. The message is passed to the message queue if $K = R$.

29 Note that message queues can bind using any valid routing key value, but most often message queues will
30 bind using their own name as routing key.

31 A default binding for each message queue MUST BE made like this, using the message queue name.

32 One suggested design for the direct exchange is a lookup table that allows a message routing key to be
33 rapidly mapped to a list of message queues. This exchange type, and a pre-declared exchange called
34 `amq.direct`, are mandatory.

35 The server MUST implement the direct exchange type and MUST pre-declare within each virtual host at
36 least two direct exchanges: one named `amq.direct`, and one with **no public name** that serves as the default
37 exchange for Publish methods.

3.1.3.2 The Fanout Exchange Type

The fanout exchange type works as follows:

1. A message queue binds to the exchange with no arguments.
2. A publisher sends the exchange a message.
3. The message is passed to the message queue unconditionally.

The fanout exchange is trivial to design and implement. This exchange type, and a pre-declared exchange called **amq.fanout**, are mandatory.

3.1.3.3 The Topic Exchange Type

The topic exchange type works as follows:

1. A message queue binds to the exchange using a routing pattern, P.
2. A publisher sends the exchange a message with the routing key R.
3. The message is passed to the message queue if R matches P.

The routing key used for a topic exchange MUST consist of words delimited by dots. Each word may contain the letters A-Z and a-z and digits 0-9.

The routing pattern follows the same rules as the routing key with the addition that * matches a single word, and # matches zero or more words. Thus the routing pattern *.stock.# matches the routing keys usd.stock and eur.stock.db but not stock.nasdaq.

One suggested design for the topic exchange is to hold the set of all known routing keys, and update this when publishers use new routing keys. It is possible to determine all bindings for a given routing key, and so to rapidly find the message queues for a message. This exchange type is optional.

The server SHOULD implement the topic exchange type and in that case, the server MUST pre-declare within each virtual host at least one topic exchange, named **amq.topic**.

3.1.3.4 The System Exchange Type

The system exchange type works as follows:

1. A publisher sends the exchange a message with the routing key S.
2. The system exchange passes this to a system service S.

System services starting with "amq." are reserved for AMQP usage. All other names may be used freely on by server implementations. This exchange type is optional.

3.1.3.5 Implementation-defined Exchange Types

All non-normative exchange types MUST be named starting with "x-". Exchange types that do not start with "x-" are reserved for future use in the AMQP standard.

3.1.4 Message Queues

A message queue is a named FIFO buffer that holds message on behalf of a set of consumer applications. Applications can freely create, share, use, and destroy message queues, within the limits of their authority.

Note that in the presence of multiple readers from a queue, or client transactions, or use of priority fields, or use of message selectors, or implementation-specific delivery optimisations the queue MAY NOT exhibit true FIFO characteristics. The only way to guarantee FIFO is to have just one consumer connected to a queue. The queue may be described as “weak-FIFO” in these cases.

Message queues may be durable, temporary, or auto-deleted. Durable message queues last until they are deleted. Temporary message queues last until the server shuts-down. Auto-deleted message queues last until they are no longer used.

Message queues hold their messages in memory, on disk, or some combination of these. Message queues are named on a per-virtual host basis.

Message queues hold messages and distribute them between one or more consumer clients. A message routed to a message queue is never sent to more than one client unless it is being resent after a failure or rejection.

A single message queue can hold different types of content at the same time and independently. That is, if Basic and File contents are sent to the same message queue, these will be delivered to consuming applications independently as requested.

3.1.5 Bindings

A binding is a relationship between a message queue and an exchange. The binding specifies routing arguments that tell the exchange which messages the queue should get.

Applications create and destroy bindings as needed to drive the flow of messages into their message queues. The lifespan of bindings depend on the message queues they are defined for - when a message queue is destroyed, its bindings are also destroyed.

The specific semantics of the Queue.Bind method depends on the exchange type.

3.1.6 Consumers

We use the term "consumer" to mean both the client application and the entity that controls how a specific client application receives messages off a message queue. When the client "starts a consumer" it creates a consumer entity in the server. When the client "cancels a consumer" it destroys a consumer entity in the server.

Consumers belong to a single client channel and cause the message queue to send messages asynchronously to the client.

3.1.7 Quality of Service

The quality of service controls how fast messages are sent. The quality of service depends on the type of content being distributed. For basic messaging, for file transfer, and for streaming, we define different quality of service semantics.

1 In general the quality of service uses the concept of "pre-fetch" to specify how many messages or how
2 many octets of data will be sent before the client acknowledges a message. The goal is to send message
3 data in advance, to reduce latency.

4 3.1.8 Acknowledgements

5 An acknowledgement is a formal signal from the client application to a message queue that it has
6 successfully processed a message. There are two possible acknowledgement models:

- 7 1. Automatic, in which the server removes a content from a message queue as soon as it delivers it to an
8 application (via the Deliver or Get-Ok methods).
- 9 2. Explicit, in which the client application must send an Ack method for each message, or batch of
10 messages, that it has processed.

11 The client layers can themselves implement explicit acknowledgements in different ways, e.g. as soon as a
12 message is received, or when the application indicates that it has processed it. These differences do not
13 affect AMQP or interoperability.

14 3.1.9 Flow Control

15 Flow control is an emergency procedure used to halt the flow of messages from a peer. It works in the
16 same way between client and server and is implemented by the Channel.Flow command. Flow control is
17 the only mechanism that can stop an over-producing publisher. A consumer can use the more elegant
18 mechanism of pre-fetch windowing, if it uses message acknowledgements (which usually means using
19 transactions).

20 3.1.10 Naming Conventions

21 These conventions govern the naming of AMQP entities. The server and client MUST respect these
22 conventions:

- 23 ♦ User defined exchange types MUST be prefixed by "x-"
- 24 ♦ Standard exchange instances are prefixed by "amq."
- 25 ♦ Standard system services are prefixed by "amq."
- 26 ♦ Standard message queues are prefixed by "amq."
- 27 ♦ All other exchange, system service, and message queue names are in application space.

28 3.2 AMQP Command Specification (Classes & Methods)

29 3.2.1 Explanatory Notes

30 The AMQP methods may define specific minimal values (such as numbers of consumers per message
31 queue) for interoperability reasons. These minima are defined in the description of each class.

32 Note conforming AMQP implementations SHOULD implement reasonably generous values for such fields,
33 the minima is only intended for use on the least capable platforms.

34 The grammars use this notation:

- 1 ◆ 'S:' indicates data or a method sent from the server to the client
- 2 ◆ 'C:' indicates data or a method sent from the client to the server
- 3 ◆ +term or +(…) expression means '1 or more instances'
- 4 ◆ *term or *(…) expression means 'zero or more instances'.

5 We define methods as being either:

- 6 ◆ a synchronous request ("syn request"). The sending peer SHOULD wait for the specific reply method, but MAY implement this asynchronously
- 7 ◆ a synchronous reply ("syn reply for XYZ")
- 8 ◆ an asynchronous request or reply ("async").

10 3.2.2 Class and Method Ids

11 These are the AMQP class and method ids. Note that these may change in new versions of AMQP and
12 implementors are strongly recommended to use the AMQP class specifications as a source for the class
13 and method ids rather than hard-coding these values.

14 These are the ID values for each class:

15 Connection 10

16	Channel	20
17	Access	30
18	Exchange	40
19	Queue	50
20	Basic	60
21	File	70
22	Stream	80
23	Tx	90
24	Dtx	100
25	Tunnel	110

26 These are the ID values for the Connection methods:

27	Connection.Start	10
28	Connection.Start_Ok	11
29	Connection.Secure	20
30	Connection.Secure_Ok	21
31	Connection.Tune	30
32	Connection.Tune_Ok	31
33	Connection.Open	40
34	Connection.Open_Ok	41
35	Connection.Redirect	50
36	Connection.Close	60
37	Connection.Close_Ok	61

38 These are the ID values for the Channel methods:

1	Channel.Open	10
2	Channel.Open_Ok	11
3	Channel.Flow	20
4	Channel.Flow_Ok	21
5	Channel.Alert	30
6	Channel.Close	40
7	Channel.Close_Ok	41

8 These are the ID values for the Access methods:

9	Access.Request	10
10	Access.Request_Ok	11

11 These are the ID values for the Exchange methods:

12	Exchange.Declare	10
13	Exchange.Declare_Ok	11
14	Exchange.Delete	20
15	Exchange.Delete_Ok	21

16 These are the ID values for the Queue methods:

17	Queue.Declare	10
18	Queue.Declare_Ok	11
19	Queue.Bind	20
20	Queue.Bind_Ok	21
21	Queue.Purge	30
22	Queue.Purge_Ok	31
23	Queue.Delete	40
24	Queue.Delete_Ok	41

25 These are the ID values for the Basic methods:

26	Basic.Qos	10
27	Basic.Qos_Ok	11
28	Basic.Consume	20
29	Basic.Consume_Ok	21
30	Basic.Cancel	30
31	Basic.Cancel_Ok	31
32	Basic.Publish	40
33	Basic.Return	50
34	Basic.Deliver	60
35	Basic.Get	70
36	Basic.Get_Ok	71
37	Basic.Get_Empty	72
38	Basic.Ack	80
39	Basic.Reject	90

40 These are the ID values for the File methods:

1	File.Qos	10
2	File.Qos_Ok	11
3	File.Consume	20
4	File.Consume_Ok	21
5	File.Cancel	30
6	File.Cancel_Ok	31
7	File.Open	40
8	File.Open_Ok	41
9	File.Stage	50
10	File.Publish	60
11	File.Return	70
12	File.Deliver	80
13	File.Ack	90
14	File.Reject	100

15 These are the ID values for the Stream methods:

16	Stream.Qos	10
17	Stream.Qos_Ok	11
18	Stream.Consume	20
19	Stream.Consume_Ok	21
20	Stream.Cancel	30
21	Stream.Cancel_Ok	31
22	Stream.Publish	40
23	Stream.Return	50
24	Stream.Deliver	60

25 These are the ID values for the Tx methods:

26	Tx.Select	10
27	Tx.Select_Ok	11
28	Tx.Commit	20
29	Tx.Commit_Ok	21
30	Tx.Rollback	30
31	Tx.Rollback_Ok	31

32 These are the ID values for the Dtx methods:

33	Dtx.Select	10
34	Dtx.Select_Ok	11
35	Dtx.Start	20
36	Dtx.Start_Ok	21

37 These are the ID values for the Tunnel methods:

38	Tunnel.Request	10
----	----------------	----

39

40 [TODO: JOH: INSERT GENERATED XML DOCUMENTATION HERE]

41

4 Technical Specifications

4.1 IANA Assigned Port Number

The standard AMQP port number has been assigned by IANA as 5672 for both TCP and UDP.

The UDP port will be used in a future multi-cast implementation.

4.2 AMQP Wire-Level Format

4.2.1 Format Protocol Grammar

We provide a complete grammar for AMQP (this is provided for reference, and you may find it more interesting to skip through to the next sections that detail the different frame types and their formats):

```

1  amqp                = protocol-header *amqp-unit
2
3  protocol-header    = literal-AMQP protocol-id protocol-version
4  literal-AMQP       = %d65.77.81.80           ; "AMQP"
5  protocol-id        = %d1.1                 ; AMQP over TCP/IP
6  protocol-version   = %d9.1                 ; 0.9 revision 1
7
8  amqp-unit          = method | oob-method | trace | heartbeat
9
10 method              = method-frame [ content ]
11 method-frame        = %d1 frame-properties method-payload frame-end
12 frame-properties    = cycle channel payload-size
13 cycle               = OCTET
14 channel             = short-integer          ; Non-zero
15 payload-size        = long-integer
16 method-payload      = class-id method-id *amqp-field
17 class-id            = %x00.01-%xFF.FF
18 method-id           = %x00.01-%xFF.FF
19 amqp-field          = BIT / OCTET / short-integer / long-integer
20                    / long-long-integer
21                    / short-string / long-string
22                    / timestamp
23                    / field-table
24 short-integer       = 2*OCTET
25 long-integer        = 4*OCTET
26 long-long-integer   = 8*OCTET
27 short-string        = OCTET *string-char      ; length + content
28 string-char         = %x01 .. %xFF
29 long-string         = long-integer *OCTET      ; length + content
30 timestamp           = long-long-integer
31 field-table         = long-integer *field-value-pair
32 field-value-pair    = field-name field-value
33 field-name          = short-string
34 field-value         = 'S' long-string
35                    / 'I' signed-integer
36                    / 'D' decimal-value
37                    / 'T' timestamp
38                    / 'F' field-table
39 signed-integer      = 4*OCTET
40 decimal-value       = decimals long-integer
41 decimals            = OCTET
42 frame-end           = %xCE
43
44 content             = %d2 content-header child-content
45                    / *content-body
46 content-header      = frame-properties header-payload frame-end
47 header-payload      = content-class content-weight content-body-size
48                    / property-flags property-list
49 content-class       = OCTET
50 content-weight      = OCTET
51 content-body-size   = long-long-integer
52 property-flags      = 15*BIT %b0 / 15*BIT %b1 property-flags

```

```

1  property-list      = amqp-field
2  child-content     = content-weight*content
3  content-body      = %d3 frame-properties body-payload frame-end
4  body-payload      = *OCTET
5
6  oob-method        = oob-method-frame [ oob-content ]
7  oob-method-frame  = %d4 frame-properties frame-end
8  oob-content       = %d5 content-header oob-child-content
9                    *oob-content-body
10 oob-child-content = content-weight*oob-content
11 oob-content-body  = %d6 frame-properties frame-end
12
13 trace             = %d7 cycle %d0 payload-size trace-payload
14                  / frame-end
15 trace-payload     = *OCTET
16
17 heartbeat         = %d8 cycle %d0 %d0 frame-end

```

We use the Augmented BNF syntax defined in IETF RFC 2234. In summary,

- ◆ The name of a rule is simply the name itself.
- ◆ Terminals are specified by one or more numeric characters with the base interpretation of those characters indicated as 'd' or 'x'.
- ◆ A rule can define a simple, ordered string of values by listing a sequence of rule names.
- ◆ A range of alternative numeric values can be specified compactly, using dash ("-") to indicate the range of alternative values.
- ◆ Elements enclosed in parentheses are treated as a single element, whose contents are strictly ordered.
- ◆ Elements separated by forward slash ("/") are alternatives.
- ◆ The operator "*" preceding an element indicates repetition. The full form is: "<a>*element", where <a> and are optional decimal values, indicating at least <a> and at most occurrences of element.
- ◆ A rule of the form: "<n>element" is equivalent to <n>*<n>element.
- ◆ Square brackets enclose an optional element sequence.

4.2.2 Protocol Header

The client **MUST** start a new connection by sending a protocol header.

This is an 8-octet sequence:

```

+---+---+---+---+---+---+---+---+
|'A'|'M'|'Q'|'P'| 1 | 1 | 9 | 1 |
+---+---+---+---+---+---+---+---+
                        8 octets

```

The protocol header consists of the upper case letters "AMQP" followed by:

1. The protocol class, which is 1 (for all AMQP protocols).
2. The protocol instance, which is 1 (for AMQP over TCP/IP).
3. The protocol major version, which is 9 (version 1.0 is 10, highest possible release is 25.5).

4. The protocol minor version, which is currently 1.

The protocol negotiation model is compatible with existing protocols such as HTTP that initiate a connection with an constant text string, and with firewalls that sniff the start of a protocol in order to decide what rules to apply to it.

An AMQP client and server agree on a protocol and version as follows:

- ◆ The client opens a new socket connection to the AMQP server and sends the protocol header.
- ◆ The server either accepts or rejects the protocol header. If it rejects the protocol header writes a valid protocol header to the socket and then closes the socket.
- ◆ Otherwise it leaves the socket open and implements the protocol accordingly.

Examples:


Client sends:	Server responds:
AMQP%d1.1.9.1	Connection.Start method
AMQP%d2.0.1.1	AMQP%d1.1.9.1<Close connection>
HTTP	AMQP%d1.1.9.1<Close connection>

Guidelines for implementers:

- ◆ An AMQP server **MUST** accept the AMQP protocol as defined by class = 1, instance = 1. Conformance test: amq_wlp_header_01.
- ◆ An AMQP server **MAY** accept non-AMQP protocols such as HTTP. Conformance test: amq_wlp_header_02.
- ◆ If the server does not recognise the first 4 octets of data on the socket, or does not support the specific protocol version that the client requests, it **MUST** write a valid protocol header to the socket, then flush the socket (to ensure the client application will receive the data) and then close the socket connection. The server **MAY** print a diagnostic message to assist debugging. Conformance test: amq_wlp_header_03.
- ◆ An AMQP client **MAY** detect the server protocol version by attempting to connect with its highest supported version and reconnecting with a lower version if it receives such information back from the server. Conformance test: amq_wlp_header_04.

4.2.3 General Frame Format

All frames start with an 8-octet header composed of a type field (octet), a cycle field (octet), a channel field (short integer) and a size field (long integer):



AMQP defines these frame types:

- ◆ Type = 1, "METHOD": method frame.
- ◆ Type = 2, "HEADER": content header frame.
- ◆ Type = 3, "BODY": content body frame.

- 1 ◆ Type = 4, "OOB-METHOD": out-of-band method frame.
- 2 ◆ Type = 5, "OOB-HEADER": out-of-band band header frame.
- 3 ◆ Type = 6, "OOB-BODY": out-of-band body frame.
- 4 ◆ Type = 7, "TRACE": trace frame.
- 5 ◆ Type = 8, "HEARTBEAT": heartbeat frame.

6 The channel number is 0 for all frames which are global to the connection and 1-65535 for frames that refer to specific channels.

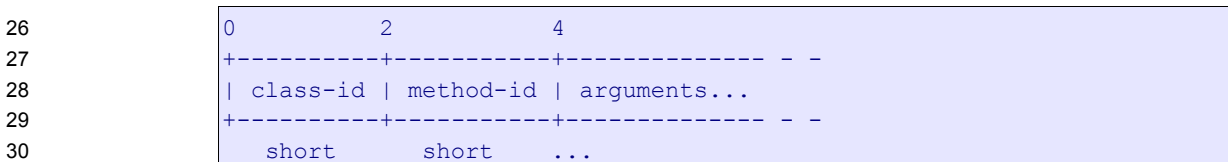
8 The size field is the size of the payload, excluding the frame-end octet. While AMQP assumes a reliable connected protocol, we use the frame end to detect framing errors caused by incorrect client or server implementations.

11 Guidelines for implementers:

- 12 ◆ If a peer receives a frame with a type that is not one of these defined types, it MUST treat this as a fatal protocol error and close the connection without sending any further data on it. Conformance test: amq_wlp_frame_01.
- 13 ◆ When a peer reads a frame it MUST check that the frame-end is valid before attempting to decode the frame. If the frame-end is not valid it MUST treat this as a fatal protocol error and close the connection without sending any further data on it. It SHOULD log information about the problem, since this indicates an error in either the server or client framing code implementation. Conformance test: amq_wlp_frame_02.
- 14 ◆ A peer MUST NOT send frames larger than the agreed-upon size. A peer that receives an oversized frame MUST signal a connection exception with reply code 501 (frame error). Conformance test: amq_wlp_frame_03.

23 4.2.4 Method Frames

24 Method frame bodies consist of an invariant list of data fields, called "arguments". All method bodies start with identifier numbers for the class and method:



31 Guidelines for implementers:

- 32 ◆ The class-id and method-id are constants that are defined in the AMQP class and method specifications.
- 33 ◆ The arguments are a set of AMQP fields that specific to each method.
- 34 ◆ Class id values from %x00.01-%xEF.FF are reserved for AMQP standard classes.
- 35 ◆ Class id values from %xF0.00-%xFF.FF (%d61440-%d65535) may be used by implementations for non-standard extension classes.

4.2.5 AMQP Data Fields

4.2.5.1 Integers

AMQP defines these integer types:

- ◆ Unsigned octet (8 bits).
- ◆ Unsigned short integers (16 bits).
- ◆ Unsigned long integers (32 bits).
- ◆ Unsigned long long integers (64 bits).

Integers and string lengths are always unsigned and held in network byte order. We make no attempt to optimise the case when two low-high systems (e.g. two Intel CPUs) talk to each other.

Guidelines for implementers:

- ◆ Implementers **MUST NOT** assume that integers encoded in a frame are aligned on memory word boundaries.

4.2.5.2 Bits

Bits are accumulated into whole octets. When two or more bits are contiguous in a frame these will be packed into one or more octets, starting from the low bit in each octet. There is no requirement that all the bit values in a frame be contiguous, but this is generally done to minimise frame sizes.

4.2.5.3 Strings

AMQP strings are variable length and represented by an integer length followed by zero or more octets of data. AMQP defines two string types:

- ◆ Short strings, stored as an 8-bit unsigned integer length followed by zero or more octets of data. Short strings can carry up to 255 octets of UTF-8 data, but may not contain binary zero octets.
- ◆ Long strings, stored as a 32-bit unsigned integer length followed by zero or more octets of data. Long strings can contain any data.

4.2.5.4 Timestamps

Time stamps are held in the 64-bit POSIX `time_t` format with an accuracy of one second. By using 64 bits we avoid future wraparound issues associated with 31-bit and 32-bit `time_t` values.

4.2.5.5 Field Tables

Field tables are long strings that contain packed name-value pairs. Each name-value pair is a structure that provides a field name, a field type, and a field value. A field can hold a tiny text string, a long string, a long signed integer, a decimal, a date and/or time, or another field table.

Guidelines for implementers:

- ◆ Field names **MUST** start with a letter, '\$' or '#' and may continue with letters, '\$' or '#', digits, or underlines, to a maximum length of 128 characters.
- ◆ The server **SHOULD** validate field names and upon receiving an invalid field name, it **SHOULD** signal a connection exception with reply code 503 (syntax error). Conformance test: amq_wlp_table_01.
- ◆ Specifically and only in field tables, integer values are signed (31 bits plus sign bit).
- ◆ Decimal values are not intended to support floating point values, but rather business values such as currency rates and amounts. The 'decimals' octet is not signed.
- ◆ A peer **MUST** handle duplicate fields by using only the first instance.

4.2.5.6 Content Framing

Certain specific methods (Publish, Deliver, etc.) carry content. Please refer to the chapter "Functional Specifications" for specifications of each method, and whether or not the method carries content. Methods that carry content do so unconditionally.

Content consists of a list of 1 or more frames as follows:

1. Exactly one content header frame that provides properties for the content.
2. Optionally, one or more child contents. A child content follows the exact rules for a content. Contents may thus be structured in a hierarchy to any level.
3. Optionally, one or more content body frames.

Content frames on a specific channel form an strict list. That is, they may be mixed with frames for different channels, but two contents may not be mixed or overlapped on a single channel, nor may content frames for a single content be mixed with method frames on the same channel.


Note that any non-content frame explicitly marks the end of the content.

Guidelines for implementers:

- ◆ A peer that receives an incomplete content **MUST** raise a connection exception with reply code 501 (frame error). Conformance test: amq_wlp_content_01.

4.2.5.7 The Content Header

A content header payload has this format:



Guidelines for implementers:

- ◆ The content class-id **MUST** match the method frame class id. The peer **MUST** respond to an invalid content class-id by raising a connection exception with reply code 501 (frame error). Conformance test: amq_wlp_content_02.
- ◆ The weight field specifies the number of child-contents that the content contains. This is zero for simple contents and non-zero for structured contents (explained below).

- ◆ The body size is a 64-bit value that defines the total size of the content body. It may be zero, indicating that there will be no content body frames.
- ◆ The property flags are an array of bits that indicate the presence or absence of each property value in sequence. The bits are ordered from most high to low - bit 15 indicates the first property.
- ◆ The property flags can specify more than 16 properties. If the last bit (0) is set, this indicates that a further property flags field follows. There are many property flags fields as needed.
- ◆ The property values are class-specific AMQP data fields.
- ◆ Bit properties are indicated ONLY by their respective property flag (1 or 0) and are never present in the property list.
- ◆ The channel number in content frames MUST NOT be zero. A peer that receives a zero channel number in a content frame MUST signal a connection exception with reply code 504 (channel error).
Conformance test: amq_wlp_content_03.

4.2.5.8 The Content Body

The content body payload is an opaque binary block followed by a frame end octet¹:

```
+-----+ +-----+
| Opaque binary payload | | frame-end |
+-----+ +-----+
```

The content body can be split into as many frames as needed. The maximum size of the frame payload is agreed upon by both peers during connection negotiation.

Guidelines for implementers:

- ◆ A peer MUST handle a content body that is split into multiple frames by storing these frames as a single set, and either retransmitting them as-is, broken into smaller frames, or concatenated into a single block for delivery to an application.

4.2.5.9 Structured Content

A structured content consists of a single top level content and multiple child contents, as complex as needed by the application. Structured contents form a hierarchy, a tree with a single root.

At any level of this tree, the weight field in the content header indicates whether the content has child contents or not. If the content has child contents, these follow immediately after the header and before the body of the parent content, e.g.:

```
[parent-header weight = 1]
  [child-header weight = 0] [child-body]
[parent-body]
```

The weight is the number of child contents at the current level.

Guidelines for implementers:

¹ Strictly this is redundant, however it does make debugging both protocol network streams and memory buffers somewhat easier.

- 1 ◆ The peer MAY support structured contents. If it does not support structured contents it MUST
2 respond to a structured content by raising a connection exception with reply code 540 (not
3 implemented). Conformance test: amq_wlp_content_04.
- 4 ◆ The peer MUST correctly detect a mismatch between the content weight and the frames that follow,
5 and report such a mismatch by raising a connection exception with reply code 501 (frame error).
6 Conformance test: amq_wlp_content_05.

7 4.2.5.10 Out-Of-Band Frames

8 The formatting of out-of-band frames follows the same specifications as for normal frames, with the
9 exception that frame payloads are sent via some unspecified transport mechanism. This could be shared
10 memory, specialised network protocols, etc.

11 The actual out-of-band transport used, and its configuration, is specified in the Channel.Open method.

12 4.2.5.11 Trace Frames

13 Trace frames are intended for a "trace handler" embedded in the recipient peer. The significance and
14 implementation of the trace handler is implementation-defined.

15 Guidelines for implementers:

- 16 ◆ Trace frames MUST have a channel number of zero. A peer that receives an invalid trace frame MUST
17 raise a connection exception with reply code 501 (frame error). Conformance test: amq_wlp_trace_01.
- 18 ◆ If the recipient of a trace frame does not have a suitable trace handler, it MUST discard the trace frame
19 without signalling any error or fault. Conformance test: amq_wlp_trace_02.

20 4.2.5.12 Heartbeat Frames


21 Heartbeat frames tell the recipient that the sender is still alive. The rate and timing of heartbeat frames is
22 negotiated during connection tuning.

23 Guidelines for implementers:

- 24 ◆ Heartbeat frames MUST have a channel number of zero. A peer that receives an invalid trace frame
25 MUST raise a connection exception with reply code 501 (frame error). Conformance test:
26 amq_wlp_heartbeat_01.
- 27 ◆ If the peer does not support heartbeating it MUST discard the heartbeat frame without signalling any
28 error or fault. Conformance test: amq_wlp_heartbeat_02.

29 4.3 Channel Multiplexing

30 AMQP permits peers to create multiple independent threads of control. Each channel acts as a virtual
31 connection that share a single socket:



8 Guidelines for implementers:

- 9 ♦ An AMQP peer SHOULD support multiple channels. The maximum number of channels is defined at connection negotiation, and a peer MAY negotiate this down to 1¹.
- 10
- 11 ♦ Each peer SHOULD balance the traffic on all open channels in a fair fashion. This balancing can be
- 12 done on a per-frame basis, or on the basis of amount of traffic per channel. A peer SHOULD NOT
- 13 allow one very busy channel to starve the progress of a less busy channel.

14 4.4 Error Handling

15 4.4.1 Exceptions

16 Using the standard 'exception' programming model, AMQP does not signal success, only failure.

17 AMQP defines two exception levels²:

- 18 1. **Channel exceptions.** These close the channel that caused the error. Channel exceptions are usually
- 19 due to 'soft' errors that do not affect the rest of the application.
- 20 2. **Connection exceptions.** These close the socket connection and are usually due to 'hard' errors that
- 21 indicate a programming fault, a bad configuration, or other case that needs intervention.

22 We document the assertions formally in the definition of each class and method.

23 4.4.2 Reply Code Format

24 We use the IETF standard format for reply codes as described in IETF RFC 821. A reply code uses three

25 digits, and the first digit provides the main feedback as to whether and how an operation completed. The

26 second and third digits provide additional information. The reply codes can be processed by client

27 applications without full knowledge of their meaning.

28 We use a standard 3-digit reply code. The first digit (the completion indicator) reports whether the request

29 succeeded or not:

- 30 1: Ready to be performed, pending some confirmation.

¹ It is expected that all but the most simplistic client or server implementation will support several channels active on each connection simultaneously and that the best implementations will support hundreds of channels in one connection should a client application require it.

² The severity of these exceptions may surprise the reader, however it is a requirement of AMQP that the system either works predictably, or not at all – to this end, fail fast and fail early will have the effect of achieving rapid convergence in the quality and interoperability of this standard as bugs and incompatibilities will be discovered quickly and corrected.

1 2: Successful.

2 3: Ready to be performed, pending more information.

3 4: Failed, but may succeed later.

4 5: Failed, requires intervention.

5 6-9: Reserved for future use.

6 The second digit (the category indicator) provides more information on failures:

7 0: Error in syntax.

8 1: The reply provides general information.

9 2: Problem with session or connection.

10 3: Problem with security.

11 4: Problem with implementation.

12 5-9: Reserved for future use.

13 The third digit (the instance indicator) distinguishes among different situations with the same
14 completion/category.

15 4.4.3 Channel Exception Reply Codes

16 When the server raises a channel exception it may use one of the following reply codes. These are all
17 associated with failures that affect the current channel but not other channels in the same connection:

- 18 ◆ 310=NOT_DELIVERED: The client asked for a specific message that is no longer available. The
19 message was delivered to another client, or was purged from the queue for some other reason.
- 20 ◆ 311=CONTENT_TOO_LARGE: The client attempted to transfer content larger than the server could
21 accept at the present time. The client may retry at a later time.
- 22 ◆ 403=ACCESS_REFUSED: The client attempted to work with a server entity to which it has no access
23 due to security settings.
- 24 ◆ 404=NOT_FOUND: The client attempted to work with a server entity that does not exist.
- 25 ◆ 405=RESOURCE_LOCKED: The client attempted to work with a server entity to which it has no
26 access because another client is working with it.

27 4.4.4 Connection Exception Reply Codes

28 When the server raises a connection exception it may use one of the following reply codes. These are all
29 associated with failures that preclude any further activity on the connection:

- 30 ◆ 320=CONNECTION_FORCED: An operator intervened to close the connection for some reason. The
31 client may retry at some later date.
- 32 ◆ 402=INVALID_PATH: The client tried to work with an unknown virtual host or cluster.

- 1 ◆ 501=FRAME_ERROR: The client sent a malformed frame that the server could not decode. This
2 strongly implies a programming error in the client.
- 3 ◆ 502=SYNTAX_ERROR: The client sent a frame that contained illegal values for one or more fields.
4 This strongly implies a programming error in the client.
- 5 ◆ 503=COMMAND_INVALID: The client sent an invalid sequence of frames, attempting to perform an
6 operation that was considered invalid by the server. This usually implies a programming error in the
7 client.
- 8 ◆ 504=CHANNEL_ERROR: The client attempted to work with a channel that had not been correctly
9 opened. This most likely indicates a fault in the client layer.
- 10 ◆ 506=RESOURCE_ERROR: The server could not complete the method because it lacked sufficient
11 resources. This may be due to the client creating too many of some type of entity.
- 12 ◆ 530=NOT_ALLOWED: The client tried to work with some entity in a manner that is prohibited by the
13 server, due to security settings or by some other criteria.
- 14 ◆ 540=NOT_IMPLEMENTED: The client tried to use functionality that is not implemented in the server.
- 15 ◆ 541=INTERNAL_ERROR: The server could not complete the method because of an internal error. The
16 server may require intervention by an operator in order to resume normal operations.

17 4.5 Limitations

18 The AMQP specifications impose these limits on future extensions of AMQP or protocols from the same
19 wire-level format:

- 20 ◆ Number of channels per connection: 16-bit channel number.
- 21 ◆ Number of protocol classes: 16-bit class id.
- 22 ◆ Number of methods per protocol class: 16-bit method id.

23 The AMQP specifications impose these limits on data:

- 24 ◆ Maximum size of a short string: 255 octets.
- 25 ◆ Maximum size of a long string or field table: 32-bit size.
- 26 ◆ Maximum size of a frame payload: 32-bit size.
- 27 ◆ Maximum size of a content: 64-bit size.
- 28 ◆ Maximum depth of a structured content: unlimited.
- 29 ◆ Maximum weight of a structured content: 16-bit weight.

30 An AMQP server or client implementation will also impose its own limits on resources such as number of
31 simultaneous connections, number of consumers per channel, number of queues, etc. These do not affect
32 interoperability and are not specified.

33 4.6 Security

34 4.6.1 Goals and Principles

35 We guard against buffer-overflow exploits by using length-specified buffers in all places. All externally-
36 provided data can be verified against maximum allowed lengths whenever any data is read.

1 Invalid data can be handled unambiguously, by closing the channel or the connection.

2 4.6.2 Denial of Service Attacks

3 AMQP handles errors by returning a reply code and then closing the channel or connection. This avoids
4 ambiguous states after errors.

5 It should be assumed that exceptional conditions during connection negotiation stage are due to an hostile
6 attempt to gain access to the server. The general response to any exceptional condition in the connection
7 negotiation is to pause that connection (presumably a thread) for a period of several seconds and then to
8 close the network connection. This includes syntax errors, over-sized data, and failed attempts to
9 authenticate. The server SHOULD log all such exceptions and flag or block clients provoking multiple
10 failures.

11

5 Conformance Tests

5.1 Introduction

The AMQP conformance tests are designed to verify how far an AMQ Protocol server actually conforms to the specifications laid out in this document. In principle, every "guideline for implementers", or "RULE" in the protocol's XML specification has a specific test that verifies whether the server conforms or not. In practice, some of the guidelines are intended for clients, and some are not testable without excessive cost.

The protocol itself cross references test by a logical label from within the protocol XML description, but the Test Sets will be documented elsewhere as developed and ratified by the AMQ Protocol governing body.

Note that tests do not test performance, stability, or scalability. The scope of the conformance tests is to measure how far an AMQP server is compatible with the protocol specifications, not how well it is built.

5.2 Design

5.2.1 "Test Sets" group Tests into meaningful capabilities

Because it is difficult for all implementations of the protocol to be at the same stage of completeness or compliance at all times, the concept of "Test Sets" is used to enable end users to easily identify the capability claims of a particular client or server implementation.

Test Sets are named groupings of related or commonly used functionality and the collection of tests which prove that functionality is compliant with some version of the AMQ Protocol.

Hence implementations can claim verifiable compliance with useful subsets of the protocol. In doing so users can have confidence in the product in question and its interoperability, and product providers can make rapid, visible, provable progress in delivering their products.

The Test Sets as a whole and the individual tests are designed as assertions. That is, each Test Set or individual test either succeeds, or exits with an assertion if it failed.

5.2.2 Wire-Level Tests

The wire-level tests check how the server:

1. Accepts the various types of valid data that the wire-level protocol defines, including frames, structured content, etc.
2. Handles incorrect data, e.g. malformed frames, incomplete content, etc.

5.2.3 Functional Tests

The functional tests check how the server:

- 1 1. Implements mandatory functionality, which is expressed in the specifications as "MUST" and "MUST
- 2 NOT".
- 3 2. Implements recommended functionality, which is expressed in the specifications as "SHOULD".
- 4 3. Implements optional functionality, which is expressed in the specifications as "MAY".
- 5 4. Handles limits, when the client creates excessive numbers of entities such as queues, consumers, etc.
- 6 5. Handles entity life-cycles: that deleted entities properly disappear, etc.

7 **5.3 Test Sets**

8 This section has still to be completed.

9

10 # End of Document #

11