



OASIS Advanced Message Queueing Protocol (AMQP) Version 1.0

Candidate OASIS Standard 01 07 August 2012

Specification URIs

This version:

<http://docs.oasis-open.org/amqp/core/v1.0/cos01/amqp-core-overview-v1.0-cos01.xml> (Authoritative)
<http://docs.oasis-open.org/amqp/core/v1.0/cos01/amqp-core-overview-v1.0-cos01.html>
<http://docs.oasis-open.org/amqp/core/v1.0/cos01/amqp-core-complete-v1.0-cos01.pdf>

Previous version:

<http://docs.oasis-open.org/amqp/core/v1.0/csprd01/amqp-core-overview-v1.0-csprd01.xml> (Authoritative)
<http://docs.oasis-open.org/amqp/core/v1.0/csprd01/amqp-core-overview-v1.0-csprd01.html>
<http://docs.oasis-open.org/amqp/core/v1.0/csprd01/amqp-core-complete-v1.0-csprd01.pdf>

Latest version:

<http://docs.oasis-open.org/amqp/core/v1.0/amqp-core-overview-v1.0.xml> (Authoritative)
<http://docs.oasis-open.org/amqp/core/v1.0/amqp-core-overview-v1.0.html>
<http://docs.oasis-open.org/amqp/core/v1.0/amqp-core-complete-v1.0.pdf>

Technical Committee:

OASIS Advanced Message Queueing Protocol (AMQP) TC

Chairs:

Ram Jeyaraman (Ram.Jeyaraman@microsoft.com), Microsoft
Angus Telfer (angus.telfer@inetco.com), INETCO Systems

Editors:

Robert Godfrey (robert.godfrey@jpmorgan.com), JPMorgan Chase & Co.
David Ingham (David.Ingham@microsoft.com), Microsoft
Rafael Schloming (rafaels@redhat.com), Red Hat

Additional artifacts:

This specification consists of the following documents:

- Part 0: Overview - Overview of the AMQP specification
<http://docs.oasis-open.org/amqp/core/v1.0/cos01/amqp-core-overview-v1.0-cos01.xml> (Authoritative)
<http://docs.oasis-open.org/amqp/core/v1.0/cos01/amqp-core-overview-v1.0-cos01.html>
- Part 1: Types - AMQP type system and encoding
<http://docs.oasis-open.org/amqp/core/v1.0/cos01/amqp-core-types-v1.0-cos01.xml> (Authoritative)
<http://docs.oasis-open.org/amqp/core/v1.0/cos01/amqp-core-types-v1.0-cos01.html>
- Part 2: Transport - AMQP transport layer
<http://docs.oasis-open.org/amqp/core/v1.0/cos01/amqp-core-transport-v1.0-cos01.xml> (Authoritative)
<http://docs.oasis-open.org/amqp/core/v1.0/cos01/amqp-core-transport-v1.0-cos01.html>
- Part 3: Messaging - AMQP Messaging Layer
<http://docs.oasis-open.org/amqp/core/v1.0/cos01/amqp-core-messaging-v1.0-cos01.xml> (Authoritative)
<http://docs.oasis-open.org/amqp/core/v1.0/cos01/amqp-core-messaging-v1.0-cos01.html>
- Part 4: Transactions - AMQP Transactions Layer
<http://docs.oasis-open.org/amqp/core/v1.0/cos01/amqp-core-transactions-v1.0-cos01.xml> (Authoritative)
<http://docs.oasis-open.org/amqp/core/v1.0/cos01/amqp-core-transactions-v1.0-cos01.html>
- Part 5: Security - AMQP Security Layers
<http://docs.oasis-open.org/amqp/core/v1.0/cos01/amqp-core-security-v1.0-cos01.xml> (Authoritative)
<http://docs.oasis-open.org/amqp/core/v1.0/cos01/amqp-core-security-v1.0-cos01.html>
- XML Document Type Definition (DTD)
<http://docs.oasis-open.org/amqp/core/v1.0/cos01/amqp.dtd> (Authoritative)

Related work:

This specification replaces or supersedes:

- AMQP v1.0 Final, 07 October 2011. <http://www.amqp.org/specification/1.0/amqp-org-download>

Abstract:

The Advanced Message Queuing Protocol (AMQP) is an open internet protocol for business messaging. It defines a binary wire-level protocol that allows for the reliable exchange of business messages between two parties. AMQP has a layered architecture and the specification is organized as a set of parts that reflects that architecture. Part 1 defines the AMQP type system and encoding. Part 2 defines the AMQP transport layer, an efficient, binary, peer-to-peer protocol for transporting messages between two processes over a network. Part 3 defines the AMQP message format, with a concrete encoding. Part 4 defines how interactions can be grouped within atomic transactions. Part 5 defines the AMQP security layers.

Status:

This document was last revised or approved by the OASIS Advanced Message Queuing Protocol (AMQP) TC on the above date. The level of approval is also listed above. Check the "Latest version" location noted above for possible later revisions of this document.

Technical Committee members should send comments on this specification to the Technical Committee's email list. Others should send comments to the Technical Committee by using the "Send A Comment" button on the Technical Committee's web page at <http://www.oasis-open.org/committees/amqp>.

For information on whether any patents have been disclosed that may be essential to implementing this specification, and any offers of patent licensing terms, please refer to the Intellectual Property Rights section of the Technical Committee web page (<http://www.oasis-open.org/committees/amqp/ipr.php>).

Citation format:

When referencing this specification the following citation format should be used:

[amqp-core-complete-v1.0]

OASIS Advanced Message Queuing Protocol (AMQP) Version 1.0. 07 August 2012. Candidate OASIS Standard 01.

<http://docs.oasis-open.org/amqp/core/v1.0/cos01/amqp-core-complete-v1.0-cos01.pdf>.

Notices:

Copyright © OASIS Open 2012. All Rights Reserved.

All capitalized terms in the following text have the meanings assigned to them in the OASIS Intellectual Property Rights Policy (the "OASIS IPR Policy"). The full Policy may be found at the OASIS website.

This document and translations of it may be copied and furnished to others, and derivative works that comment on or otherwise explain it or assist in its implementation may be prepared, copied, published, and distributed, in whole or in part, without restriction of any kind, provided that the above copyright notice and this section are included on all such copies and derivative works. However, this document itself may not be modified in any way, including by removing the copyright notice or references to OASIS, except as needed for the purpose of developing any document or deliverable produced by an OASIS Technical Committee (in which case the rules applicable to copyrights, as set forth in the OASIS IPR Policy, must be followed) or as required to translate it into languages other than English.

The limited permissions granted above are perpetual and will not be revoked by OASIS or its successors or assigns.

This document and the information contained herein is provided on an "AS IS" basis and OASIS DISCLAIMS ALL WARRANTIES, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO ANY WARRANTY THAT THE USE OF THE INFORMATION HEREIN WILL NOT INFRINGE ANY OWNERSHIP RIGHTS OR ANY IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.

OASIS requests that any OASIS Party or any other party that believes it has patent claims that would necessarily be infringed by implementations of this OASIS Committee Specification or OASIS Standard, to notify OASIS TC Administrator and provide an indication of its willingness to grant patent licenses to such patent claims in a manner consistent with the IPR Mode of the OASIS Technical Committee that produced this specification.

OASIS invites any party to contact the OASIS TC Administrator if it is aware of a claim of ownership of any patent claims that would necessarily be infringed by implementations of this specification by a patent holder that is not willing to provide a license to such patent claims in a manner consistent with the IPR Mode of the OASIS Technical Committee that produced this specification. OASIS may include such claims on its website, but disclaims any obligation to do so.

OASIS takes no position regarding the validity or scope of any intellectual property or other rights that might be claimed to pertain to the implementation or use of the technology described in this document or the extent to which any license under such rights might or might not be available; neither does it represent that it has made any effort to identify any such rights. Information on OASIS' procedures with respect to rights in any document or deliverable produced by an OASIS Technical Committee can be found on the OASIS website. Copies of claims of rights made available for publication and any assurances of licenses to be made available, or the result of an attempt made to obtain a general license or permission for the use of such proprietary rights by implementers or users of this OASIS Committee Specification or OASIS Standard, can be obtained from the OASIS TC Administrator. OASIS makes no representation that any information or list of intellectual property rights will at any time be complete, or that any claims in such list are, in fact, Essential Claims.

The name "OASIS" is a trademark of OASIS, the owner and developer of this specification, and should be used only to refer to the organization and its official outputs. OASIS welcomes reference to, and implementation and use of, specifications, while reserving the right to enforce its marks against misleading uses. Please see <http://www.oasis-open.org/policies-guidelines/trademark> for above guidance.

Contents

0	Overview	7
0.1	Introduction	7
0.1.1	Terminology	7
0.1.2	Normative References	7
0.1.3	Non-normative References	9
0.2	Conformance	10
0.3	Acknowledgements	11
0.4	Revision History	13
1	Types	16
1.1	Type System	16
1.1.1	Primitive Types	16
1.1.2	Described Types	16
1.1.3	Composite Types	16
1.1.4	Restricted Types	17
1.2	Type Encodings	17
1.2.1	Fixed Width	19
1.2.2	Variable Width	20
1.2.3	Compound	20
1.2.4	Array	20
1.3	Type Notation	21
1.3.1	Primitive Type Notation	21
1.3.2	Composite Type Notation	22
1.3.3	Descriptor Notation	22
1.3.4	Field Notation	23
1.3.5	Restricted Type Notation	23
1.4	Composite Type Representation	24
1.5	Descriptor Values	25
1.6	Primitive Type Definitions	25
1.6.1	null	25
1.6.2	boolean	25
1.6.3	ubyte	26
1.6.4	ushort	26
1.6.5	uint	26
1.6.6	ulong	26
1.6.7	byte	26
1.6.8	short	27
1.6.9	int	27
1.6.10	long	27
1.6.11	float	27
1.6.12	double	27
1.6.13	decimal32	28
1.6.14	decimal64	28
1.6.15	decimal128	28
1.6.16	char	28
1.6.17	timestamp	28
1.6.18	uuid	29
1.6.19	binary	29
1.6.20	string	29

1.6.21	symbol	29
1.6.22	list	29
1.6.23	map	30
1.6.24	array	30
2	Transport	31
2.1	Transport	31
2.1.1	Conceptual Model	31
2.1.2	Communication Endpoints	32
2.1.3	Protocol Frames	34
2.2	Version Negotiation	35
2.3	Framing	36
2.3.1	Frame Layout	37
2.3.2	AMQP Frames	38
2.4	Connections	38
2.4.1	Opening A Connection	39
2.4.2	Pipelined Open	40
2.4.3	Closing A Connection	40
2.4.4	Simultaneous Close	40
2.4.5	Idle Timeout Of A Connection	41
2.4.6	Connection States	41
2.4.7	Connection State Diagram	42
2.5	Sessions	43
2.5.1	Establishing A Session	44
2.5.2	Ending A Session	44
2.5.3	Simultaneous End	45
2.5.4	Session Errors	45
2.5.5	Session States	46
2.5.6	Session Flow Control	47
2.6	Links	48
2.6.1	Naming A Link	48
2.6.2	Link Handles	49
2.6.3	Establishing Or Resuming A Link	49
2.6.4	Detaching And Reattaching A Link	52
2.6.5	Link Errors	52
2.6.6	Closing A Link	53
2.6.7	Flow Control	53
2.6.8	Synchronous Get	55
2.6.9	Asynchronous Notification	56
2.6.10	Stopping A Link	56
2.6.11	Messages	57
2.6.12	Transferring A Message	57
2.6.13	Resuming Deliveries	60
2.6.14	Transferring Large Messages	61
2.7	Performatives	62
2.7.1	Open	62
2.7.2	Begin	64
2.7.3	Attach	65
2.7.4	Flow	67
2.7.5	Transfer	69
2.7.6	Disposition	71
2.7.7	Detach	72
2.7.8	End	72
2.7.9	Close	73
2.8	Definitions	73

2.8.1	Role	73
2.8.2	Sender Settle Mode	73
2.8.3	Receiver Settle Mode	74
2.8.4	Handle	74
2.8.5	Seconds	74
2.8.6	Milliseconds	74
2.8.7	Delivery Tag	74
2.8.8	Delivery Number	74
2.8.9	Transfer Number	75
2.8.10	Sequence No	75
2.8.11	Message Format	75
2.8.12	IETF Language Tag	75
2.8.13	Fields	75
2.8.14	Error	76
2.8.15	AMQP Error	76
2.8.16	Connection Error	77
2.8.17	Session Error	78
2.8.18	Link Error	78
2.8.19	Constant Definitions	79

3	Messaging	81
3.1	Introduction	81
3.2	Message Format	81
3.2.1	Header	82
3.2.2	Delivery Annotations	83
3.2.3	Message Annotations	84
3.2.4	Properties	84
3.2.5	Application Properties	86
3.2.6	Data	86
3.2.7	AMQP Sequence	86
3.2.8	AMQP Value	86
3.2.9	Footer	86
3.2.10	Annotations	87
3.2.11	Message ID ULong	87
3.2.12	Message ID UUID	87
3.2.13	Message ID Binary	87
3.2.14	Message ID String	87
3.2.15	Address String	87
3.2.16	Constant Definitions	87
3.3	Distribution Nodes	88
3.4	Delivery State	88
3.4.1	Received	89
3.4.2	Accepted	89
3.4.3	Rejected	90
3.4.4	Released	90
3.4.5	Modified	90
3.4.6	Resuming Deliveries Using Delivery States	91
3.5	Sources and Targets	95
3.5.1	Filtering Messages	95
3.5.2	Distribution Modes	95
3.5.3	Source	96
3.5.4	Target	97
3.5.5	Terminus Durability	98
3.5.6	Terminus Expiry Policy	99
3.5.7	Standard Distribution Mode	99

3.5.8	Filter Set	100
3.5.9	Node Properties	100
3.5.10	Delete On Close	100
3.5.11	Delete On No Links	100
3.5.12	Delete On No Messages	101
3.5.13	Delete On No Links Or Messages	101
4	Transactions	102
4.1	Transactional Messaging	102
4.2	Declaring a Transaction	102
4.3	Discharging a Transaction	103
4.4	Transactional Work	104
4.4.1	Transactional Posting	104
4.4.2	Transactional Retirement	105
4.4.3	Transactional Acquisition	106
4.4.4	Interaction Of Settlement With Transactions	107
4.4.4.1	Transactional Posting	107
4.4.4.2	Transactional Retirement	108
4.4.4.3	Transactional Acquisition	108
4.5	Coordination	108
4.5.1	Coordinator	108
4.5.2	Declare	109
4.5.3	Discharge	109
4.5.4	Transaction ID	110
4.5.5	Declared	110
4.5.6	Transactional State	110
4.5.7	Transaction Capability	111
4.5.8	Transaction Error	111
5	Security	112
5.1	Security Layers	112
5.2	TLS	112
5.2.1	Alternative Establishment	113
5.2.2	Constant Definitions	113
5.3	SASL	113
5.3.1	SASL Frames	114
5.3.2	SASL Negotiation	114
5.3.3	Security Frame Bodies	115
5.3.3.1	SASL Mechanisms	115
5.3.3.2	SASL Init	115
5.3.3.3	SASL Challenge	116
5.3.3.4	SASL Response	116
5.3.3.5	SASL Outcome	117
5.3.3.6	SASL Code	117
5.3.4	Constant Definitions	117
	XML Document Type Definition (DTD)	119

Part 0: Overview

0.1 Introduction

The Advanced Message Queuing Protocol is an open internet protocol for business messaging.

AMQP is comprised of several layers. The lowest level defines an efficient, binary, peer-to-peer protocol for transporting messages between two processes over a network. Above this, the messaging layer defines an abstract message format, with concrete standard encoding. Every compliant AMQP process **MUST** be able to send and receive messages in this standard encoding.

0.1.1 Terminology

The key words “**MUST**”, “**MUST NOT**”, “**REQUIRED**”, “**SHALL**”, “**SHALL NOT**”, “**SHOULD**”, “**SHOULD NOT**”, “**RECOMMENDED**”, “**MAY**”, and “**OPTIONAL**” in this specification are to be interpreted as described in IETF RFC 2119 [RFC2119].

The authoritative form of the AMQP specification consists of a set of XML source documents. These documents are transformed into PDF and HTML representations for readability. The machine readable version of the AMQP DTD describes the XML used for the authoritative source documents. This DTD includes the definition of the syntax used in the excerpts of XML presented in the PDF and HTML representations.

0.1.2 Normative References

[ASCII]

American National Standards Institute, Inc., *American National Standard for Information Systems, Coded Character Sets - 7-Bit American National Standard Code for Information Interchange (7-Bit ASCII)*, ANSI X3.4-1986, March 26, 1986.

[BCP47]

A. Phillips, Ed., M. Davis, Ed., *Tags for Identifying Languages*. IETF BCP: 47, September 2009.
<http://www.ietf.org/rfc/bcp/bcp47.txt>

[IANAHTTTPPARAMS]

IANA (Internet Assigned Numbers Authority), *Hypertext Transfer Protocol (HTTP) Parameters*.
<http://www.iana.org/assignments/http-parameters/http-parameters.xml>

[IANAPEN]

IANA (Internet Assigned Numbers Authority), *Private Enterprise Numbers*.
<http://www.iana.org/assignments/enterprise-numbers>

[IANASUBTAG]

IANA (Internet Assigned Numbers Authority), *Language Subtag Registry*.
<http://www.iana.org/assignments/language-subtag-registry>

[IEEE754]

Standard for Floating-Point Arithmetic. IEEE 754-2008, August 2008.
<http://ieeexplore.ieee.org/servlet/opac?punumber=4610933>

[IEEE1003]

The Single UNIX Specification, Version 4. IEEE Std 1003.1-2008, December 2008.
<http://www.unix.org/version4/>

[RFC1982]

R. Elz, R. Bush, *Serial Number Arithmetic*. IETF RFC 1982, August 1996.

<http://www.ietf.org/rfc/rfc1982.txt>

[RFC2046]

N. Freed, N. Borenstein, *Multipurpose Internet Mail Extensions (MIME) Part Two: Media Types*. IETF RFC 2046, November 1996.

<http://www.ietf.org/rfc/rfc2046.txt>

[RFC2119]

S. Bradner, *Key words for use in RFCs to Indicate Requirement Levels*. IETF RFC 2119, March 1997.

<http://www.ietf.org/rfc/rfc2119.txt>

[RFC2234]

D. Crocker, Ed., P. Overell, *Augmented BNF for Syntax Specifications: ABNF*. IETF RFC 2234, November 1997.

<http://www.ietf.org/rfc/rfc2234.txt>

[RFC2616]

R. Fielding, J. Gettys, J. Mogul, H. Frystyk, L. Masinter, P. Leach, T. Berners-Lee, *Hypertext Transfer Protocol – HTTP/1.1*. IETF RFC 2616, June 1999.

<http://www.ietf.org/rfc/rfc2616.txt>

[RFC4122]

P. Leach, M. Mealling, R. Salz, *A Universally Unique Identifier (UUID) URN Namespace*. IETF RFC 4122, July 2005.

<http://www.ietf.org/rfc/rfc4122.txt>

[RFC4366]

S. Blake-Wilson, M. Nystrom, D. Hopwood, J. Mikkelsen, T. Wright, *Transport Layer Security (TLS) Extensions*. IETF RFC 4366, April 2006.

<http://www.ietf.org/rfc/rfc4366.txt>

[RFC4422]

A. Melnikov, Ed., K. Zeilenga, Ed., *Simple Authentication and Security Layer (SASL)*. IETF RFC 4422, June 2006.

<http://www.ietf.org/rfc/rfc4422.txt>

[RFC4616]

K. Zeilenga, Ed., *The PLAIN Simple Authentication and Security Layer (SASL) Mechanism*. IETF RFC 4616, August 2006.

<http://www.ietf.org/rfc/rfc4616.txt>

[RFC5246]

T. Dierks, E. Rescorla., *The Transport Layer Security (TLS) Protocol Version 1.2*. IETF RFC 5246, August 2008.

<http://www.ietf.org/rfc/rfc5246.txt>

[RFC5646]

A. Phillips, Ed., M. Davis, Ed., *Tags for Identifying Languages*. IETF RFC 5646, September 2009.

<http://www.ietf.org/rfc/rfc5646.txt>

[RFC5802]

C. Newman, A. Menon-Sen, A. Melnikov, N. Williams, *Salted Challenge Response Authentication Mechanism (SCRAM) SASL and GSS-API Mechanisms*. IETF RFC 5802, July 2010.

<http://www.ietf.org/rfc/rfc5802.txt>

[UNICODE6]

The Unicode Consortium. The Unicode Standard, Version 6.0.0, (Mountain View, CA: The Unicode Consortium, 2011. ISBN 978-1-936213-01-6)

<http://www.unicode.org/versions/Unicode6.0.0/>

0.1.3 Non-normative References

[AMQPCONNCAP]

AMQP Capabilities Registry: Connection Capabilities

<http://www.amqp.org/specification/1.0/connection-capabilities>

[AMQPCONNPROP]

AMQP Capabilities Registry: Connection Properties

<http://www.amqp.org/specification/1.0/connection-properties>

[AMQPDELANN]

AMQP Capabilities Registry: Delivery Annotations

<http://www.amqp.org/specification/1.0/delivery-annotations>

[AMQPDISTMODE]

AMQP Capabilities Registry: Distribution Modes

<http://www.amqp.org/specification/1.0/distribution-modes>

[AMQPFILTERS]

AMQP Capabilities Registry: Filters

<http://www.amqp.org/specification/1.0/filters>

[AMQPFOOTER]

AMQP Capabilities Registry: Footer

<http://www.amqp.org/specification/1.0/footer>

[AMQPLINKCAP]

AMQP Capabilities Registry: Link Capabilities

<http://www.amqp.org/specification/1.0/link-capabilities>

[AMQPLINKPROP]

AMQP Capabilities Registry: Link Properties

<http://www.amqp.org/specification/1.0/link-properties>

[AMQPLINKSTATEPROP]

AMQP Capabilities Registry: Link State Properties

<http://www.amqp.org/specification/1.0/link-state-properties>

[AMQPMESSANN]

AMQP Capabilities Registry: Message Annotations

<http://www.amqp.org/specification/1.0/message-annotations>

[AMQPNODEPROP]

AMQP Capabilities Registry: Node Properties

<http://www.amqp.org/specification/1.0/node-properties>

[AMQPSESSCAP]

AMQP Capabilities Registry: Session Capabilities

<http://www.amqp.org/specification/1.0/session-capabilities>

[AMQPSESSPROP]

AMQP Capabilities Registry: Session Properties

<http://www.amqp.org/specification/1.0/session-properties>

[AMQPSOURCECAP]

AMQP Capabilities Registry: Source Capabilities

<http://www.amqp.org/specification/1.0/source-capabilities>

[AMQPTARGETCAP]

AMQP Capabilities Registry: Target Capabilities

<http://www.amqp.org/specification/1.0/target-capabilities>

0.2 Conformance

AMQP defines a wire level protocol for business messaging. The definition allows for all common business messaging behaviors. AMQP does not define a wire-level distinction between “clients” and “brokers”, the protocol is symmetric. It is expected and encouraged that implementations of AMQP will have different capabilities. Expectations of the capabilities of a “client library” are different from expectations of a “broker” which are themselves different from the capabilities of a “router”. As relevant profiles emerge (where appropriate and applicable) these will be formalised.

A conformant implementation **MUST** perform protocol negotiation (see Part 2: section 2.2), and then parse, process, and produce frames in accordance with the format and semantics defined in parts 1 through 5 of this specification.

Conformant implementations **MUST NOT** require the use of any extensions defined outside this document in order to interoperate with any other conformant implementation.

Part 1 of this document defines the type system and type encodings that every conformant implementation **MUST** implement.

Part 2 defines the peer-to-peer transport protocol which operates over TCP. Every conformant implementation of AMQP over TCP **MUST** implement Part 2. Future standards mapping AMQP to protocols other than TCP **MAY** modify or replace Part 2 when AMQP is being used over that protocol. A conformant implementation **MUST** implement Part 2 or a mapping of AMQP to some non-TCP protocol.

Part 2 admits behaviors that might not be appropriate for every implementation. For example a “client library” might not allow for its communication partner to spontaneously attempt to initiate a connection and request messages. Where an implementation does not allow for a behavior the implementation **MUST** respond according to the rules defined within Part 2 of the specification.

Part 3 of this document defines the AMQP Messaging Layer. Every conformant implementation which processes messages **MUST** implement this part of the specification.

Some implementations might not process messages (for example, an implementation acting as a “router” which looks only at the routing information carried by the AMQP Transport layer). Such implementations do not actively implement Part 3, but **MUST NOT** act in ways which violate the rules of this part of the specification.

The Messaging layer admits behaviors that might not be appropriate for all implementations (and within an implementation all behaviors might not be available for all configurations). Where a behavior is not admitted, the implementation **MUST** respond according to the rules defined within this specification.

Part 4 defines the requirements for transactional messaging. Transactional messaging defines two roles, that of the *transactional resource* and that of the *transaction controller*. A conformant implementation **SHOULD** be capable of operating in one of these roles but **MAY** be unable to operate in either (for instance a simplistic client library might have no ability to act as a transaction controller and would not be expected to act as a transactional resource).

It is **RECOMMENDED** that implementations designed to act as messaging intermediaries support the ability to act as a transactional resource. It is **RECOMMENDED** that implementations or re-usable libraries provide Application Programming Interfaces to enable them to act as transactional controllers.

Where a behavior is not admitted, the rules defined in part 4 regarding responses to non-admitted behaviors **MUST** be followed.

Part 5 defines Security Layers to provide an authenticated and/or encrypted transport. Implementations **SHOULD** allow the configuration of appropriate levels of security for the domain in which they are to be deployed.

Conformant implementations acting in the TCP server role are strongly **RECOMMENDED** to implement Part 5: section 5.2 (or Part 5: 5.2.1 Alternative Establishment). Implementations acting in the TCP server role are strongly **RECOMMENDED** to implement Part 5: section 5.3 and to support commonly used SASL mechanisms. In particular such implementations **SHOULD** support the PLAIN [RFC4616] and SCRAM-SHA1 [RFC5802] mechanisms.

Conformant implementations acting in the TCP client role SHOULD be capable of being configured to connect to an implementation in the TCP server role that is following the recommendations above.

0.3 Acknowledgements

The following individuals have contributed significantly towards the creation of this specification and are gratefully acknowledged:

- Michael Bridgen (VMware)
- Xin Chen (Microsoft)
- Raphael Cohn (Individual)
- Allan Cornish (INETCO Systems)
- Robert Godfrey (JPMorgan Chase & Co.)
- David Ingham (Microsoft)
- Andreas Moravec (Deutsche Boerse)
- Andreas Mueller (IIT Software)
- Simon MacMullen (VMware)
- John O'Hara (Individual)
- Rafael Schloming (Red Hat)
- Gordon Sim (Red Hat)
- Angus Telfer (INETCO Systems)

The following individuals were members of the OASIS AMQP Technical Committee during the creation of this specification and their contributions are gratefully acknowledged:

- Sanjay Aiyagari (VMware)
- Brian Albers (Kaazing)
- Matthew Arrott (Individual)
- Roger Bass (Traxian)
- Allan Beck (JPMorgan Chase & Co.)
- Mark Blair (Credit Suisse)
- Paul Blanz (Goldman Sachs)
- Laurie Bryson (JPMorgan Chase & Co.)
- Raphael Cohn (Individual)
- Allan Cornish (INETCO Systems)
- Robin Cover (OASIS)
- Raphael Delaporte (Zenika)
- Brad Drysdale (Kaazing)
- Helena Edelson (VMware)
- Chet Ensign (OASIS)
- John Fallows (Kaazing)

- Paul Fremantle (WSO2)
- Robert Gemmell (JPMorgan Chase & Co.)
- Rob Godfrey (JPMorgan Chase & Co.)
- Steve Hodge (Bank of America)
- Steve Huston (Individual)
- David Ingham (Microsoft)
- Stavros Isaiadis (Bank of America)
- Ram Jeyaraman (Microsoft)
- Bill Kahlert (HCL America)
- James Kirkland (Red Hat)
- Hanno Klein (Deutsche Boerse)
- Theo Kramer (Flame Computing Enterprises)
- Alex Kritikos (Software AG)
- Stanley Lewis (Progress Software)
- Mark Little (Red Hat)
- Colin MacNaughton (Progress Software)
- Mansour Mazinani (SITA SC)
- Shawn McAllister (Solace Systems)
- Donald McGarry (Mitre Corporation)
- Dale Moberg (Axway Software)
- Andreas Moravec (Deutsche Boerse)
- Andreas Mueller (IIT Software)
- Suryanarayanan Nagarajan (Software AG)
- Bob Natale (Mitre Corporation)
- John O'Hara (Individual)
- Kimberly Palko (Red Hat)
- Kenneth Peebles (Red Hat)
- Pierre Queinnec (Zenika)
- Makesh Rao (Cisco Systems)
- Alexis Richardson (VMware)
- Rafael Schloming (Red Hat)
- Dee Schur (OASIS)
- Gordon Sim (Red Hat)
- Sunjeet Singh (INETCO Systems)
- Bruce Snyder (VMware)
- Angus Telfer (INETCO Systems)
- Carl Trieloff (Red Hat)

- Asir Vedamuthu (Microsoft)
- Eamon Walshe (StormMQ)
- Denis Weerasiri (WSO2)
- Jeffrey Wheeler (Huawei Technologies Co.)
- Wang Xuan (Primeton Technologies)
- Prasad Yendluri (Software AG)

0.4 Revision History

2012-08-07 : Candidate OASIS Standard 01

AMQP-74 : HTML Rendering contains broken cross reference links for fields of primitive types

2012-07-18 : Committee Specification 1

AMQP-75 : Minor corrections to the front matter

2012-05-29 : Committee Specification Public Review Draft 2

2012-05-29 : Committee Specification Draft 2

2012-05-28 : Working Draft 9

AMQP-71 : Update acknowledgements to reflect current TC member list

AMQP-72 : In the PDF, add explicit links to all Parts of the document (including DTD) and state authoritative version

AMQP-73 : add clarity to document composition and XML excerpts

2012-05-15 : Working Draft 8

AMQP-62 : Better explain the role of the XML notation used in the specification and its relation to the wire format

AMQP-63 : Improve the conformance section

AMQP-64 : Fix typo in 2.1, second sentence needs initial capital letter

AMQP-65 : Better explain the relationship between links, connections, channels and sessions

AMQP-66 : Remove non-normative use of RFC2119 keywords

AMQP-67 : Improve presentation of the Revision History of the document

AMQP-69 : Fix rendering of sub- and superscript in HTML

2012-02-21 : Committee Specification Public Review Draft 1

2012-02-21 : Committee Specification Draft 1

2012-02-20 : Working Draft 7

- AMQP-56 : Fix links to non-normative AMQP Capabilities and Related Work
- AMQP-61 : Fixes to front-matter presentation

2012-02-15 : Working Draft 6

- AMQP-58 : Fix capitalization of headings on front page material
- AMQP-59 : Add missing acknowledgments
- AMQP-60 : Correct PDF ToC entry for DTD

2012-02-13 : Working Draft 5

- AMQP-50 : Section 3.2.16 is rendered differently in the html and pdf versions of the draft.
- AMQP-51 : Add missing acknowledgments
- AMQP-52 : Appendix A DTD is missing in the PDF document.
- AMQP-53 : Add Revision History
- AMQP-54 : Fix capitalization differences between HTML and PDF on front page
- AMQP-55 : Change labeling of DTD on front page
- AMQP-57 : Unresolved cross references due to removal of redundant subsection

2012-02-02 : Working Draft 4

- AMQP-43 : Review comments from Ram Jeyaraman

2012-02-01 : Working Draft 3

- AMQP-29 : Add introductory description of the spec's XML schema
- AMQP-36 : Clarify that delivery-count is not a count (in 2.6.7 and 2.7.4)
- AMQP-37 : Improve description of snd-settle-mode and rcv-settle-mode
- AMQP-38 : Provide more information as to how an incomplete settled map is resolved
- AMQP-39 : Clarify settlement on multi-transfer deliveries
- AMQP-41 : Rephrase language around keys in annotations
- AMQP-42 : Clarify use of error field in rejected outcome
- AMQP-44 : Specification should refer to AMQP.org document as related (superseded) work
- AMQP-47 : Messaging: Incorporate feedback from Steve Huston's review

2012-01-24 : Working Draft 2

- AMQP-7 : Typo in SaslInit.hostname description
- AMQP-8 : Slight contradiction in description of hostname in SaslInit and TLS SNI
- AMQP-9 : Could we add a repository for node-properties?
- AMQP-10 : Could we have a repository for additional supported-dist-modes?
- AMQP-11 : Message Properties.content-type MIME types need adjustment to fit new section types
- AMQP-12 : Labelling of headers and section numbers as per the AMQP WG document...
- AMQP-13 : Typo: Transactions in 4.4.1 transactional Posting "Settles" - > "Settle"
- AMQP-14 : PDF Document should use Sans-Serif font to be consistent with OASIS Standards
- AMQP-15 : Move copyright statement to second line of footer to be in line with OASIS style
- AMQP-16 : References to INETCO should be "INETCO Systems Limited"
- AMQP-17 : Typo in AMQP Open.hostname description

- AMQP-18 : Add figure titles to the HTML presentation
- AMQP-19 : Fix generated (sub-)section titles in the PDF where they contain a dash (“-”)
- AMQP-20 : PDF Presentation: use correct “open” and “close” double quotes
- AMQP-21 : Change presentation of cross-references to show only section number, not title
- AMQP-22 : Transport: Incorporate feedback from Steve Huston's review
- AMQP-23 : Abstract/Introduction: Incorporate feedback from Steve Huston's review
- AMQP-24 : Update XML transformation to render “n byte” as “n-byte”
- AMQP-25 : Types: Incorporate feedback from Steve Huston's review
- AMQP-26 : Transport: Incorporate feedback from Steve Huston's review (2)
- AMQP-27 : add missing diagram titles
- AMQP-28 : Confused use of the terms two's-complement“ and ”signed“
- AMQP-30 : Clarify version negotiation
- AMQP-31 : Figure numbering incorrect in the HTML presentation
- AMQP-32 : Clarify open behavior
- AMQP-33 : Tighten conformance language for 2.4.5
- AMQP-34 : Add description of HDR_EXCH state in 2.4.6
- AMQP-35 : Clarifications to 2.6.12
- AMQP-40 : Incorrect title for IEEE1003 reference
- AMQP-45 : Transactions: Incorporate feedback from Steve Huston's review
- AMQP-46 : Security: Incorporate feedback from Steve Huston's review
- AMQP-48 : Remove capitalization of certain terms within the specification
- AMQP-49 : Description of “dynamic” field refers to non-existent concepts of “session name” and “client-id”

2011-11-08 : Working Draft 1

- AMQP-1 : Make the definition of flow.echo more explicit
- AMQP-2 : Make it explicit that flow.properties is only for a link
- AMQP-3 : specification/overview.xml has a contributor's name in all caps
- AMQP-4 : specification/overview.xml is slightly inconsistent in use of company legal entity names and abbreviations thereof

Part 1: Types

1.1 Type System

The AMQP type system defines a set of commonly used primitive types for interoperable data representation. AMQP values can be annotated with additional semantic information beyond that associated with the primitive type. This allows for the association of an AMQP value with an external type that is not present as an AMQP primitive. For example, a URL is commonly represented as a string, however not all strings are valid URLs, and many programming languages and/or applications define a specific type to represent URLs. The AMQP type system would allow for the definition of a code with which to annotate strings when the value is intended to represent a URL.

1.1.1 Primitive Types

The AMQP type system defines a standard set of primitive types for representing both common scalar values and common collections. The scalar types include booleans, integral numbers, floating point numbers, timestamps, UUIDs, characters, strings, binary data, and symbols. The collection types include arrays (monomorphic), lists (polymorphic), and maps.

1.1.2 Described Types

The primitive types defined by AMQP can directly represent many of the basic types present in most popular programming languages, and therefore can be trivially used to exchange basic data. In practice, however, even the simplest applications have their own set of custom types used to model concepts within the application's domain. In messaging applications these custom types need to be externalized for transmission.

AMQP provides a means to do this by allowing any AMQP type to be annotated with a *descriptor*. A *descriptor* forms an association between a custom type, and an AMQP type. This association indicates that the AMQP type is actually a *representation* of the custom type. The resulting combination of the AMQP type and its descriptor is referred to as a *described type*.

A described type contains two distinct kinds of type information. It identifies both an AMQP type and a custom type (as well as the relationship between them), and so can be understood at two different levels. An application with intimate knowledge of a given domain can understand described types as the custom types they represent, thereby decoding and processing them according to the complete semantics of the domain. An application with no intimate knowledge can still understand the described types as AMQP types, decoding and processing them as such.

1.1.3 Composite Types

AMQP defines a number of *composite types* used for encoding structured data such as frame bodies. A composite type defines a composite value where each constituent value is identified by a well-known named *field*. Each composite type definition includes an ordered sequence of fields, each with a specified name, type, and multiplicity. Composite type definitions also include one or more descriptors (symbolic and/or numeric) for identifying their defined representations. Composite types are formally defined using the XML notation defined in section 1.3.

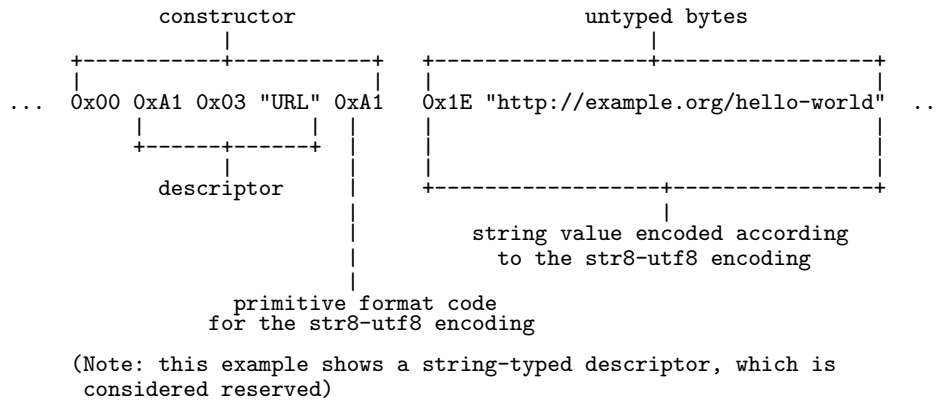


Figure 1.2: Described Format Code (URL)

The descriptor portion of a described format code is itself any valid AMQP encoded value, including other described values. The formal BNF for constructors is given below.

```

constructor = format-code
             / %x00 descriptor constructor

format-code = fixed / variable / compound / array
fixed       = empty / fixed-one / fixed-two / fixed-four
             / fixed-eight / fixed-sixteen
variable    = variable-one / variable-four
compound    = compound-one / compound-four
array       = array-one / array-four

descriptor  = value
value       = constructor untyped-bytes
untyped-bytes = *OCTET ; this is not actually *OCTET, the
                  ; valid byte sequences are restricted
                  ; by the constructor

; fixed width format codes
empty      = %x40-4E / %x4F %x00-FF
fixed-one  = %x50-5E / %x5F %x00-FF
fixed-two  = %x60-6E / %x6F %x00-FF
fixed-four = %x70-7E / %x7F %x00-FF
fixed-eight = %x80-8E / %x8F %x00-FF
fixed-sixteen = %x90-9E / %x9F %x00-FF

; variable width format codes
variable-one = %xA0-AE / %xAF %x00-FF
variable-four = %xB0-BE / %xBF %x00-FF

; compound format codes
compound-one = %xC0-CE / %xCF %x00-FF
compound-four = %xD0-DE / %xDF %x00-FF

; array format codes
array-one = %xE0-EE / %xEF %x00-FF
array-four = %xF0-FE / %xFF %x00-FF
  
```

Figure 1.3: Constructor BNF

Format codes map to one of four different categories: fixed width, variable width, compound and array. Values encoded within each category share the same basic structure parameterized by width. The subcategory within a format-code identifies both the category and width.

Fixed Width The size of fixed-width data is determined based solely on the subcategory of the format code for the fixed width value.




Figure 1.6: Layout of Fixed Width Data Encodings

1.2.2 Variable Width

All variable width encodings consist of a size in octets followed by *size* octets of encoded data. The width of the size for a specific variable width encoding can be computed from the subcategory of the format code:




Figure 1.7: Layout of Variable Width Data Encodings

1.2.3 Compound

All compound encodings consist of a size and a count followed by *count* encoded items. The width of the size and count for a specific compound encoding can be computed from the category of the format code:




Figure 1.8: Layout of Compound Data Encodings

1.2.4 Array

All array encodings consist of a size followed by a count followed by an element constructor followed by *count* elements of encoded data formatted as REQUIRED by the element constructor:




Figure 1.9: Layout of Array Encodings

1.3 Type Notation

Types are formally specified using an XML notation described below. Please note that this XML notation is used only to define the type and is never used to represent encoded data. All encoded AMQP data is binary as described in section 1.2.

The type element formally defines a semantic type used to exchange data on the wire. Every type definition has the following general form. The child elements present depend upon the class of the type definition (primitive, composite, or restricted):

```
<type class="primitive|composite|restricted" name="..." label="..." provides="...">
  ...
</type>
```

Figure 1.10: Type Definitions

The attributes of a type definition specify the following:

- name** The name of the type.
- label** A short description of the type.
- class** A string identifying what class of type is being defined: “primitive”, “composite”, or “restricted”.
- provides** A comma separated list of archetypes (see field element).

There are three different classes of types identified by the “class” attribute: primitive, composite, and restricted.

1.3.1 Primitive Type Notation

A “primitive” type definition will contain one or more encoding elements that describe how the data is formatted on the wire. Primitive types do not contain the descriptor, field, or choice elements.

```
<type class="primitive" name="..." label="..." provides="...">
  <encoding ... > ... </encoding>
  ...
  <encoding ... > ... </encoding>
</type>
```

Figure 1.11: Primitive Type Definitions

The encoding element defines how a primitive type is encoded. The specification defines 4 general categories of encodings: fixed, variable, compound, and array. A specific encoding provides further details of how data is formatted within its general category.

```
<type class="primitive" ... >
  ...
  <encoding name="..." code="0xNN" category="fixed|variable|compound|array" width="N" label="..."/>
  ...
</type>
```

Figure 1.12: Encoding Definitions

name	The name of the encoding.
label	A short description of the encoding.
code	The numeric value that prefixes the encoded data on the wire.
category	The category of the encoding: “fixed”, “variable”, “compound”, or “array”.
width	The width of the encoding or the size/count prefixes depending on the category.

1.3.2 Composite Type Notation

A “composite” type definition specifies a new kind of record type, i.e., a type composed of a fixed number of fields whose values are each of a specific type. A composite type does not have a new encoding, but is sent on the wire as a list annotated by a specific descriptor value that identifies it as a representation of the defined composite type. A composite type definition will contain a descriptor and zero or more field elements. Composite types do not contain encoding or choice elements. The source attribute of a composite type will always be “list”, other values are reserved for future use.

```
<type class="composite" name="..." label="..." provides="...">
  <descriptor name="..." code="..."/>
  <field name="..." ... > ... </field>
  ...
  <field name="..." ... > ... </field>
</type>
```

Figure 1.13: Composite Type Definitions

1.3.3 Descriptor Notation

The descriptor element specifies what annotation is used to identify encoded values as representations of a described type.

```
<type class="composite|restricted" ...>
  <!-- domain-id:descriptor-id -->
  <descriptor name="domain:name" code="0xNNNNNNNN:0xNNNNNNNN"/>
  ...
</type>
```

Figure 1.14: Descriptor Definitions

name	A symbolic name for the representation.
code	The numeric value.

1.3.4 Field Notation

The field element identifies a field within a composite type. Every field has the following attributes:

```
<type class="composite" ...>
  <field name="..." type="..." requires="..." default="..." label="..."
    mandatory="true|false"
    multiple="true|false" />
  ...
</type>
```

Figure 1.15: Field Definitions

name	A name that uniquely identifies the field within the type.
type	The type of the field. This attribute defines the range of values that are permitted to appear in this field. It might name a specific type, in which case the values are restricted to that type, or it might be the special character "*", in which case a value of any type is permitted. In the latter case the range of values might be further restricted by the requires attribute.
requires	A comma separated list of archetypes. Field values are restricted to types providing at least one of the specified archetypes.
default	A default value for the field if no value is encoded.
label	A short description of the field.
mandatory	"true" iff a non null value for the field is always encoded.
multiple	"true" iff the field can have multiple values of its specified type.

1.3.5 Restricted Type Notation

A "restricted" type definition specifies a new kind of type whose values are restricted to a subset of the values that might be represented with another type. The source attribute identifies the base type from which a restricted type is derived. A restricted type can have a descriptor element in which case it is identified by a descriptor on the wire. The values permitted by a restricted type can be specified either via documentation, or directly limited to a fixed number of values by the choice elements contained within the definition.

```
<type class="restricted" name="..." label="..." provides="...">
  <descriptor name="..." code="..." />
  <choice name="..." value="..." />
  ...
  <choice name="..." value="..." />
</type>
```

Figure 1.16: Restricted Type Definitions

The choice element identifies a legal value for a restricted type. The choice element has the following attributes:

```
<type class="restricted" ...>
  ...
  <choice name="..." value="..." />
  ...
</type>
```

Figure 1.17: Choice Definitions

name	A name for the value.
value	The permitted value.

1.4 Composite Type Representation

AMQP composite types are represented as a described list. Each element in the list is positionally correlated with the fields listed in the composite type definition. The permitted element values are determined by the type specification and multiplicity of the corresponding field definitions. When the trailing elements of the list representation are null, they MAY be omitted. The descriptor of the list indicates the specific composite type being represented.

```
<type class="composite" name="book" label="example composite type">
  <doc>
    <p>An example composite type.</p>
  </doc>

  <descriptor name="example:book:list" code="0x00000003:0x00000002"/>
  <field name="title" type="string" mandatory="true" label="title of the book"/>
  <field name="authors" type="string" multiple="true"/>
  <field name="isbn" type="string" label="the ISBN code for the book"/>
</type>
```

Figure 1.18: Example Composite Type

The *mandatory* attribute of a field description controls whether a null element value is permitted in the representation.

The *multiple* attribute of a field description controls whether multiple element values are permitted in the representation. A single element of the type specified in the field description is always permitted. Multiple values are represented by the use of an array where the type of the elements in the array is the type defined in the field definition. Note that a null value and a zero-length array (with a correct type for its elements) both describe an absence of a value and MUST be treated as semantically identical.

A field which is defined as both multiple and mandatory MUST contain at least one value (i.e. for such a field both *null* and an array with no entries are invalid).

The described list shown below is an example composite value of the *book* type defined above. A trailing null element corresponding to the absence of an ISBN value is depicted in the example, but can optionally be omitted according to the encoding rules.


```
<type name="boolean" class="primitive">
  <encoding code="0x56" category="fixed" width="1"
    label="boolean with the octet 0x00 being false and octet 0x01 being true"/>
  <encoding name="true" code="0x41" category="fixed" width="0" label="the boolean value true"/>
  <encoding name="false" code="0x42" category="fixed" width="0" label="the boolean value false"/>
</type>
```

1.6.3 ubyte

Integer in the range 0 to $2^8 - 1$ inclusive.

```
<type name="ubyte" class="primitive">
  <encoding code="0x50" category="fixed" width="1" label="8-bit unsigned integer"/>
</type>
```

1.6.4 ushort

Integer in the range 0 to $2^{16} - 1$ inclusive.

```
<type name="ushort" class="primitive">
  <encoding code="0x60" category="fixed" width="2"
    label="16-bit unsigned integer in network byte order"/>
</type>
```

1.6.5 uint

Integer in the range 0 to $2^{32} - 1$ inclusive.

```
<type name="uint" class="primitive">
  <encoding code="0x70" category="fixed" width="4"
    label="32-bit unsigned integer in network byte order"/>
  <encoding name="smalluint" code="0x52" category="fixed" width="1"
    label="unsigned integer value in the range 0 to 255 inclusive"/>
  <encoding name="uint0" code="0x43" category="fixed" width="0" label="the uint value 0"/>
</type>
```

1.6.6 ulong

Integer in the range 0 to $2^{64} - 1$ inclusive.

```
<type name="ulong" class="primitive">
  <encoding code="0x80" category="fixed" width="8"
    label="64-bit unsigned integer in network byte order"/>
  <encoding name="smallulong" code="0x53" category="fixed" width="1"
    label="unsigned long value in the range 0 to 255 inclusive"/>
  <encoding name="ulong0" code="0x44" category="fixed" width="0" label="the ulong value 0"/>
</type>
```

1.6.7 byte

Integer in the range $-(2^7)$ to $2^7 - 1$ inclusive.

```
<type name="byte" class="primitive">
  <encoding code="0x51" category="fixed" width="1"
    label="8-bit two's-complement integer"/>
</type>
```

1.6.8 short

Integer in the range $-(2^{15})$ to $2^{15} - 1$ inclusive.

```
<type name="short" class="primitive">
  <encoding code="0x61" category="fixed" width="2"
    label="16-bit two's-complement integer in network byte order"/>
</type>
```

1.6.9 int

Integer in the range $-(2^{31})$ to $2^{31} - 1$ inclusive.

```
<type name="int" class="primitive">
  <encoding code="0x71" category="fixed" width="4"
    label="32-bit two's-complement integer in network byte order"/>
  <encoding name="smallint" code="0x54" category="fixed" width="1"
    label="8-bit two's-complement integer"/>
</type>
```

1.6.10 long

Integer in the range $-(2^{63})$ to $2^{63} - 1$ inclusive.

```
<type name="long" class="primitive">
  <encoding code="0x81" category="fixed" width="8"
    label="64-bit two's-complement integer in network byte order"/>
  <encoding name="smalllong" code="0x55" category="fixed" width="1"
    label="8-bit two's-complement integer"/>
</type>
```

1.6.11 float

32-bit floating point number (IEEE 754-2008 binary32).

```
<type name="float" class="primitive">
  <encoding name="ieee-754" code="0x72" category="fixed" width="4" label="IEEE 754-2008 binary32"/>
</type>
```

A 32-bit floating point number (IEEE 754-2008 binary32 [IEEE754]).

1.6.12 double

64-bit floating point number (IEEE 754-2008 binary64).

```
<type name="double" class="primitive">
  <encoding name="ieee-754" code="0x82" category="fixed" width="8" label="IEEE 754-2008 binary64"/>
</type>
```

A 64-bit floating point number (IEEE 754-2008 binary64 [IEEE754]).

1.6.13 decimal32

32-bit decimal number (IEEE 754-2008 decimal32).

```
<type name="decimal32" class="primitive">
  <encoding name="ieee-754" code="0x74" category="fixed" width="4"
    label="IEEE 754-2008 decimal32 using the Binary Integer Decimal encoding"/>
</type>
```

A 32-bit decimal number (IEEE 754-2008 decimal32 [IEEE754]).

1.6.14 decimal64

64-bit decimal number (IEEE 754-2008 decimal64).

```
<type name="decimal64" class="primitive">
  <encoding name="ieee-754" code="0x84" category="fixed" width="8"
    label="IEEE 754-2008 decimal64 using the Binary Integer Decimal encoding"/>
</type>
```

A 64-bit decimal number (IEEE 754-2008 decimal64 [IEEE754]).

1.6.15 decimal128

128-bit decimal number (IEEE 754-2008 decimal128).

```
<type name="decimal128" class="primitive">
  <encoding name="ieee-754" code="0x94" category="fixed" width="16"
    label="IEEE 754-2008 decimal128 using the Binary Integer Decimal encoding"/>
</type>
```

A 128-bit decimal number (IEEE 754-2008 decimal128 [IEEE754]).

1.6.16 char

A single Unicode character.

```
<type name="char" class="primitive">
  <encoding name="utf32" code="0x73" category="fixed" width="4"
    label="a UTF-32BE encoded Unicode character"/>
</type>
```

A UTF-32BE encoded Unicode character [UNICODE6].

1.6.17 timestamp

An absolute point in time.

```
<type name="timestamp" class="primitive">
  <encoding name="ms64" code="0x83" category="fixed" width="8"
    label="64-bit two's-complement integer representing milliseconds since the unix epoch"/>
</type>
```

Represents an approximate point in time using the Unix time `t` [IEEE1003] encoding of UTC, but with a precision of milliseconds. For example, 1311704463521 represents the moment 2011-07-26T18:21:03.521Z.

1.6.18 uuid

A universally unique identifier as defined by RFC-4122 section 4.1.2 .

```
<type name="uuid" class="primitive">
  <encoding code="0x98" category="fixed" width="16"
    label="UUID as defined in section 4.1.2 of RFC-4122"/>
</type>
```

UUID is defined in section 4.1.2 of RFC-4122 [RFC4122].

1.6.19 binary

A sequence of octets.

```
<type name="binary" class="primitive">
  <encoding name="vbin8" code="0xa0" category="variable" width="1"
    label="up to 2^8 - 1 octets of binary data"/>
  <encoding name="vbin32" code="0xb0" category="variable" width="4"
    label="up to 2^32 - 1 octets of binary data"/>
</type>
```

1.6.20 string

A sequence of Unicode characters.

```
<type name="string" class="primitive">
  <encoding name="str8-utf8" code="0xa1" category="variable" width="1"
    label="up to 2^8 - 1 octets worth of UTF-8 Unicode (with no byte order mark)"/>
  <encoding name="str32-utf8" code="0xb1" category="variable" width="4"
    label="up to 2^32 - 1 octets worth of UTF-8 Unicode (with no byte order mark)"/>
</type>
```

A string represents a sequence of Unicode characters as defined by the Unicode V6.0.0 standard [UNICODE6].

1.6.21 symbol

Symbolic values from a constrained domain.

```
<type name="symbol" class="primitive">
  <encoding name="sym8" code="0xa3" category="variable" width="1"
    label="up to 2^8 - 1 seven bit ASCII characters representing a symbolic value"/>
  <encoding name="sym32" code="0xb3" category="variable" width="4"
    label="up to 2^32 - 1 seven bit ASCII characters representing a symbolic value"/>
</type>
```

Symbols are values from a constrained domain. Although the set of possible domains is open-ended, typically the both number and size of symbols in use for any given application will be small, e.g. small enough that it is reasonable to cache all the distinct values. Symbols are encoded as ASCII characters [ASCII].

1.6.22 list

A sequence of polymorphic values.

```

<type name="list" class="primitive">
  <encoding name="list0" code="0x45" category="fixed" width="0"
    label="the empty list (i.e. the list with no elements)"/>
  <encoding name="list8" code="0xc0" category="compound" width="1"
    label="up to 2^8 - 1 list elements with total size less than 2^8 octets"/>
  <encoding name="list32" code="0xd0" category="compound" width="4"
    label="up to 2^32 - 1 list elements with total size less than 2^32 octets"/>
</type>

```

1.6.23 map

A polymorphic mapping from distinct keys to values.

```

<type name="map" class="primitive">
  <encoding name="map8" code="0xc1" category="compound" width="1"
    label="up to 2^8 - 1 octets of encoded map data"/>
  <encoding name="map32" code="0xd1" category="compound" width="4"
    label="up to 2^32 - 1 octets of encoded map data"/>
</type>

```

A map is encoded as a compound value where the constituent elements form alternating key value pairs.




Figure 1.20: Layout of Map Encoding

Map encodings MUST contain an even number of items (i.e. an equal number of keys and values). A map in which there exist two identical key values is invalid. Unless known to be otherwise, maps MUST be considered to be ordered, that is, the order of the key-value pairs is semantically important and two maps which are different only in the order in which their key-value pairs are encoded are not equal.

1.6.24 array

A sequence of values of a single type.

```

<type name="array" class="primitive">
  <encoding name="array8" code="0xe0" category="array" width="1"
    label="up to 2^8 - 1 array elements with total size less than 2^8 octets"/>
  <encoding name="array32" code="0xf0" category="array" width="4"
    label="up to 2^32 - 1 array elements with total size less than 2^32 octets"/>
</type>

```

Part 2: Transport

2.1 Transport

2.1.1 Conceptual Model

An AMQP network consists of *nodes* connected via *links*. Nodes are named entities responsible for the safe storage and/or delivery of *messages*. Messages can originate from, terminate at, or be relayed by nodes.

A link is a unidirectional route between two nodes. A link attaches to a node at a *terminus*. There are two kinds of terminus: *sources* and *targets*. A terminus is responsible for tracking the state of a particular stream of incoming or outgoing messages. Sources track outgoing messages and targets track incoming messages. Messages only travel along a link if they meet the entry criteria at the source.

As a message travels through an AMQP network, the responsibility for safe storage and delivery of the message is transferred between the nodes it encounters. The *link protocol* (defined in section 2.6) manages the transfer of responsibility between the source and target.




Figure 2.1: Message Transfer between Nodes

Nodes exist within a *container*. Examples of containers are *brokers* and *client* applications. Each container MAY hold many nodes. Examples of AMQP nodes are *producers*, *consumers*, and *queues*. Producers and consumers are the elements within an application that generate and process messages. Queues are entities that store and forward messages.




Figure 2.2: Class Diagram of Concrete Containers and Nodes

2.1.2 Communication Endpoints

The AMQP transport specification defines a peer-to-peer protocol for transferring messages between nodes in an AMQP network. This portion of the specification is not concerned with the internal workings of any sort of node, and only deals with the mechanics of unambiguously transferring a message from one node to another.

In order for communication to occur between nodes in different containers a connection needs be established.




Figure 2.3: An AMQP Connection

An AMQP connection consists of a full-duplex, reliably ordered sequence of *frames*. A frame is the unit of work carried on the wire. Connections have a negotiated maximum frame size allowing byte streams to be easily defragmented into complete frame bodies representing the independently parsable units formally defined in section 2.7. The precise requirement for a connection is that if the n^{th} frame arrives, all frames prior to n MUST also have arrived. It is assumed connections are transient and can fail for a variety of reasons resulting in the loss of an unknown number of frames, but they are still subject to the aforementioned ordered reliability criteria. This is similar to the guarantee that TCP or SCTP provides for byte streams, and the specification defines a framing system used to parse a byte stream into a sequence of frames for use in establishing an AMQP connection (see section 2.3).

An AMQP connection is divided into a negotiated number of independent unidirectional *channels*. Each frame is marked with the channel number indicating its parent channel, and the frame sequence for each channel is multiplexed into a single frame sequence for the connection.

An AMQP *session* correlates two unidirectional channels to form a bidirectional, sequential conversation between two containers. Sessions provide a flow control scheme based on the number of *transfer frames* transmitted. Since frames have a maximum size for a given connection, this provides flow control based on the number of bytes transmitted and can be used to optimize performance.




Figure 2.4: An AMQP Session

A single connection MAY have multiple independent sessions active simultaneously, up to the negotiated channel limit. Both connections and sessions are modeled by each peer as *endpoints* that store local and last known remote state regarding the connection or session in question.




Figure 2.5: Session & Connection Endpoints

In order to transfer messages between nodes (e.g., to move messages from a queue to a consumer) a *link* needs to be established between the nodes. A link is a unidirectional route between two nodes. A link attaches to a node at a terminus. There are two kinds of terminus: sources and targets. A terminus is responsible for tracking the state of a particular stream of incoming or outgoing messages. Sources track outgoing messages and targets track incoming messages. Messages only travel along a link if they meet the entry criteria at the source.

Links provide a credit-based flow control scheme based on the number of messages transmitted, allowing applications to control which nodes to receive messages from at a given point (e.g., to explicitly fetch a message from a given queue).




Figure 2.6: An AMQP Link

Sessions provide the context for communication between sources and targets. A *link endpoint* associates a terminus with a *session endpoint*. Within a session, the link protocol (defined in section 2.6) is used to establish links between sources and targets and to transfer messages across them. A single session can be simultaneously associated with any number of links.

Links are named, and the state at the termini can live longer than the connection on which they were established.



Figure 2.7: Connection Loss

The retained state at the termini can be used to reestablish the link on a new connection (and session) with precise control over delivery guarantees (e.g., ensuring “exactly once” delivery).




Figure 2.8: Link Recovery

The diagram below shows the relationship between the three AMQP communication endpoints, Connection, Session and Link.

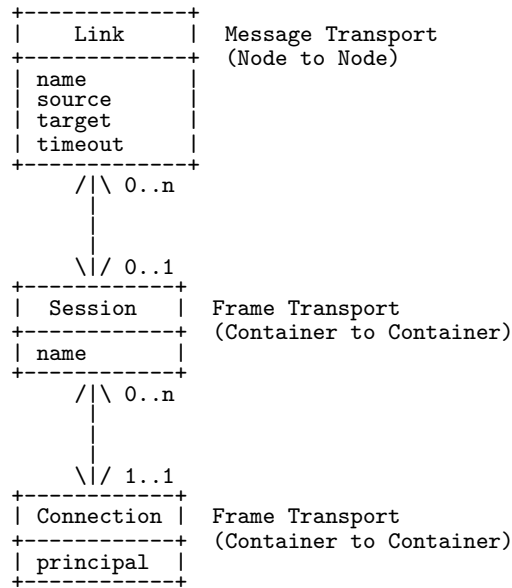


Figure 2.9: Class Diagram of Communication Endpoints

2.1.3 Protocol Frames

The protocol consists of nine frame body types that are formally defined in section 2.7. The following table lists the frame bodies and defines which endpoints handle them.

Frame	Connection	Session	Link
open	H		
begin	I	H	
attach		I	H
flow		I	H
transfer		I	H
disposition		I	H
detach		I	H
end	I	H	
close	H		

Key:

H: handled by the endpoint

I: intercepted (endpoint examines the frame, but delegates further processing to another endpoint)

Figure 2.10: Frame Dispatch Table

2.2 Version Negotiation

Prior to sending any frames on a connection, each peer MUST start by sending a protocol header that indicates the protocol version used on the connection. The protocol header consists of the upper case ASCII letters “AMQP” followed by a protocol id of zero, followed by three unsigned bytes representing the major, minor, and revision of the protocol version (currently 1 (MAJOR), 0 (MINOR), 0 (REVISION)). In total this is an 8-octet sequence:

4 OCTETS	1 OCTET	1 OCTET	1 OCTET	1 OCTET
"AMQP"	%d0	major	minor	revision

Figure 2.11: Protocol Header Layout

Any data appearing beyond the protocol header MUST match the version indicated by the protocol header. If the incoming and outgoing protocol headers do not match, both peers MUST close their outgoing stream and SHOULD read the incoming stream until it is terminated.

The AMQP peer which acted in the role of the TCP client (i.e. the peer that actively opened the connection) MUST immediately send its outgoing protocol header on establishment of the TCP connection. The AMQP peer which acted in the role of the TCP server MAY elect to wait until receiving the incoming protocol header before sending its own outgoing protocol header. This permits a multi protocol server implementation to choose the correct protocol version to fit each client.

Two AMQP peers agree on a protocol version as follows (where the words “client” and “server” refer to the roles being played by the peers at the TCP connection level):

- When the client opens a new socket connection to a server, it MUST send a protocol header with the client’s preferred protocol version.
- If the requested protocol version is supported, the server MUST send its own protocol header with the requested version to the socket, and then proceed according to the protocol definition.
- If the requested protocol version is **not** supported, the server MUST send a protocol header with a **supported** protocol version and then close the socket.
- When choosing a protocol version to respond with, the server SHOULD choose the highest supported version that is less than or equal to the requested version. If no such version exists, the server SHOULD respond with the highest supported version.

- If the server cannot parse the protocol header, the server **MUST** send a valid protocol header with a supported protocol version and then close the socket.
- Note that if the server only supports a single protocol version, it is consistent with the above rules for the server to send its protocol header prior to receiving anything from the client and to subsequently close the socket if the client's protocol header does not match the server's.

Based on this behavior a client can discover which protocol versions a server supports by attempting to connect with its highest supported version and reconnecting with a version less than or equal to the version received back from the server.

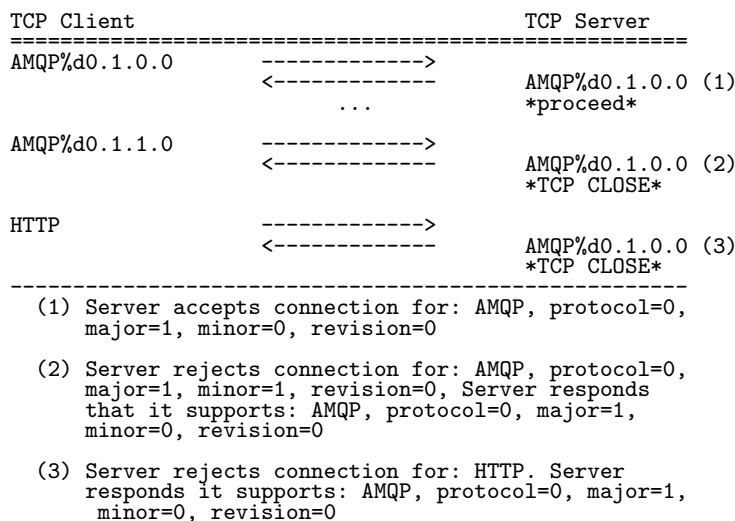


Figure 2.12: Version Negotiation Examples

Please note that the above examples use the literal notation defined in RFC 2234 [RFC2234] for non alphanumeric values.

The protocol id is not a part of the protocol version and thus the rule above regarding the highest supported version does not apply. A client might request use of a protocol id that is unacceptable to a server - for example, it might request a raw AMQP connection when the server is configured to require a TLS or SASL security layer (See Part 5: section 5.1). In this case, the server **MUST** send a protocol header with an **acceptable** protocol id (and version) and then close the socket. It **MAY** choose any protocol id.




Figure 2.13: Protocol ID Rejection Example

2.3 Framing

Frames are divided into three distinct areas: a fixed width frame header, a variable width extended header, and a variable width frame body.

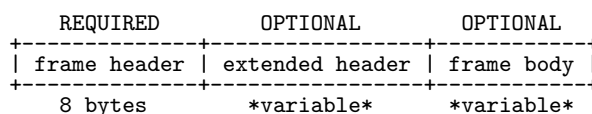


Figure 2.14: Frame Layout

- frame header** The frame header is a fixed size (8 byte) structure that precedes each frame. The frame header includes mandatory information necessary to parse the rest of the frame including size and type information.
- extended header** The extended header is a variable width area preceding the frame body. This is an extension point defined for future expansion. The treatment of this area depends on the frame type.
- frame body** The frame body is a variable width sequence of bytes the format of which depends on the frame type.

2.3.1 Frame Layout

The diagram below shows the details of the general frame layout for all frame types.

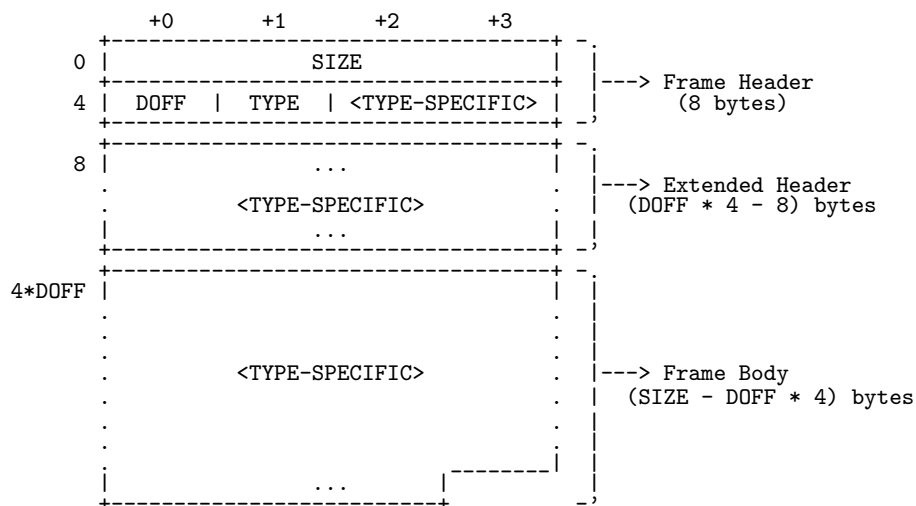


Figure 2.15: General Frame Layout

- SIZE** Bytes 0-3 of the frame header contain the frame size. This is an unsigned 32-bit integer that **MUST** contain the total frame size of the frame header, extended header, and frame body. The frame is malformed if the size is less than the size of the frame header (8 bytes).
- DOFF** Byte 4 of the frame header is the data offset. This gives the position of the body within the frame. The value of the data offset is an unsigned, 8-bit integer specifying a count of 4-byte words. Due to the mandatory 8-byte frame header, the frame is malformed if the value is less than 2.
- TYPE** Byte 5 of the frame header is a type code. The type code indicates the format and purpose of the frame. The subsequent bytes in the frame header **MAY** be interpreted differently depending on the type of the frame. A type code of 0x00 indicates that the frame is an AMQP frame. A type code of 0x01 indicates that the frame is a SASL frame, see Part 5: section 5.3.

2.3.2 AMQP Frames

Bytes 6 and 7 of an AMQP frame contain the channel number (see section 2.1). The frame body is defined as a *performative* followed by an opaque *payload*. The performative **MUST** be one of those defined in section 2.7 and is encoded as a described type in the AMQP type system. The remaining bytes in the frame body form the payload for that frame. The presence and format of the payload is defined by the semantics of the given performative.




Figure 2.16: AMQP Frame Layout

An AMQP frame with no body **MAY** be used to generate artificial traffic as needed to satisfy any negotiated idle timeout interval (see subsection 2.4.5).

2.4 Connections

AMQP connections are divided into a number of unidirectional channels. A connection endpoint contains two kinds of channel endpoints: incoming and outgoing. A connection endpoint maps incoming frames other than `open` and `close` to an incoming channel endpoint based on the incoming channel number, as well as relaying frames produced by outgoing channel endpoints, marking them with the associated outgoing channel number before sending them.

This requires connection endpoints to contain two mappings. One from incoming channel number to incoming channel endpoint, and one from outgoing channel endpoint, to outgoing channel number.

2.4.2 Pipelined Open

For applications that use many short-lived connections, it MAY be desirable to pipeline the connection negotiation process. A peer MAY do this by starting to send subsequent frames before receiving the partner's connection header or open frame. This is permitted so long as the pipelined frames are known *a priori* to conform to the capabilities and limitations of its partner. For example, this can be accomplished by keeping the use of the connection within the capabilities and limits expected of all AMQP implementations as defined by the specification of the open frame.




Figure 2.20: Pipelined Connection Open Sequence

The use of pipelined frames by a peer cannot be distinguished by the peer's partner from non-pipelined use so long as the pipelined frames conform to the partner's capabilities and limitations.

2.4.3 Closing A Connection

Prior to closing a connection, each peer MUST write a `close` frame with a code indicating the reason for closing. This frame MUST be the last thing ever written onto a connection. After writing this frame the peer SHOULD continue to read from the connection until it receives the partner's `close` frame (in order to guard against erroneously or maliciously implemented partners, a peer SHOULD implement a timeout to give its partner a reasonable time to receive and process the close before giving up and simply closing the underlying transport mechanism). A `close` frame MAY be received on any channel up to the maximum channel number negotiated in open. However, implementations SHOULD send it on channel 0, and MUST send it on channel 0 if pipelined in a single batch with the corresponding `open`.




Figure 2.21: Synchronous Connection Close Sequence

Implementations SHOULD NOT expect to be able to reuse open TCP sockets after `close` performatives have been exchanged. There is no requirement for an implementation to read from a socket after a `close` performative has been received.

2.4.4 Simultaneous Close

Normally one peer will initiate the connection close, and the partner will send its close in response. However, because both endpoints MAY simultaneously choose to close the connection for independent reasons, it is possible

for a simultaneous close to occur. In this case, the only potentially observable difference from the perspective of each endpoint is the code indicating the reason for the close.




Figure 2.22: Simultaneous Connection Close Sequence

2.4.5 Idle Timeout Of A Connection

Connections are subject to an idle timeout threshold. The timeout is triggered by a local peer when no frames are received after a threshold value is exceeded. The idle timeout is measured in milliseconds, and starts from the time the last frame is received. If the threshold is exceeded, then a peer SHOULD try to gracefully close the connection using a `close` frame with an error explaining why. If the remote peer does not respond gracefully within a threshold to this, then the peer MAY close the TCP socket.

Each peer has its own (independent) idle timeout. At connection open each peer communicates the maximum period between activity (frames) on the connection that it desires from its partner. The `open` frame carries the idle-time-out field for this purpose. To avoid spurious timeouts, the value in idle-time-out SHOULD be half the peer's actual timeout threshold.

If a peer can not, for any reason support a proposed idle timeout, then it SHOULD close the connection using a `close` frame with an error explaining why. There is no requirement for peers to support arbitrarily short or long idle timeouts.

The use of idle timeouts is in addition to any network protocol level control. Implementations SHOULD make use of TCP keep-alive wherever possible in order to be good citizens.

If a peer needs to satisfy the need to send traffic to prevent idle timeout, and has nothing to send, it MAY send an empty frame, i.e., a frame consisting solely of a frame header, with no frame body. Implementations MUST be prepared to handle empty frames arriving on any valid channel, though implementations SHOULD use channel 0 when sending empty frames, and MUST use channel 0 if a maximum channel number has not yet been negotiated (i.e., before an `open` frame has been received). Apart from this use, empty frames have no meaning.

Empty frames can only be sent after the `open` frame is sent. As they are a frame, they MUST NOT be sent after the `close` frame has been sent.

As an alternative to using an empty frame to prevent an idle timeout, if a connection is in a permissible state, an implementation MAY choose to send a flow frame for a valid session.

If during operation a peer exceeds the remote peer's idle timeout's threshold, e.g., because it is heavily loaded, it SHOULD gracefully close the connection by using a `close` frame with an error explaining why.

2.4.6 Connection States

START	In this state a connection exists, but nothing has been sent or received. This is the state an implementation would be in immediately after performing a <code>socket connect</code> or <code>socket accept</code> .
HDR_RCVD	In this state the connection header has been received from the peer but a connection header has not been sent.
HDR_SENT	In this state the connection header has been sent to the peer but no connection header has been received.

HDR_EXCH	In this state the connection header has been sent to the peer and a connection header has been received from the peer.
OPEN_PIPE	In this state both the connection header and the <code>open</code> frame have been sent but nothing has been received.
OC_PIPE	In this state, the connection header, the <code>open</code> frame, any pipelined connection traffic, and the <code>close</code> frame have been sent but nothing has been received.
OPEN_RCVD	In this state the connection headers have been exchanged. An <code>open</code> frame has been received from the peer but an <code>open</code> frame has not been sent.
OPEN_SENT	In this state the connection headers have been exchanged. An <code>open</code> frame has been sent to the peer but no <code>open</code> frame has yet been received.
CLOSE_PIPE	In this state the connection headers have been exchanged. An <code>open</code> frame, any pipelined connection traffic, and the <code>close</code> frame have been sent but no <code>open</code> frame has yet been received from the peer.
OPENED	In this state the connection header and the <code>open</code> frame have been both sent and received.
CLOSE_RCVD	In this state a <code>close</code> frame has been received indicating that the peer has initiated an AMQP close. No further frames are expected to arrive on the connection; however, frames can still be sent. If desired, an implementation MAY do a TCP half-close at this point to shut down the read side of the connection.
CLOSE_SENT	In this state a <code>close</code> frame has been sent to the peer. It is illegal to write anything more onto the connection, however there could potentially still be incoming frames. If desired, an implementation MAY do a TCP half-close at this point to shutdown the write side of the connection.
DISCARDING	The DISCARDING state is a variant of the CLOSE_SENT state where the <code>close</code> is triggered by an error. In this case any incoming frames on the connection MUST be silently discarded until the peer's <code>close</code> frame is received.
END	In this state it is illegal for either endpoint to write anything more onto the connection. The connection can be safely closed and discarded.

2.4.7 Connection State Diagram

The graph below depicts a complete state diagram for each endpoint. The boxes represent states, and the arrows represent state transitions. Each arrow is labeled with the action that triggers that particular transition.




Figure 2.23: Connection State Diagram

State	Legal Sends	Legal Receives	Legal Connection Actions
START	HDR	HDR	
HDR_RCVD	HDR	OPEN	
HDR_SENT	OPEN	HDR	
HDR_EXCH	OPEN	OPEN	
OPEN_RCVD	OPEN	*	
OPEN_SENT	**	OPEN	
OPEN_PIPE	**	HDR	
CLOSE_PIPE	-	OPEN	TCP Close for Write
OC_PIPE	-	HDR	TCP Close for Write
OPENED	*	*	
CLOSE_RCVD	*	-	TCP Close for Read
CLOSE_SENT	-	*	TCP Close for Write
DISCARDING	-	*	TCP Close for Write
END	-	-	TCP Close

* = any frames
 - = no frames
 ** = any frame known a priori to conform to the peer's capabilities and limitations

Figure 2.24: Connection State Table

2.5 Sessions

A session is a bidirectional sequential conversation between two containers that provides a grouping for related links. Sessions serve as the context for link communication. Any number of links of any directionality can be attached to a given session. However, a link MUST NOT be attached to more than one session at a time.




Figure 2.25: Instance Diagram of Session/Link attachment

Messages transferred on a link are sequentially identified within the session. A session can be viewed as multiplexing link traffic, much like a connection multiplexes session traffic. However, unlike the sessions on a connection, links on a session are not entirely independent since they share a common delivery sequence scoped to the session. This common sequence allows endpoints to efficiently refer to sets of deliveries regardless of the originating link. This is of particular benefit when a single application is receiving messages along a large number of different links. In this case the session provides *aggregation* of otherwise independent links into a single stream that can be efficiently acknowledged by the receiving application.

2.5.1 Establishing A Session

Sessions are established by creating a session endpoint, assigning it to an unused channel number, and sending a `begin` announcing the association of the session endpoint with the outgoing channel. Upon receiving the `begin` the partner will check the `remote-channel` field and find it empty. This indicates that the `begin` is referring to remotely initiated session. The partner will therefore allocate an unused outgoing channel for the remotely initiated session and indicate this by sending its own `begin` setting the `remote-channel` field to the incoming channel of the remotely initiated session.

To make it easier to monitor AMQP sessions, it is RECOMMENDED that implementations always assign the lowest available unused channel number.

The `remote-channel` field of a `begin` frame MUST be empty for a locally initiated session, and MUST be set when announcing the endpoint created as a result of a remotely initiated session.




Figure 2.26: Session Begin Sequence

2.5.2 Ending A Session

Sessions end automatically when the connection is closed or interrupted. Sessions are explicitly ended when either endpoint chooses to end the session. When a session is explicitly ended, an `end` frame is sent to announce the disassociation of the endpoint from its outgoing channel, and to carry error information when relevant.




Figure 2.27: Session End Sequence

2.5.3 Simultaneous End

Due to the potentially asynchronous nature of sessions, it is possible that both peers simultaneously decide to end a session. If this happens, it will appear to each peer as though their partner's spontaneously initiated end frame is actually an answer to the peers initial end frame.




Figure 2.28: Simultaneous Session End Sequence

2.5.4 Session Errors

When a session is unable to process input, it MUST indicate this by issuing an END with an appropriate `error` indicating the cause of the problem. It MUST then proceed to discard all incoming frames from the remote endpoint until receiving the remote endpoint's corresponding end frame.




Figure 2.29: Session Error Sequence

2.5.5 Session States

- UNMAPPED** In the UNMAPPED state, the session endpoint is not mapped to any incoming or outgoing channels on the connection endpoint. In this state an endpoint cannot send or receive frames.
- BEGIN_SENT** In the BEGIN_SENT state, the session endpoint is assigned an outgoing channel number, but there is no entry in the incoming channel map. In this state the endpoint MAY send frames but cannot receive them.
- BEGIN_RCVD** In the BEGIN_RCVD state, the session endpoint has an entry in the incoming channel map, but has not yet been assigned an outgoing channel number. The endpoint MAY receive frames, but cannot send them.
- MAPPED** In the MAPPED state, the session endpoint has both an outgoing channel number and an entry in the incoming channel map. The endpoint MAY both send and receive frames.
- END_SENT** In the END_SENT state, the session endpoint has an entry in the incoming channel map, but is no longer assigned an outgoing channel number. The endpoint MAY receive frames, but cannot send them.
- END_RCVD** In the END_RCVD state, the session endpoint is assigned an outgoing channel number, but there is no entry in the incoming channel map. The endpoint MAY send frames, but cannot receive them.
- DISCARDING** The DISCARDING state is a variant of the END_SENT state where the `end` is triggered by an error. In this case any incoming frames on the session MUST be silently discarded until the peer's `end` frame is received.

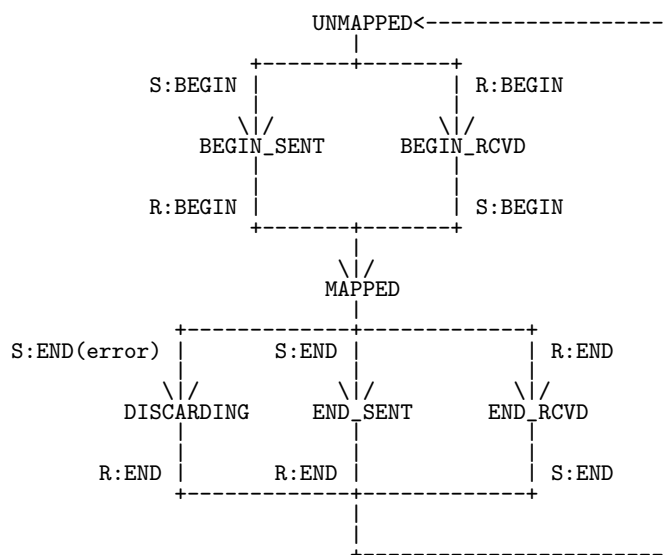


Figure 2.30: State Transitions

There is no obligation to retain a session endpoint after it transitions to the UNMAPPED state.

2.5.6 Session Flow Control

The session endpoint assigns each outgoing `transfer` frame an implicit `transfer-id` from a session scoped sequence. Each session endpoint maintains the following state to manage incoming and outgoing `transfer` frames:

next-incoming-id The `next-incoming-id` identifies the expected `transfer-id` of the next incoming `transfer` frame.

incoming-window The `incoming-window` defines the maximum number of incoming `transfer` frames that the endpoint can currently receive. This identifies a current maximum incoming `transfer-id` that can be computed by subtracting one from the sum of `incoming-window` and `next-incoming-id`.

next-outgoing-id The `next-outgoing-id` is the `transfer-id` to assign to the next `transfer` frame. The `next-outgoing-id` MAY be initialized to an arbitrary value and is incremented after each successive `transfer` according to RFC-1982 [RFC1982] serial number arithmetic.

outgoing-window The `outgoing-window` defines the maximum number of outgoing `transfer` frames that the endpoint can currently send. This identifies a current maximum outgoing `transfer-id` that can be computed by subtracting one from the sum of `outgoing-window` and `next-outgoing-id`.

remote-incoming-window

The `remote-incoming-window` reflects the maximum number of outgoing transfers that can be sent without exceeding the remote endpoint's `incoming-window`. This value MUST be decremented after every `transfer` frame is sent, and recomputed when informed of the remote session endpoint state.

remote-outgoing-window

The `remote-outgoing-window` reflects the maximum number of incoming transfers that MAY arrive without exceeding the remote endpoint's `outgoing-window`. This value MUST be decremented after every incoming `transfer` frame is received, and recomputed when informed of the remote session endpoint state. When this window shrinks, it is an indication of outstanding transfers. Settling outstanding transfers can cause the window to grow.

Once initialized, this state is updated by various events that occur in the lifespan of a session and its associated links:

sending a transfer Upon sending a transfer, the sending endpoint will increment its next-outgoing-id, decrement its remote-incoming-window, and MAY (depending on policy) decrement its outgoing-window.

receiving a transfer

Upon receiving a transfer, the receiving endpoint will increment the next-incoming-id to match the implicit transfer-id of the incoming transfer plus one, as well as decrementing the remote-outgoing-window, and MAY (depending on policy) decrement its incoming-window.

receiving a flow

When the endpoint receives a `flow` frame from its peer, it MUST update the *next-incoming-id* directly from the *next-outgoing-id* of the frame, and it MUST update the *remote-outgoing-window* directly from the *outgoing-window* of the frame.

The *remote-incoming-window* is computed as follows:

$$\text{next-incoming-id}_{\text{flow}} + \text{incoming-window}_{\text{flow}} - \text{next-outgoing-id}_{\text{endpoint}}$$

If the *next-incoming-id* field of the `flow` frame is not set, then *remote-incoming-window* is computed as follows:

$$\text{initial-outgoing-id}_{\text{endpoint}} + \text{incoming-window}_{\text{flow}} - \text{next-outgoing-id}_{\text{endpoint}}$$

2.6 Links

A link provides a unidirectional transport for messages between a source and a target. The primary responsibility of a source or target (a terminus) is to maintain a record of the status of each active delivery attempt until such a time as it is safe to forget. These are referred to as *unsettled* deliveries. When a terminus forgets the state associated with a *delivery-tag*, it is considered *settled*. Settling a delivery at a terminus is an idempotent idempotent, i.e., a delivery can transition from unsettled to settled, but never the reverse. Each delivery attempt is assigned a unique *delivery-tag* at the source. The status of an active delivery attempt is known as the *delivery state* of the delivery.

Link endpoints interface between a terminus and a session endpoint, and maintain additional state used for active communication between the local and remote endpoints. Therefore there are two types of endpoint: *senders* and *receivers*. When the sending application submits a message to the sender for transport, it also supplies the *delivery-tag* used by the source to track the delivery state. The link endpoint assigns each message a unique *delivery-id* from a session scoped sequence. These *delivery-ids* are used to efficiently reference subsets of the outstanding deliveries on a session.

Termini can exist beyond their associated link endpoints, so it is possible for a session to terminate and the termini to remain. A link is said to be *suspended* if the termini exist, but have no associated link endpoints. The process of associating new link endpoints with existing termini and re-establishing communication is referred to as *resuming* a link.

The original link endpoint state is not necessary for resumption of a link. Only the unsettled delivery state maintained at the termini is necessary for link resume, and this need not be stored directly. The form of *delivery-tags* is intentionally left open-ended so that they and their related delivery state can, if desired, be (re)constructed from application state, thereby minimizing or eliminating the need to retain additional protocol-specific state in order to resume a link.

2.6.1 Naming A Link

Links are named so that they can be recovered when communication is interrupted. Link names MUST uniquely identify the link amongst all links of the same direction between the two participating containers. Link names are

only used when attaching a link, so they can be arbitrarily long without a significant penalty.

A link's name uniquely identifies the link from the container of the source to the container of the target node, i.e., if the container of the source node is A, and the container of the target node is B, the link can be globally identified by the (ordered) tuple $(A,B,<name>)$. Consequently, a link can only be active in one connection at a time. If an attempt is made to attach the link subsequently when it is not suspended, then the link can be 'stolen', i.e., the second attach succeeds and the first attach MUST then be closed with a link error of `stolen`. This behavior ensures that in the event of a connection failure occurring and being noticed by one party, that re-establishment has the desired effect.

2.6.2 Link Handles

Each link endpoint is assigned a numeric handle used by the peer as a shorthand to refer to the link in all frames that reference the link (`attach`, `detach`, `flow`, `transfer`, `disposition`). This handle is assigned by the initial `attach` frame and remains in use until the link is detached. The two endpoints are not REQUIRED to use the same handle. This means a peer is free to independently chose its handle when a link endpoint is associated with the session. The locally chosen handle is referred to as the *output handle*. The remotely chosen handle is referred to as the *input handle*.

At an endpoint, a link is considered to be *attached* when the link endpoint exists and has both input and output handles assigned at an active session endpoint. A link is considered to be *detached* when the link endpoint exists, but is not assigned either input or output handles. A link can be considered *half attached* (or *half detached*) when only one of the input or output handles is assigned.




Figure 2.31: Link Handles

2.6.3 Establishing Or Resuming A Link

Links are established and/or resumed by creating a link endpoint associated with a local terminus, assigning it to an unused handle, and sending an `attach` frame. This frame carries the state of the newly created link endpoint, including the local and remote termini, one being the source and one being the target depending on the directionality of the link endpoint. On receipt of the `attach`, the remote session endpoint creates a corresponding link endpoint and informs its application of the attaching link. The application attempts to locate the terminus previously associated with the link. This terminus is associated with the link endpoint and can be updated if its properties do not match those sent by the remote link endpoint. If no such terminus exists, the application MAY choose to create one using the properties supplied by the remote link endpoint. The link endpoint is then mapped to an unused handle, and an `attach` frame is issued carrying the state of the newly created endpoint. Note that if the application chooses not to create a terminus, the session endpoint will still create a link endpoint and issue

an attach indicating that the link endpoint has no associated local terminus. In this case, the session endpoint MUST immediately detach the newly created link endpoint.




Figure 2.32: Establishing a Link

If there is no pre-existing terminus, and the peer does not wish to create a new one, this is indicated by setting the local terminus (source or target as appropriate) to null.




Figure 2.33: Refusing a Link

If either end of the link is already associated with a terminus, the attach frame MUST include its unsettled delivery state.

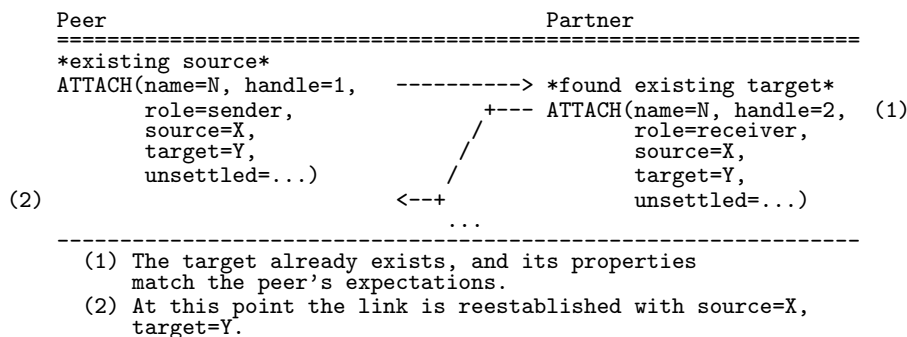


Figure 2.34: Resuming a Link

Note that it is possible that the expected terminus properties do not match the actual terminus properties reported by the remote endpoint. In this case, the link is always considered to be between the source as described by the

sender, and the target as described by the receiver. This can happen both when establishing and when resuming a link.

When a link is established, it is possible for an endpoint not to have all the capabilities necessary to create the terminus exactly matching the expectations of the peer. If this happens, the endpoint MAY adjust the properties in order to succeed in creating the terminus. In this case the endpoint MUST report the actual properties of the terminus as created.

When resuming a link, it is possible that the properties of the source and target have changed while the link was suspended. When this happens, the termini properties communicated in the source and target fields of the `attach` frames could be in conflict. In this case, the sender is considered to hold the authoritative version of the source properties, the receiver is considered to hold the authoritative version of the target properties. As above, the resulting link is constructed to be between the source as described by the sender, and the target as described by the receiver. Once the link is resumed, either peer is free to continue if the updated properties are acceptable, or, if not, detach the link.

Note that a peer MUST take responsibility for verifying that the remote terminus meets its requirements. The remote peer SHOULD NOT attempt to preempt whether the terminus will meet the requirements of its partner. This is equally true both for creating and resuming links.

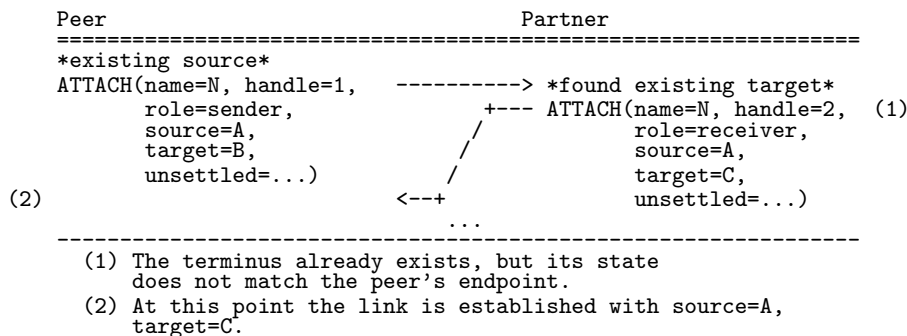


Figure 2.35: Resuming an altered Link

It is possible to resume a link even if one of the termini has lost nearly all its state. All that is necessary is the link name and direction. This is referred to as *recovering* a link. This is done by creating a new link endpoint with an empty source or target for incoming or outgoing links respectively. The full link state is then constructed from the authoritative source or target supplied by the other endpoint once the link is established. If the remote peer has no record of the link, then no terminus will be located, and local terminus (source or target as appropriate) field in the `attach` frame will be null.

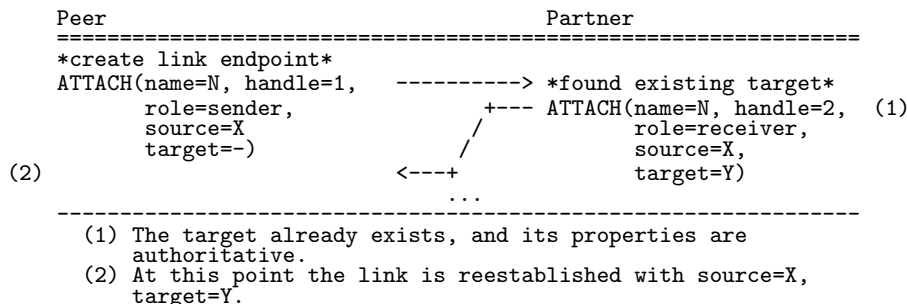


Figure 2.36: Recovering a Link

terminated with an `errant-link session-error`. Since the link endpoint has been destroyed, the peer cannot reattach, and MUST resume the link in order to restore communication. In order to disambiguate the resume request from a pipelined re-attach the resuming `attach` performative MUST contain a non-null value for its `unsettled` field. Receipt of a pipelined `attach` MUST result in the session being terminated with an `errant-link session-error`.

2.6.6 Closing A Link

A peer closes a link by sending the `detach` frame with the `handle` for the specified link, and the `closed` flag set to true. The partner will destroy the corresponding link endpoint, and reply with its own `detach` frame with the `closed` flag set to true.

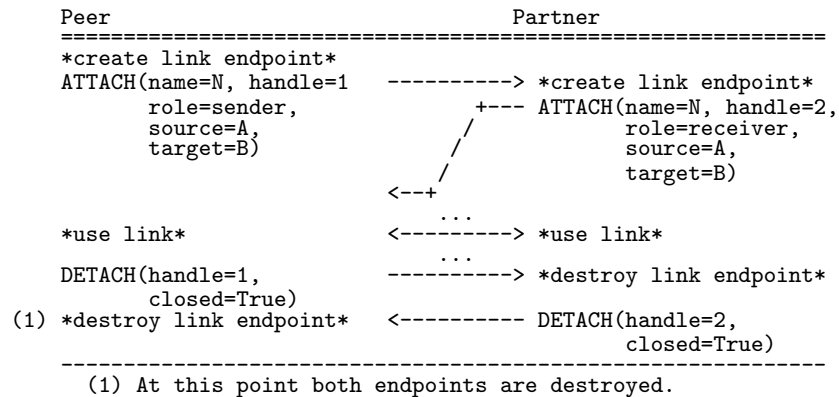


Figure 2.39: Closing a Link

Note that one peer MAY send a closing `detach` while its partner is sending a non-closing `detach`. In this case, the partner MUST signal that it has closed the link by reattaching and then sending a closing `detach`.

2.6.7 Flow Control

Once attached, a link is subject to flow control of message transfers. Link endpoints maintain the following flow control state. This state defines when it is legal to send transfers on an attached link, as well as indicating when certain interesting conditions occur, such as insufficient messages to consume the currently available `link-credit`, or insufficient `link-credit` to send available messages:

delivery-count The `delivery-count` is initialized by the sender when a link endpoint is created, and is incremented whenever a message is sent. Only the sender MAY independently modify this field. The receiver's value is calculated based on the last known value from the sender and any subsequent messages received on the link. Note that, despite its name, the `delivery-count` is not a count but a sequence number initialized at an arbitrary point by the sender.

link-credit The `link-credit` variable defines the current maximum legal amount that the `delivery-count` can be increased by. This identifies a `delivery-limit` that can be computed by adding the `link-credit` to the `delivery-count`.

Only the receiver can independently choose a value for this field. The sender's value MUST always be maintained in such a way as to match the `delivery-limit` identified by the receiver. This means that the sender's `link-credit` variable MUST be set according to this formula when flow information is given by the receiver:

$$\text{link-credit}_{\text{snd}} := \text{delivery-count}_{\text{rcv}} + \text{link-credit}_{\text{rcv}} - \text{delivery-count}_{\text{snd}}$$

In the event that the receiver does not yet know the *delivery-count*, i.e., $delivery-count_{RCV}$ is unspecified, the sender MUST assume that the $delivery-count_{RCV}$ is the first $delivery-count_{snd}$ sent from sender to receiver, i.e., the $delivery-count_{snd}$ specified in the flow state carried by the initial *attach* frame from the sender to the receiver.

Additionally, whenever the sender increases *delivery-count*, it MUST decrease *link-credit* by the same amount in order to maintain the *delivery-limit* identified by the receiver.

available

The *available* variable is controlled by the sender, and indicates to the receiver, that the sender could make use of the indicated amount of *link-credit*. Only the sender can independently modify this field. The receiver's value is calculated based on the last known value from the sender and any subsequent incoming messages received. The sender MAY transfer messages even if the available variable is zero. If this happens, the receiver MUST maintain a floor of zero in its calculation of the value of available.

drain

The drain flag indicates how the sender SHOULD behave when insufficient messages are available to consume the current link-credit. If set, the sender will (after sending all available messages) advance the delivery-count as much as possible, consuming all link-credit, and send the flow state to the receiver. Only the receiver can independently modify this field. The sender's value is always the last known value indicated by the receiver.

If the link-credit is less than or equal to zero, i.e., the delivery-count is the same as or greater than the delivery-limit, a sender MUST NOT send more messages. If the link-credit is reduced by the receiver when transfers are in-flight, the receiver MAY either handle the excess messages normally or detach the link with a transfer-limit-exceeded error code.

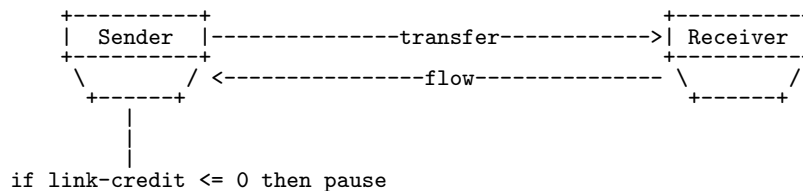


Figure 2.40: Flow Control

If the sender's drain flag is set and there are no available messages, the sender MUST advance its delivery-count until link-credit is zero, and send its updated *flow* state to the receiver.

The delivery-count is an absolute value. While the value itself is conceptually unbounded, it is encoded as a 32-bit integer that wraps around and compares according to RFC-1982 [RFC1982] serial number arithmetic.

The initial flow state of a link endpoint is determined as follows. The *link-credit* and *available* variables are initialized to zero. The *drain* flag is initialized to false. The sender MAY choose an arbitrary point to initialize the *delivery-count*. This value is communicated in the initial *attach* frame. The receiver initializes its *delivery-count* upon receiving the sender's *attach*.




Figure 2.41: Flow State & related Frames

The flow control semantics defined in this section provide the primitives necessary to implement a wide variety of flow control strategies. Additionally, by manipulating the link-credit and drain flag, a receiver can provide a variety of different higher level behaviors often useful to applications, including synchronous blocking fetch, synchronous fetch with a timeout, asynchronous notifications, and stopping/pausing.

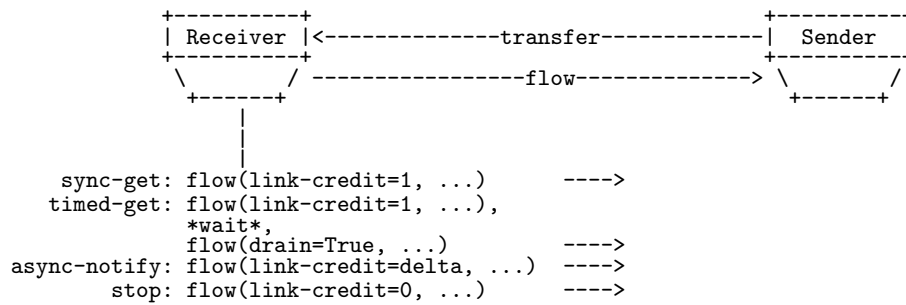


Figure 2.42: Flow Control Usage Patterns

2.6.8 Synchronous Get

A synchronous get of a message from a link is accomplished by incrementing the link-credit, sending the updated flow state, and waiting indefinitely for a transfer to arrive.




Figure 2.43: Synchronous Get

Synchronous get with a timeout is accomplished by incrementing the link-credit, sending the updated flow state and waiting for the link-credit to be consumed. When the desired time has elapsed the receiver then sets the drain flag and sends the newly updated flow state again, while continuing to wait for the link-credit to be consumed. Even if no messages are available, this condition will be met promptly because of the drain flag. Once the link-credit is consumed, the receiver can unambiguously determine whether a message has arrived or whether the operation has timed out.

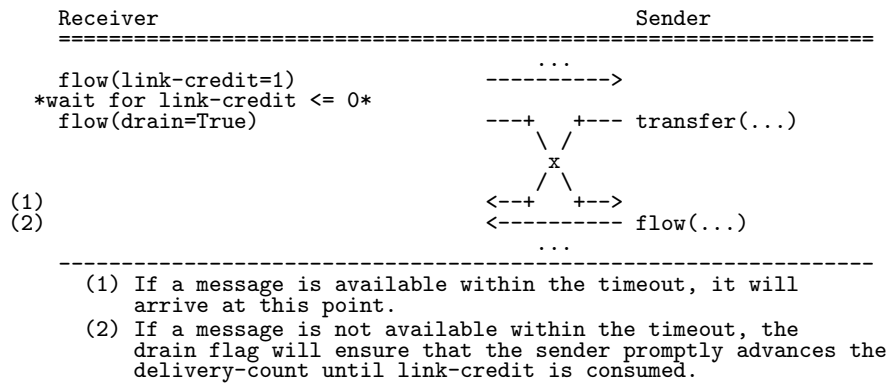


Figure 2.44: Synchronous Get w/ Timeout

2.6.9 Asynchronous Notification

Asynchronous notification can be accomplished as follows. The receiver maintains a target amount of link-credit for that link. As transfer arrive on the link, the sender’s link-credit decreases as the delivery-count increases. When the sender’s link-credit falls below a threshold, the flow state MAY be sent to increase the sender’s link-credit back to the desired target amount.

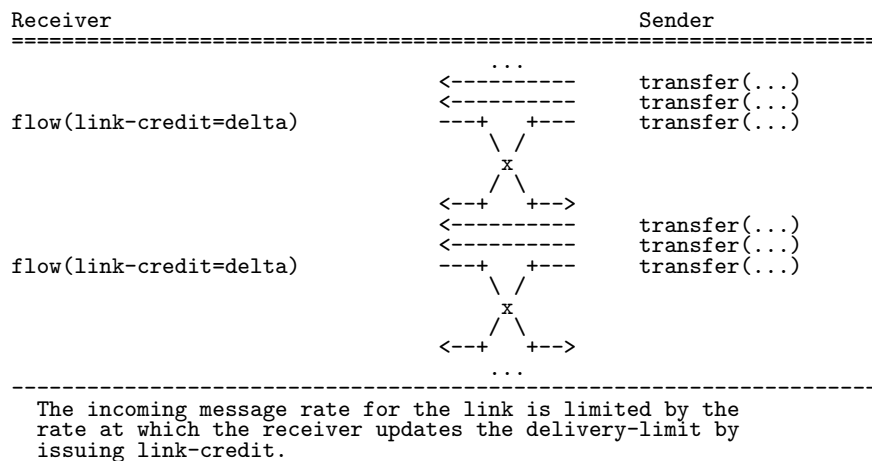


Figure 2.45: Asynchronous Notification

2.6.10 Stopping A Link

Stopping the transfers on a given link is accomplished by updating the link-credit to be zero and sending the updated flow state. It is possible that some transfers could be in flight at the time the flow state is sent, so incoming transfers could still arrive on the link. The echo field of the flow frame MAY be used to request the sender’s flow state be echoed back. This MAY be used to determine when the link has finally quiesced.




Figure 2.46: Stopping Incoming Messages

2.6.11 Messages

The transport layer assumes as little as possible about messages and allows alternative message representations to be layered above. Message data is carried as the payload in frames containing the `transfer` performative. Messages can be fragmented across several `transfer` frames as indicated by the more flag of the `transfer` performative.

2.6.12 Transferring A Message

When an application initiates a message transfer, it assigns a `delivery-tag` used to track the state of the delivery while the message is in transit. A delivery is considered *unsettled* at the sender/receiver from the point at which it was sent/received until it has been *settled* by the sending/receiving application. Each delivery **MUST** be identified by a `delivery-tag` chosen by the sending application. The `delivery-tag` **MUST** be unique amongst all deliveries that could be considered *unsettled* by either end of the link.

Upon initiating a transfer, the application will supply the sending link endpoint (Sender) with the message data and its associated `delivery-tag`. The sender will create an entry in its `unsettled` map, and send a `transfer` frame that includes the `delivery-tag`, the delivery's initial state, and its associated message data. For brevity on the wire, the `delivery-tag` is also associated with a `delivery-id` assigned by the session. The `delivery-id` is then used to refer to the `delivery-tag` in all subsequent interactions on that session.

The following diagrams illustrate the fundamentals involved in transferring a message. For normative semantics please refer to the definitions of the `transfer` and `disposition` performatives. For simplicity the `delivery-id` is omitted in the following diagrams and the `delivery-tag` is itself used directly. These diagrams also assume that this interaction takes place in the context of a single established link, and as such omit other details that would be present on the wire in practice such as the channel number, link handle, fragmentation flags, etc., focusing only on the essential aspects of message transfer.




Figure 2.47: Initial Transfer

Upon receiving the transfer, the receiving link endpoint (receiver) will create an entry in its own `unsettled` map and make the transferred message data available to the application to process.




Figure 2.48: Initial Receipt

Once notified of the received message data, the application processes the message, indicating the updated delivery state to the link endpoint as desired. Applications MAY wish to classify delivery states as *terminal* or *non-terminal* depending on whether an endpoint will ever update the state further once it has been reached. In some cases (e.g., large messages or transactions), the receiving application MAY wish to indicate non-terminal delivery states to the sender. This is done via the disposition frame.




Figure 2.49: Indication of Non-Terminal State

Once the receiving application has finished processing the message, it indicates to the link endpoint a *terminal* delivery state that reflects the outcome of the application processing (successful or otherwise) and thus the outcome which the receiver wishes to occur at the sender. This state is communicated back to the sender via the disposition frame.

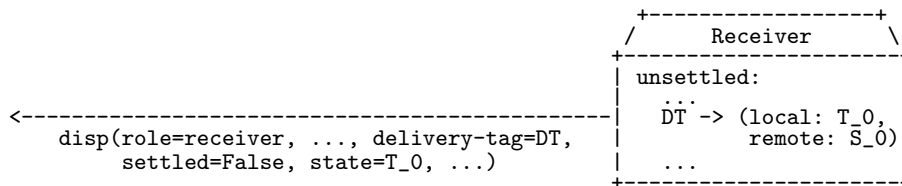


Figure 2.50: Indication of Presumptive Terminal State

Upon receiving the updated delivery state from the receiver, the sender will, if it has not already spontaneously attained a terminal state (e.g., through the expiry of the TTL at the sender), update its view of the state and communicate this back to the sending application.




Figure 2.51: Receipt of Terminal State

The sending application will then typically perform some action based on this terminal state and then settle the delivery, causing the sender to remove the delivery-tag from its unsettled map. The sender will then send its final delivery state along with an indication that the delivery is settled at the sender. Note that this amounts to the sender announcing that it is forever forgetting everything about the delivery-tag in question, and as such it is only possible to make such an announcement once, since after the sender forgets, it has no way of remembering to make the announcement again. If this frame gets lost due to an interruption in communication, the receiver will find out that the sender has settled the delivery upon link recovery. When the sender re-attaches the receiver will examine the unsettled state of the sender (i.e., what has **not** been forgotten) and from this can derive that the delivery in question has been settled (since its tag will not be in the unsettled state).




Figure 2.52: Indication of Settlement

When the receiver finds out that the sender has settled the delivery, the receiver will update its view of the remote state to indicate this, and then notify the receiving application.




Figure 2.53: Receipt of Settlement

The application can then perform some final action, e.g., remove the delivery-tag from a set kept for de-duplication, and then notify the receiver that the delivery is settled. The receiver will then remove the delivery-tag from its unsettled map. Note that because the receiver knows that the delivery is already settled at the sender, it makes no effort to notify the other endpoint that it is settling the delivery.



Figure 2.54: Final Settlement

As alluded to above, it is possible for the sending application to transition a delivery to a terminal state at the sender spontaneously (i.e., not as a consequence of a disposition that has been received from the receiver). In this case the sender **SHOULD** send a disposition to the receiver, but not settle until the receiver confirms, via a disposition in the opposite direction, that it has updated the state at its endpoint.

This set of exchanges illustrates the basic principals of message transfer. While a delivery is unsettled the endpoints exchange the current state of the delivery. Eventually both endpoints reach a terminal state as indicated by the application. This triggers the other application to take some final action and settle the delivery, and once one endpoint settles, this usually triggers the application at the other endpoint to settle.

This basic pattern can be modified in a variety of ways to achieve different guarantees. For example if the sending application settles the delivery *before* sending it, this results in an *at-most-once* guarantee. The sender has indicated up front with his initial transmission that he has forgotten everything about this delivery and will therefore make no further attempts to send it. If this delivery makes it to the receiver, the receiver clearly has no obligation to respond with updates of the receiver's delivery state, as they would be meaningless and ignored by the sender.




Figure 2.55: At-Most-Once

Similarly, if the basic scenario is modified such that the receiving application chooses to settle immediately upon processing the message rather than waiting for the sender to settle first, that yields an *at-least-once* guarantee. If the disposition frame indicated below is lost, then upon link recovery the sender will not see the delivery-tag in the receiver's unsettled map and will therefore assume the delivery was lost and resend it, resulting in duplicate processing of the message at the receiver.




Figure 2.56: At-Least-Once

As one might guess, the scenario presented initially where the sending application settles when the receiver reaches a terminal state, and the receiving application settles when the sender settles, results in an *exactly-once* guarantee. More generally if the receiver settles prior to the sender, it is possible for duplicate messages to occur, except in the case where the sender settles before the initial transmission. Similarly, if the sender settles before the receiver reaches a terminal state, it is possible for messages to be lost.

The sender and receiver policy regarding settling can either be preconfigured for the entire link, thereby allowing for optimized endpoint choices, or can be determined on an ad-hoc basis for each delivery. An application **MAY** also choose to settle at an endpoint independently of its delivery state, for example the sending application **MAY** choose to settle a delivery due to the message ttl expiring regardless of whether the receiver has reached a terminal state.

2.6.13 Resuming Deliveries

When a suspended link having unsettled deliveries is resumed, the *unsettled* field from the `attach` frame will carry the delivery-tags and delivery state of all deliveries considered unsettled by the issuing link endpoint. The set of delivery tags and delivery states contained in the unsettled maps from both endpoints can be divided into three categories:

Deliveries that only the source considers unsettled

Deliveries in this category **MAY** be resumed at the discretion of the sending application. If the sending application marks the resend attempt as a resumed delivery then it **MUST** be ignored by the receiver. (This allows the sender to pipeline resumes without risk of duplication at the sender).

Deliveries that only the target considers unsettled

Deliveries in this category **MUST** be ignored by the sender, and **MUST** be considered settled by the receiver.

Deliveries that both the source and target consider unsettled

Deliveries in this category **MUST** be resumed by the sender.

Note that in the case where an endpoint indicates that the unsettled map is incomplete, the absence of an entry in the unsettled map is not an indication of settlement. In this case the two endpoints **MUST** reduce the levels of unsettled state as much as they can by the sender resuming and/or settling transfers that it observes that the receiver considers unsettled. Upon completion of this reduction of state, the two parties **MUST** suspend and re-attempt to resume the link. Only when both sides have complete unsettled maps can new unsettled state be created by the sending of non-resuming transfers.

A delivery is resumed much the same way it is initially transferred with the following exceptions:

- The resume flag of the `transfer` frame **MUST** be set to true when resuming a delivery.
- The sender **MAY** omit message data when the delivery state of the receiver indicates retransmission is unnecessary.

Note that unsettled delivery-tags do **NOT** have any valid delivery-ids associated until they are resumed, as the delivery-ids from their original link endpoints are meaningless to the new link endpoints.

2.6.14 Transferring Large Messages

Each `transfer` frame can carry an arbitrary amount of message data up to the limit imposed by the maximum frame size. For messages that are too large to fit within the maximum frame size, additional data **MAY** be transferred in additional `transfer` frames by setting the more flag on all but the last `transfer` frame. When a message is split up into multiple `transfer` frames in this manner, messages being transferred along different links **MAY** be interleaved. However, messages transferred along a single link **MUST NOT** be interleaved.

The sender **MAY** indicate an aborted attempt to deliver a message by setting the abort flag on the last `transfer`. In this case the receiver **MUST** discard the message data that was transferred prior to the abort.




Figure 2.57: Outgoing Fragmentation State Diagram

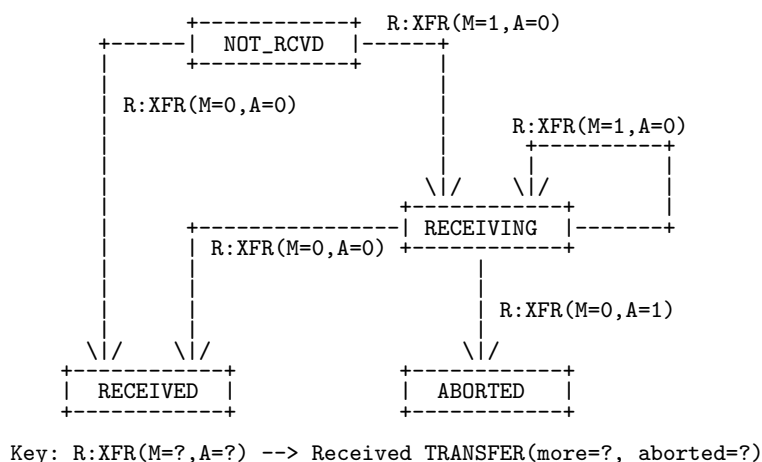


Figure 2.58: Incoming Fragmentation State Diagram

2.7 Performatives

2.7.1 Open

Negotiate connection parameters.

```
<type name="open" class="composite" source="list" provides="frame">
  <descriptor name="amqp:open:list" code="0x00000000:0x00000010"/>
  <field name="container-id" type="string" mandatory="true"/>
  <field name="hostname" type="string"/>
  <field name="max-frame-size" type="uint" default="4294967295"/>
  <field name="channel-max" type="ushort" default="65535"/>
  <field name="idle-time-out" type="milliseconds"/>
  <field name="outgoing-locales" type="ietf-language-tag" multiple="true"/>
  <field name="incoming-locales" type="ietf-language-tag" multiple="true"/>
  <field name="offered-capabilities" type="symbol" multiple="true"/>
  <field name="desired-capabilities" type="symbol" multiple="true"/>
  <field name="properties" type="fields"/>
</type>
```

The first frame sent on a connection in either direction **MUST** contain an open performative. Note that the connection header which is sent first on the connection is **not** a frame.

The fields indicate the capabilities and limitations of the sending peer.

Field Details

container-id	<i>the id of the source container</i>
hostname	<i>the name of the target host</i>

The name of the host (either fully qualified or relative) to which the sending peer is connecting. It is not mandatory to provide the hostname. If no hostname is provided the receiving peer SHOULD select a default based on its own configuration. This field can be used by AMQP proxies to determine the correct back-end service to connect the client to.

This field MAY already have been specified by the `sasl-init` frame, if a SASL layer is used, or, the server name indication extension as described in RFC-4366, if a TLS layer is used, in which case this field SHOULD be null or contain the same value. It is undefined what a different value to that already specified means.

`max-frame-size` *proposed maximum frame size*

The largest frame size that the sending peer is able to accept on this connection. If this field is not set it means that the peer does not impose any specific limit. A peer MUST NOT send frames larger than its partner can handle. A peer that receives an oversized frame MUST close the connection with the framing-error error-code.

Both peers MUST accept frames of up to 512 (MIN-MAX-FRAME-SIZE) octets.

`channel-max` *the maximum channel number that can be used on the connection*

The channel-max value is the highest channel number that can be used on the connection. This value plus one is the maximum number of sessions that can be simultaneously active on the connection. A peer MUST not use channel numbers outside the range that its partner can handle. A peer that receives a channel number outside the supported range MUST close the connection with the framing-error error-code.

`idle-time-out` *idle time-out*

The idle timeout REQUIRED by the sender (see subsection 2.4.5). A value of zero is the same as if it was not set (null). If the receiver is unable or unwilling to support the idle time-out then it SHOULD close the connection with an error explaining why (e.g., because it is too small).

If the value is not set, then the sender does not have an idle time-out. However, senders doing this SHOULD be aware that implementations MAY choose to use an internal default to efficiently manage a peer's resources.

`outgoing-locales` *locales available for outgoing text*

A list of the locales that the peer supports for sending informational text. This includes connection, session and link error descriptions. A peer MUST support at least the *en-US* locale (see subsection 2.8.12 IETF Language Tag). Since this value is always supported, it need not be supplied in the outgoing-locales. A null value or an empty list implies that only *en-US* is supported.

`incoming-locales` *desired locales for incoming text in decreasing level of preference*

A list of locales that the sending peer permits for incoming informational text. This list is ordered in decreasing level of preference. The receiving partner will choose the first (most preferred) incoming locale from those which it supports. If none of the requested locales are supported, *en-US* will be chosen. Note that *en-US* need not be supplied in this list as it is always the fallback. A peer MAY determine which of the permitted incoming locales is chosen by examining the partner's supported locales as specified in the outgoing-locales field. A null value or an empty list implies that only *en-US* is supported.

`offered-capabilities` *extension capabilities the sender supports*

If the receiver of the offered-capabilities requires an extension capability which is not present in the offered-capability list then it MUST close the connection.

A registry of commonly defined connection capabilities and their meanings is maintained [AMQP-CONNCAP].

`desired-capabilities` *extension capabilities the sender can use if the receiver supports them*

The desired-capability list defines which extension capabilities the sender MAY use if the receiver offers them (i.e., they are in the offered-capabilities list received by the sender of the desired-capabilities). The sender MUST NOT attempt to use any capabilities it did not declare in the desired-capabilities field. If the receiver of the desired-capabilities offers extension capabilities which are not present in the desired-capabilities list it received, then it can be sure those (undesired) capabilities will not be used on the connection.

properties *connection properties*

The properties map contains a set of fields intended to indicate information about the connection and its container.

A registry of commonly defined connection properties and their meanings is maintained [AMQP-CONNPROP].

2.7.2 Begin

Begin a session on a channel.

```
<type name="begin" class="composite" source="list" provides="frame">
  <descriptor name="amqp:begin:list" code="0x00000000:0x00000011"/>
  <field name="remote-channel" type="ushort"/>
  <field name="next-outgoing-id" type="transfer-number" mandatory="true"/>
  <field name="incoming-window" type="uint" mandatory="true"/>
  <field name="outgoing-window" type="uint" mandatory="true"/>
  <field name="handle-max" type="handle" default="4294967295"/>
  <field name="offered-capabilities" type="symbol" multiple="true"/>
  <field name="desired-capabilities" type="symbol" multiple="true"/>
  <field name="properties" type="fields"/>
</type>
```

Indicate that a session has begun on the channel.

Field Details

remote-channel *the remote channel for this session*

If a session is locally initiated, the remote-channel MUST NOT be set. When an endpoint responds to a remotely initiated session, the remote-channel MUST be set to the channel on which the remote session sent the begin.

next-outgoing-id *the transfer-id of the first transfer id the sender will send*

See subsection 2.5.6.

incoming-window *the initial incoming-window of the sender*

See subsection 2.5.6.

outgoing-window *the initial outgoing-window of the sender*

See subsection 2.5.6.

handle-max *the maximum handle value that can be used on the session*

The handle-max value is the highest handle value that can be used on the session. A peer MUST NOT attempt to attach a link using a handle value outside the range that its partner can handle. A peer that receives a handle outside the supported range MUST close the connection with the framing-error error-code.

offered-capabilities *the extension capabilities the sender supports*

A registry of commonly defined session capabilities and their meanings is maintained [AMQPSESSCAP].

`desired-capabilities` *the extension capabilities the sender can use if the receiver supports them*

The sender MUST NOT attempt to use any capability other than those it has declared in `desired-capabilities` field.

`properties` *session properties*

The properties map contains a set of fields intended to indicate information about the session and its container.

A registry of commonly defined session properties and their meanings is maintained [AMQPSESSPROP].

2.7.3 Attach

Attach a link to a session.

```
<type name="attach" class="composite" source="list" provides="frame">
  <descriptor name="amqp:attach:list" code="0x00000000:0x00000012"/>
  <field name="name" type="string" mandatory="true"/>
  <field name="handle" type="handle" mandatory="true"/>
  <field name="role" type="role" mandatory="true"/>
  <field name="snd-settle-mode" type="sender-settle-mode" default="mixed"/>
  <field name="rcv-settle-mode" type="receiver-settle-mode" default="first"/>
  <field name="source" type="*" requires="source"/>
  <field name="target" type="*" requires="target"/>
  <field name="unsettled" type="map"/>
  <field name="incomplete-unsettled" type="boolean" default="false"/>
  <field name="initial-delivery-count" type="sequence-no"/>
  <field name="max-message-size" type="ulong"/>
  <field name="offered-capabilities" type="symbol" multiple="true"/>
  <field name="desired-capabilities" type="symbol" multiple="true"/>
  <field name="properties" type="fields"/>
</type>
```

The attach frame indicates that a link endpoint has been attached to the session.

Field Details

`name` *the name of the link*

This name uniquely identifies the link from the container of the source to the container of the target node, e.g., if the container of the source node is A, and the container of the target node is B, the link MAY be globally identified by the (ordered) tuple (A,B,<name>).

`handle` *the handle for the link while attached*

The numeric handle assigned by the the peer as a shorthand to refer to the link in all performatives that reference the link until the it is detached. See subsection 2.6.2.

The handle MUST NOT be used for other open links. An attempt to attach using a handle which is already associated with a link MUST be responded to with an immediate `close` carrying a `handle-in-use session-error`.

To make it easier to monitor AMQP link attach frames, it is RECOMMENDED that implementations always assign the lowest available handle to this field.

`role` *role of the link endpoint*

The role being played by the peer, i.e., whether the peer is the sender or the receiver of messages on the link.

`snd-settle-mode` *settlement policy for the sender*

The delivery settlement policy for the sender. When set at the receiver this indicates the desired value for the settlement mode at the sender. When set at the sender this indicates the actual settlement mode in use. The sender SHOULD respect the receiver's desired settlement mode if the receiver initiates the attach exchange and the sender supports the desired mode.

`rcv-settle-mode` *the settlement policy of the receiver*

The delivery settlement policy for the receiver. When set at the sender this indicates the desired value for the settlement mode at the receiver. When set at the receiver this indicates the actual settlement mode in use. The receiver SHOULD respect the sender's desired settlement mode if the sender initiates the attach exchange and the receiver supports the desired mode.

`source` *the source for messages*

If no source is specified on an outgoing link, then there is no source currently attached to the link. A link with no source will never produce outgoing messages.

`target` *the target for messages*

If no target is specified on an incoming link, then there is no target currently attached to the link. A link with no target will never permit incoming messages.

`unsettled` *unsettled delivery state*

This is used to indicate any unsettled delivery states when a suspended link is resumed. The map is keyed by delivery-tag with values indicating the delivery state. The local and remote delivery states for a given delivery-tag MUST be compared to resolve any in-doubt deliveries. If necessary, deliveries MAY be resent, or resumed based on the outcome of this comparison. See subsection 2.6.13.

If the local unsettled map is too large to be encoded within a frame of the agreed maximum frame size then the session MAY be ended with the `frame-size-too-small` error. The endpoint SHOULD make use of the ability to send an incomplete unsettled map (see below) to avoid sending an error.

The unsettled map MUST NOT contain null valued keys.

When reattaching (as opposed to resuming), the unsettled map MUST be null.

`incomplete-unsettled`

If set to true this field indicates that the unsettled map provided is not complete. When the map is incomplete the recipient of the map cannot take the absence of a delivery tag from the map as evidence of settlement. On receipt of an incomplete unsettled map a sending endpoint MUST NOT send any new deliveries (i.e. deliveries where resume is not set to true) to its partner (and a receiving endpoint which sent an incomplete unsettled map MUST detach with an error on receiving a transfer which does not have the resume flag set to true).

Note that if this flag is set to true then the endpoints MUST detach and reattach at least once in order to send new deliveries. This flag can be useful when there are too many entries in the unsettled map to fit within a single frame. An endpoint can attach, resume, settle, and detach until enough unsettled state has been cleared for an attach where this flag is set to false.

`initial-delivery-count` *the sender's initial value for delivery-count*

This MUST NOT be null if role is sender, and it is ignored if the role is receiver. See subsection 2.6.7.

`max-message-size` *the maximum message size supported by the link endpoint*

This field indicates the maximum message size supported by the link endpoint. Any attempt to deliver a message larger than this results in a `message-size-exceeded` `link-error`. If this field is zero or unset, there is no maximum size imposed by the link endpoint.

`offered-capabilities` *the extension capabilities the sender supports*

A registry of commonly defined link capabilities and their meanings is maintained [AMQ-PLINKCAP].

`desired-capabilities` *the extension capabilities the sender can use if the receiver supports them*

The sender MUST NOT attempt to use any capability other than those it has declared in `desired-capabilities` field.

`properties` *link properties*

The `properties` map contains a set of fields intended to indicate information about the link and its container.

A registry of commonly defined link properties and their meanings is maintained [AMQ-PLINKPROP].

2.7.4 Flow

Update link state.

```
<type name="flow" class="composite" source="list" provides="frame">
  <descriptor name="amqp:flow:list" code="0x00000000:0x00000013"/>
  <field name="next-incoming-id" type="transfer-number"/>
  <field name="incoming-window" type="uint" mandatory="true"/>
  <field name="next-outgoing-id" type="transfer-number" mandatory="true"/>
  <field name="outgoing-window" type="uint" mandatory="true"/>
  <field name="handle" type="handle"/>
  <field name="delivery-count" type="sequence-no"/>
  <field name="link-credit" type="uint"/>
  <field name="available" type="uint"/>
  <field name="drain" type="boolean" default="false"/>
  <field name="echo" type="boolean" default="false"/>
  <field name="properties" type="fields"/>
</type>
```

Updates the flow state for the specified link.

Field Details

`next-incoming-id`

Identifies the expected transfer-id of the next incoming `transfer` frame. This value MUST be set if the peer has received the `begin` frame for the session, and MUST NOT be set if it has not. See subsection 2.5.6 for more details.

`incoming-window`

Defines the maximum number of incoming `transfer` frames that the endpoint can currently receive. See subsection 2.5.6 for more details.

`next-outgoing-id`

The transfer-id that will be assigned to the next outgoing `transfer` frame. See subsection 2.5.6 for more details.

outgoing-window

Defines the maximum number of outgoing `transfer` frames that the endpoint could potentially currently send, if it was not constrained by restrictions imposed by its peer's incoming-window. See subsection 2.5.6 for more details.

handle

If set, indicates that the flow frame carries flow state information for the local link endpoint associated with the given handle. If not set, the flow frame is carrying only information pertaining to the session endpoint.

If set to a handle that is not currently associated with an attached link, the recipient **MUST** respond by ending the session with an `unattached-handle` session error.

delivery-count *the endpoint's value for the delivery-count sequence number*

See subsection 2.6.7 for the definition of delivery-count.

When the handle field is not set, this field **MUST NOT** be set.

When the handle identifies that the flow state is being sent from the sender link endpoint to receiver link endpoint this field **MUST** be set to the current delivery-count of the link endpoint.

When the flow state is being sent from the receiver endpoint to the sender endpoint this field **MUST** be set to the last known value of the corresponding sending endpoint. In the event that the receiving link endpoint has not yet seen the initial `attach` frame from the sender this field **MUST NOT** be set.

link-credit *the current maximum number of messages that can be received*

The current maximum number of messages that can be handled at the receiver endpoint of the link. Only the receiver endpoint can independently set this value. The sender endpoint sets this to the last known value seen from the receiver. See subsection 2.6.7 for more details.

When the handle field is not set, this field **MUST NOT** be set.

available *the number of available messages*

The number of messages awaiting credit at the link sender endpoint. Only the sender can independently set this value. The receiver sets this to the last known value seen from the sender. See subsection 2.6.7 for more details.

When the handle field is not set, this field **MUST NOT** be set.

drain *indicates drain mode*

When flow state is sent from the sender to the receiver, this field contains the actual drain mode of the sender. When flow state is sent from the receiver to the sender, this field contains the desired drain mode of the receiver. See subsection 2.6.7 for more details.

When the handle field is not set, this field **MUST NOT** be set.

echo *request state from partner*

If set to true then the receiver **SHOULD** send its state at the earliest convenient opportunity.

If set to true, and the handle field is not set, then the sender only requires session endpoint state to be echoed, however, the receiver **MAY** fulfil this requirement by sending a flow performative carrying link-specific state (since any such flow also carries session state).

If a sender makes multiple requests for the same state before the receiver can reply, the receiver **MAY** send only one flow in return.

Note that if a peer responds to echo requests with flows which themselves have the echo field set to true, an infinite loop could result if its partner adopts the same policy (therefore such a policy **SHOULD** be avoided).

properties *link state properties*

A registry of commonly defined link state properties and their meanings is maintained [AMQP-LINKSTATEPROP].

When the handle field is not set, this field **MUST NOT** be set.

2.7.5 Transfer

Transfer a message.

```
<type name="transfer" class="composite" source="list" provides="frame">
  <descriptor name="amqp:transfer:list" code="0x00000000:0x00000014"/>
  <field name="handle" type="handle" mandatory="true"/>
  <field name="delivery-id" type="delivery-number"/>
  <field name="delivery-tag" type="delivery-tag"/>
  <field name="message-format" type="message-format"/>
  <field name="settled" type="boolean"/>
  <field name="more" type="boolean" default="false"/>
  <field name="rcv-settle-mode" type="receiver-settle-mode"/>
  <field name="state" type="*" requires="delivery-state"/>
  <field name="resume" type="boolean" default="false"/>
  <field name="aborted" type="boolean" default="false"/>
  <field name="batchable" type="boolean" default="false"/>
</type>
```

The transfer frame is used to send messages across a link. Messages **MAY** be carried by a single transfer up to the maximum negotiated frame size for the connection. Larger messages **MAY** be split across several transfer frames.

Field Details

handle

Specifies the link on which the message is transferred.

delivery-id *alias for delivery-tag*

The delivery-id **MUST** be supplied on the first transfer of a multi-transfer delivery. On continuation transfers the delivery-id **MAY** be omitted. It is an error if the delivery-id on a continuation transfer differs from the delivery-id on the first transfer of a delivery.

delivery-tag

Uniquely identifies the delivery attempt for a given message on this link. This field **MUST** be specified for the first transfer of a multi-transfer message and can only be omitted for continuation transfers. It is an error if the delivery-tag on a continuation transfer differs from the delivery-tag on the first transfer of a delivery.

message-format *indicates the message format*

This field **MUST** be specified for the first transfer of a multi-transfer message and can only be omitted for continuation transfers. It is an error if the message-format on a continuation transfer differs from the message-format on the first transfer of a delivery.

settled

If not set on the first (or only) transfer for a (multi-transfer) delivery, then the settled flag MUST be interpreted as being false. For subsequent transfers in a multi-transfer delivery if the settled flag is left unset then it MUST be interpreted as true if and only if the value of the settled flag on any of the preceding transfers was true; if no preceding transfer was sent with settled being true then the value when unset MUST be taken as false.

If the negotiated value for `snd-settle-mode` at attachment is `settled`, then this field MUST be true on at least one transfer frame for a delivery (i.e., the delivery MUST be settled at the sender at the point the delivery has been completely transferred).

If the negotiated value for `snd-settle-mode` at attachment is `unsettled`, then this field MUST be false (or unset) on every transfer frame for a delivery (unless the delivery is aborted).

`more` *indicates that the message has more content*

Note that if both the `more` and `aborted` fields are set to true, the `aborted` flag takes precedence. That is, a receiver SHOULD ignore the value of the `more` field if the transfer is marked as aborted. A sender SHOULD NOT set the `more` flag to true if it also sets the `aborted` flag to true.

`rcv-settle-mode`

If `first`, this indicates that the receiver MUST settle the delivery once it has arrived without waiting for the sender to settle first.

If `second`, this indicates that the receiver MUST NOT settle until sending its disposition to the sender and receiving a settled disposition from the sender.

If not set, this value is defaulted to the value negotiated on link attach.

If the negotiated link value is `first`, then it is illegal to set this field to `second`.

If the message is being sent settled by the sender, the value of this field is ignored.

The (implicit or explicit) value of this field does not form part of the transfer state, and is not retained if a link is suspended and subsequently resumed.

`state` *the state of the delivery at the sender*

When set this informs the receiver of the state of the delivery at the sender. This is particularly useful when transfers of unsettled deliveries are resumed after resuming a link. Setting the state on the transfer can be thought of as being equivalent to sending a disposition immediately before the transfer performative, i.e., it is the state of the delivery (not the transfer) that existed at the point the frame was sent.

Note that if the transfer performative (or an earlier disposition performative referring to the delivery) indicates that the delivery has attained a terminal state, then no future transfer or disposition sent by the sender can alter that terminal state.

`resume` *indicates a resumed delivery*

If true, the resume flag indicates that the transfer is being used to reassociate an unsettled delivery from a dissociated link endpoint. See subsection 2.6.13 for more details.

The receiver MUST ignore resumed deliveries that are not in its local unsettled map. The sender MUST NOT send resumed transfers for deliveries not in its local unsettled map.

If a resumed delivery spans more than one transfer performative, then the resume flag MUST be set to true on the first transfer of the resumed delivery. For subsequent transfers for the same delivery the resume flag MAY be set to true, or MAY be omitted.

In the case where the exchange of unsettled maps makes clear that all message data has been successfully transferred to the receiver, and that only the final state (and potentially settlement) at the sender needs to be conveyed, then a resumed delivery MAY carry no payload and instead act solely as a vehicle for carrying the terminal state of the delivery at the sender.

`aborted` *indicates that the message is aborted*

Aborted messages SHOULD be discarded by the recipient (any payload within the frame carrying the performative MUST be ignored). An aborted message is implicitly settled.

`batchable` *batchable hint*

If true, then the issuer is hinting that there is no need for the peer to urgently communicate updated delivery state. This hint MAY be used to artificially increase the amount of batching an implementation uses when communicating delivery states, and thereby save bandwidth.

If the message being delivered is too large to fit within a single frame, then the setting of `batchable` to true on any of the `transfer` performatives for the delivery is equivalent to setting `batchable` to true for all the `transfer` performatives for the delivery.

The `batchable` value does not form part of the transfer state, and is not retained if a link is suspended and subsequently resumed.

2.7.6 Disposition

Inform remote peer of delivery state changes.

```
<type name="disposition" class="composite" source="list" provides="frame">
  <descriptor name="amqp:disposition:list" code="0x00000000:0x00000015"/>
  <field name="role" type="role" mandatory="true"/>
  <field name="first" type="delivery-number" mandatory="true"/>
  <field name="last" type="delivery-number"/>
  <field name="settled" type="boolean" default="false"/>
  <field name="state" type="*" requires="delivery-state"/>
  <field name="batchable" type="boolean" default="false"/>
</type>
```

The disposition frame is used to inform the remote peer of local changes in the state of deliveries. The disposition frame MAY reference deliveries from many different links associated with a session, although all links MUST have the directionality indicated by the specified *role*.

Note that it is possible for a disposition sent from sender to receiver to refer to a delivery which has not yet completed (i.e., a delivery which is spread over multiple frames and not all frames have yet been sent). The use of such interleaving is discouraged in favor of carrying the modified state on the next `transfer` performative for the delivery.

The disposition performative MAY refer to deliveries on links that are no longer attached. As long as the links have not been closed or detached with an error then the deliveries are still “live” and the updated state MUST be applied.

Field Details

`role` *directionality of disposition*

The role identifies whether the disposition frame contains information about *sending* link endpoints or *receiving* link endpoints.

`first` *lower bound of deliveries*

Identifies the lower bound of delivery-ids for the deliveries in this set.

`last` *upper bound of deliveries*

Identifies the upper bound of delivery-ids for the deliveries in this set. If not set, this is taken to be the same as *first*.

`settled` *indicates deliveries are settled*

If true, indicates that the referenced deliveries are considered settled by the issuing endpoint.

`state` *indicates state of deliveries*

Communicates the state of all the deliveries referenced by this disposition.

`batchable` *batchable hint*

If true, then the issuer is hinting that there is no need for the peer to urgently communicate the impact of the updated delivery states. This hint MAY be used to artificially increase the amount of batching an implementation uses when communicating delivery states, and thereby save bandwidth.

2.7.7 Detach

Detach the link endpoint from the session.

```
<type name="detach" class="composite" source="list" provides="frame">
  <descriptor name="amqp:detach:list" code="0x00000000:0x00000016"/>
  <field name="handle" type="handle" mandatory="true"/>
  <field name="closed" type="boolean" default="false"/>
  <field name="error" type="error"/>
</type>
```

Detach the link endpoint from the session. This unmaps the handle and makes it available for use by other links.

Field Details

`handle` *the local handle of the link to be detached*

`closed` *if true then the sender has closed the link*

See subsection 2.6.6.

`error` *error causing the detach*

If set, this field indicates that the link is being detached due to an error condition. The value of the field SHOULD contain details on the cause of the error.

2.7.8 End

End the session.

```
<type name="end" class="composite" source="list" provides="frame">
  <descriptor name="amqp:end:list" code="0x00000000:0x00000017"/>
  <field name="error" type="error"/>
</type>
```

Indicates that the session has ended.

Field Details

`error` *error causing the end*

If set, this field indicates that the session is being ended due to an error condition. The value of the field SHOULD contain details on the cause of the error.

2.7.9 Close

Signal a connection close.

```
<type name="close" class="composite" source="list" provides="frame">
  <descriptor name="amqp:close:list" code="0x00000000:0x00000018"/>
  <field name="error" type="error"/>
</type>
```

Sending a close signals that the sender will not be sending any more frames (or bytes of any other kind) on the connection. Orderly shutdown requires that this frame **MUST** be written by the sender. It is illegal to send any more frames (or bytes of any other kind) after sending a close frame.

Field Details

error *error causing the close*

If set, this field indicates that the connection is being closed due to an error condition. The value of the field **SHOULD** contain details on the cause of the error.

2.8 Definitions

2.8.1 Role

Link endpoint role.

```
<type name="role" class="restricted" source="boolean">
  <choice name="sender" value="false"/>
  <choice name="receiver" value="true"/>
</type>
```

Valid Values

false	sender
true	receiver

2.8.2 Sender Settle Mode

Settlement policy for a sender.

```
<type name="sender-settle-mode" class="restricted" source="ubyte">
  <choice name="unsettled" value="0"/>
  <choice name="settled" value="1"/>
  <choice name="mixed" value="2"/>
</type>
```

Valid Values

0	The sender will send all deliveries initially unsettled to the receiver.
1	The sender will send all deliveries settled to the receiver.

- 2** The sender MAY send a mixture of settled and unsettled deliveries to the receiver.

2.8.3 Receiver Settle Mode

Settlement policy for a receiver.

```
<type name="receiver-settle-mode" class="restricted" source="ubyte">
  <choice name="first" value="0"/>
  <choice name="second" value="1"/>
</type>
```

Valid Values

- 0** The receiver will spontaneously settle all incoming transfers.
- 1** The receiver will only settle after sending the `disposition` to the sender and receiving a `disposition` indicating settlement of the delivery from the sender.

2.8.4 Handle

The handle of a link.

```
<type name="handle" class="restricted" source="uint"/>
```

An alias established by the `attach` frame and subsequently used by endpoints as a shorthand to refer to the link in all outgoing frames. The two endpoints MAY potentially use different handles to refer to the same link. Link handles MAY be reused once a link is closed for both send and receive.

2.8.5 Seconds

A duration measured in seconds.

```
<type name="seconds" class="restricted" source="uint"/>
```

2.8.6 Milliseconds

A duration measured in milliseconds.

```
<type name="milliseconds" class="restricted" source="uint"/>
```

2.8.7 Delivery Tag

```
<type name="delivery-tag" class="restricted" source="binary"/>
```

A delivery-tag can be up to 32 octets of binary data.

2.8.8 Delivery Number

```
<type name="delivery-number" class="restricted" source="sequence-no"/>
```


2.8.14 Error

Details of an error.

```
<type name="error" class="composite" source="list">
  <descriptor name="amqp:error:list" code="0x00000000:0x0000001d"/>
  <field name="condition" type="symbol" requires="error-condition" mandatory="true"/>
  <field name="description" type="string"/>
  <field name="info" type="fields"/>
</type>
```

Field Details

condition *error condition*

A symbolic value indicating the error condition.

description *descriptive text about the error condition*

This text supplies any supplementary details not indicated by the condition field. This text can be logged as an aid to resolving issues.

info *map carrying information about the error condition*

2.8.15 AMQP Error

Shared error conditions.

```
<type name="amqp-error" class="restricted" source="symbol" provides="error-condition">
  <choice name="internal-error" value="amqp:internal-error"/>
  <choice name="not-found" value="amqp:not-found"/>
  <choice name="unauthorized-access" value="amqp:unauthorized-access"/>
  <choice name="decode-error" value="amqp:decode-error"/>
  <choice name="resource-limit-exceeded" value="amqp:resource-limit-exceeded"/>
  <choice name="not-allowed" value="amqp:not-allowed"/>
  <choice name="invalid-field" value="amqp:invalid-field"/>
  <choice name="not-implemented" value="amqp:not-implemented"/>
  <choice name="resource-locked" value="amqp:resource-locked"/>
  <choice name="precondition-failed" value="amqp:precondition-failed"/>
  <choice name="resource-deleted" value="amqp:resource-deleted"/>
  <choice name="illegal-state" value="amqp:illegal-state"/>
  <choice name="frame-size-too-small" value="amqp:frame-size-too-small"/>
</type>
```

Valid Values

amqp:internal-error

An internal error occurred. Operator intervention might be necessary to resume normal operation.

amqp:not-found

A peer attempted to work with a remote entity that does not exist.

amqp:unauthorized-access

A peer attempted to work with a remote entity to which it has no access due to security settings.

amqp:decode-error

Data could not be decoded.

amqp:resource-limit-exceeded

A peer exceeded its resource allocation.

amqp:not-allowed

The peer tried to use a frame in a manner that is inconsistent with the semantics defined in the specification.

amqp:invalid-field

An invalid field was passed in a frame body, and the operation could not proceed.

amqp:not-implemented

The peer tried to use functionality that is not implemented in its partner.

amqp:resource-locked

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

amqp:precondition-failed

The client made a request that was not allowed because some precondition failed.

amqp:resource-deleted

A server entity the client is working with has been deleted.

amqp:illegal-state

The peer sent a frame that is not permitted in the current state.

amqp:frame-size-too-small

The peer cannot send a frame because the smallest encoding of the performative with the currently valid values would be too large to fit within a frame of the agreed maximum frame size. When transferring a message the message data can be sent in multiple `transfer` frames thereby avoiding this error. Similarly when attaching a link with a large unsettled map the endpoint MAY make use of the `incomplete-unsettled` flag to avoid the need for overly large frames.

2.8.16 Connection Error

Symbols used to indicate connection error conditions.

```
<type name="connection-error" class="restricted" source="symbol" provides="error-condition">
  <choice name="connection-forced" value="amqp:connection:forced"/>
  <choice name="framing-error" value="amqp:connection:framing-error"/>
  <choice name="redirect" value="amqp:connection:redirect"/>
</type>
```

Valid Values

amqp:connection:forced

An operator intervened to close the connection for some reason. The client could retry at some later date.

amqp:connection:framing-error

A valid frame header cannot be formed from the incoming byte stream.

amqp:connection:redirect

The container is no longer available on the current connection. The peer SHOULD attempt reconnection to the container using the details provided in the info map.

hostname the hostname of the container. This is the value that SHOULD be supplied in the *hostname* field of the open frame, and during the SASL and TLS negotiation (if used).

network-host the DNS hostname or IP address of the machine hosting the container.

port the port number on the machine hosting the container.

2.8.17 Session Error

Symbols used to indicate session error conditions.

```
<type name="session-error" class="restricted" source="symbol" provides="error-condition">
  <choice name="window-violation" value="amqp:session>window-violation"/>
  <choice name="errant-link" value="amqp:session:errant-link"/>
  <choice name="handle-in-use" value="amqp:session:handle-in-use"/>
  <choice name="unattached-handle" value="amqp:session:unattached-handle"/>
</type>
```

Valid Values

amqp:session>window-violation

The peer violated incoming window for the session.

amqp:session:errant-link

Input was received for a link that was detached with an error.

amqp:session:handle-in-use

An attach was received using a handle that is already in use for an attached link.

amqp:session:unattached-handle

A frame (other than attach) was received referencing a handle which is not currently in use of an attached link.

2.8.18 Link Error

Symbols used to indicate link error conditions.

```
<type name="link-error" class="restricted" source="symbol" provides="error-condition">
  <choice name="detach-forced" value="amqp:link:detach-forced"/>
  <choice name="transfer-limit-exceeded" value="amqp:link:transfer-limit-exceeded"/>
  <choice name="message-size-exceeded" value="amqp:link:message-size-exceeded"/>
  <choice name="redirect" value="amqp:link:redirect"/>
  <choice name="stolen" value="amqp:link:stolen"/>
</type>
```

Valid Values**amqp:link:detach-forced**

An operator intervened to detach for some reason.

amqp:link:transfer-limit-exceeded

The peer sent more message transfers than currently allowed on the link.

amqp:link:message-size-exceeded

The peer sent a larger message than is supported on the link.

amqp:link:redirect

The address provided cannot be resolved to a terminus at the current container. The info map MAY contain the following information to allow the client to locate the attach to the terminus.

hostname	the hostname of the container hosting the terminus. This is the value that SHOULD be supplied in the <i>hostname</i> field of the <code>open</code> frame, and during SASL and TLS negotiation (if used).
network-host	the DNS hostname or IP address of the machine hosting the container.
port	the port number on the machine hosting the container.
address	the address of the terminus at the container.

amqp:link:stolen

The link has been attached elsewhere, causing the existing attachment to be forcibly closed.

2.8.19 Constant Definitions

PORT	5672	the IANA assigned port number for AMQP. The standard AMQP port number that has been assigned by IANA for TCP, UDP, and SCTP. There are currently no UDP or SCTP mappings defined for AMQP. The port number is reserved for future transport mappings to these protocols.
SECURE-PORT	5671	the IANA assigned port number for secure AMQP (amqps). The standard AMQP port number that has been assigned by IANA for secure TCP using TLS. Implementations listening on this port SHOULD NOT expect a protocol handshake before TLS is negotiated.
MAJOR	1	major protocol version.
MINOR	0	minor protocol version.
REVISION	0	protocol revision.

MIN-MAX-FRAME-SIZE 512 the lower bound for the agreed maximum frame size (in bytes).
During the initial connection negotiation, the two peers MUST agree upon a maximum frame size. This constant defines the minimum value to which the maximum frame size can be set. By defining this value, the peers can guarantee that they can send frames of up to this size until they have agreed a definitive maximum frame size for that connection.

Part 3: Messaging

3.1 Introduction

The messaging layer builds on top of the concepts described in Part 1 and Part 2. The transport layer defines a number of extension points suitable for use in a variety of different messaging applications. The messaging layer specifies a standardized use of these to provide interoperable messaging capabilities. This standard covers:


- message format
 - properties for the bare message
 - formats for structured and unstructured sections in the bare message
 - headers and footers for the annotated message
- delivery states for messages traveling between nodes
- states for messages stored at a distribution node
- sources and targets
 - default disposition of transfers
 - supported outcomes
 - filtering of messages from a node
 - distribution-mode for access to messages stored at a distribution node
 - on-demand node creation

3.2 Message Format

The term message is used with various connotations in the messaging world. The sender might like to think of the message as an immutable payload handed off to the messaging infrastructure for delivery. The receiver often thinks of the message as not only that immutable payload from the sender, but also various annotations supplied by the messaging infrastructure along the way. To avoid confusion we formally define the term *bare message* to mean the message as supplied by the sender and the term *annotated message* to mean the message as seen at the receiver.

An *annotated message* consists of the bare message plus sections for annotation at the head and tail of the bare message. There are two classes of annotations: annotations that travel with the message indefinitely, and annotations that are consumed by the next node.

The *bare message* consists of three sections: standard properties, application-properties, and application-data (the body).



The bare message is immutable within the AMQP network. That is, none of the sections can be changed by any node acting as an AMQP intermediary. If a section of the bare message is omitted, one **MUST NOT** be inserted by an intermediary. The exact encoding of sections of the bare message **MUST NOT** be modified. This preserves message hashes, HMACs and signatures based on the binary encoding of the bare message.

The exact structure of a message, together with its encoding, is defined by the message format. This document defines the structure and semantics of message format 0 (MESSAGE-FORMAT). Altogether a message consists of the following sections:

- Zero or one header.
- Zero or one delivery-annotations.
- Zero or one message-annotations.
- Zero or one properties.
- Zero or one application-properties.
- The body consists of one of the following three choices: one or more data sections, one or more amqp-sequence sections, or a single amqp-value section.
- Zero or one footer.

3.2.1 Header

Transport headers for a message.

```
<type name="header" class="composite" source="list" provides="section">
  <descriptor name="amqp:header:list" code="0x00000000:0x00000070"/>
  <field name="durable" type="boolean" default="false"/>
  <field name="priority" type="ubyte" default="4"/>
  <field name="ttl" type="milliseconds"/>
  <field name="first-acquirer" type="boolean" default="false"/>
  <field name="delivery-count" type="uint" default="0"/>
</type>
```

The header section carries standard delivery details about the transfer of a message through the AMQP network. If the header section is omitted the receiver **MUST** assume the appropriate default values (or the meaning implied by no value being set) for the fields within the `header` unless other target or node specific defaults have otherwise been set.

Field Details

`durable` *specify durability requirements*

Durable messages MUST NOT be lost even if an intermediary is unexpectedly terminated and restarted. A target which is not capable of fulfilling this guarantee MUST NOT accept messages where the durable header is set to true: if the source allows the rejected outcome then the message SHOULD be rejected with the `precondition-failed` error, otherwise the link MUST be detached by the receiver with the same error.

`priority` *relative message priority*

This field contains the relative message priority. Higher numbers indicate higher priority messages. Messages with higher priorities MAY be delivered before those with lower priorities.

An AMQP intermediary implementing distinct priority levels MUST do so in the following manner:

- If n distinct priorities are implemented and n is less than 10 – priorities 0 to $(5 - \text{ceiling}(n/2))$ MUST be treated equivalently and MUST be the lowest effective priority. The priorities $(4 + \text{floor}(n/2))$ and above MUST be treated equivalently and MUST be the highest effective priority. The priorities $(5 - \text{ceiling}(n/2))$ to $(4 + \text{floor}(n/2))$ inclusive MUST be treated as distinct priorities.
- If n distinct priorities are implemented and n is 10 or greater – priorities 0 to $(n - 1)$ MUST be distinct, and priorities n and above MUST be equivalent to priority $(n - 1)$.

Thus, for example, if 2 distinct priorities are implemented, then levels 0 to 4 are equivalent, and levels 5 to 9 are equivalent and levels 4 and 5 are distinct. If 3 distinct priorities are implemented the 0 to 3 are equivalent, 5 to 9 are equivalent and 3, 4 and 5 are distinct.

This scheme ensures that if two priorities are distinct for a server which implements m separate priority levels they are also distinct for a server which implements n different priority levels where $n > m$.

`ttl` *time to live in ms*

Duration in milliseconds for which the message is to be considered “live”. If this is set then a message expiration time will be computed based on the time of arrival at an intermediary. Messages that live longer than their expiration time will be discarded (or dead lettered). When a message is transmitted by an intermediary that was received with a `ttl`, the transmitted message’s header SHOULD contain a `ttl` that is computed as the difference between the current time and the formerly computed message expiration time, i.e., the reduced `ttl`, so that messages will eventually die if they end up in a delivery loop.

`first-acquirer`

If this value is true, then this message has not been acquired by any other link (see section 3.3). If this value is false, then this message MAY have previously been acquired by another link or links.

`delivery-count` *the number of prior unsuccessful delivery attempts*

The number of unsuccessful previous attempts to deliver this message. If this value is non-zero it can be taken as an indication that the delivery might be a duplicate. On first delivery, the value is zero. It is incremented upon an outcome being settled at the sender, according to rules defined for each outcome.

3.2.2 Delivery Annotations

```
<type name="delivery-annotations" class="restricted" source="annotations" provides="section">
  <descriptor name="amqp:delivery-annotations:map" code="0x00000000:0x00000071"/>
</type>
```

The `delivery-annotations` section is used for delivery-specific non-standard properties at the head of the message. Delivery annotations convey information from the sending peer to the receiving peer. If the recipient does not understand the annotation it cannot be acted upon and its effects (such as any implied propagation) cannot be acted upon. Annotations might be specific to one implementation, or common to multiple implementations. The

capabilities negotiated on link `attach` and on the `source` and `target` SHOULD be used to establish which annotations a peer supports. A registry of defined annotations and their meanings is maintained [AMQPDELANN]. The symbolic key “rejected” is reserved for the use of communicating error information regarding rejected messages. Any values associated with the “rejected” key MUST be of type `error`.

If the `delivery-annotations` section is omitted, it is equivalent to a `delivery-annotations` section containing an empty map of annotations.

3.2.3 Message Annotations

```
<type name="message-annotations" class="restricted" source="annotations" provides="section">
  <descriptor name="amqp:message-annotations:map" code="0x00000000:0x00000072"/>
</type>
```

The `message-annotations` section is used for properties of the message which are aimed at the infrastructure and SHOULD be propagated across every delivery step. Message annotations convey information about the message. Intermediaries MUST propagate the annotations unless the annotations are explicitly augmented or modified (e.g., by the use of the `modified` outcome).

The capabilities negotiated on link `attach` and on the `source` and `target` can be used to establish which annotations a peer understands; however, in a network of AMQP intermediaries it might not be possible to know if every intermediary will understand the annotation. Note that for some annotations it might not be necessary for the intermediary to understand their purpose, i.e., they could be used purely as an attribute which can be filtered on.

A registry of defined annotations and their meanings is maintained [AMQPMESSANN].

If the `message-annotations` section is omitted, it is equivalent to a `message-annotations` section containing an empty map of annotations.

3.2.4 Properties

Immutable properties of the message.

```
<type name="properties" class="composite" source="list" provides="section">
  <descriptor name="amqp:properties:list" code="0x00000000:0x00000073"/>
  <field name="message-id" type="*" requires="message-id"/>
  <field name="user-id" type="binary"/>
  <field name="to" type="*" requires="address"/>
  <field name="subject" type="string"/>
  <field name="reply-to" type="*" requires="address"/>
  <field name="correlation-id" type="*" requires="message-id"/>
  <field name="content-type" type="symbol"/>
  <field name="content-encoding" type="symbol"/>
  <field name="absolute-expiry-time" type="timestamp"/>
  <field name="creation-time" type="timestamp"/>
  <field name="group-id" type="string"/>
  <field name="group-sequence" type="sequence-no"/>
  <field name="reply-to-group-id" type="string"/>
</type>
```

The `properties` section is used for a defined set of standard properties of the message. The `properties` section is part of the bare message; therefore, if retransmitted by an intermediary, it MUST remain unaltered.

Field Details

<code>message-id</code>	<i>application message identifier</i>
-------------------------	---------------------------------------

Message-id, if set, uniquely identifies a message within the message system. The message producer is usually responsible for setting the message-id in such a way that it is assured to be globally unique. A broker MAY discard a message as a duplicate if the value of the message-id matches that of a previously received message sent to the same node.

`user-id` *creating user id*

The identity of the user responsible for producing the message. The client sets this value, and it MAY be authenticated by intermediaries.

`to` *the address of the node the message is destined for*

The to field identifies the node that is the intended destination of the message. On any given transfer this might not be the node at the receiving end of the link.

`subject` *the subject of the message*

A common field for summary information about the message content and purpose.

`reply-to` *the node to send replies to*

The address of the node to send replies to.

`correlation-id` *application correlation identifier*

This is a client-specific id that can be used to mark or identify messages between clients.

`content-type` *MIME content type*

The RFC-2046 [RFC2046] MIME type for the message's application-data section (body). As per RFC-2046 [RFC2046] this can contain a charset parameter defining the character encoding used: e.g., 'text/plain; charset="utf-8"'.
For clarity, as per section 7.2.1 of RFC-2616 [RFC2616], where the content type is unknown the content-type SHOULD NOT be set. This allows the recipient the opportunity to determine the actual type. Where the section is known to be truly opaque binary data, the content-type SHOULD be set to application/octet-stream.

When using an application-data section with a section code other than *data*, content-type SHOULD NOT be set.

`content-encoding` *MIME content type*

The content-encoding property is used as a modifier to the content-type. When present, its value indicates what additional content encodings have been applied to the application-data, and thus what decoding mechanisms need to be applied in order to obtain the media-type referenced by the content-type header field.

Content-encoding is primarily used to allow a document to be compressed without losing the identity of its underlying content type.

Content-encodings are to be interpreted as per section 3.5 of RFC 2616 [RFC2616]. Valid content-encodings are registered at IANA [IANAHTTPPARAMS].

The content-encoding MUST NOT be set when the application-data section is other than *data*. The binary representation of all other application-data section types is defined completely in terms of the AMQP type system.

Implementations MUST NOT use the *identity* encoding. Instead, implementations SHOULD NOT set this property. Implementations SHOULD NOT use the *compress* encoding, except as to remain compatible with messages originally sent with other protocols, e.g. HTTP or SMTP.

Implementations SHOULD NOT specify multiple content-encoding values except as to be compatible with messages originally sent with other protocols, e.g. HTTP or SMTP.

`absolute-expiry-time` *the time when this message is considered expired*

An absolute time when this message is considered to be expired.

<code>creation-time</code>	<i>the time when this message was created</i>
	An absolute time when this message was created.
<code>group-id</code>	<i>the group this message belongs to</i>
	Identifies the group the message belongs to.
<code>group-sequence</code>	<i>the sequence-no of this message within its group</i>
	The relative position of this message within its group.
<code>reply-to-group-id</code>	<i>the group the reply message belongs to</i>
	This is a client-specific id that is used so that client can send replies to this message to a specific group.

3.2.5 Application Properties

```
<type name="application-properties" class="restricted" source="map" provides="section">
  <descriptor name="amqp:application-properties:map" code="0x00000000:0x00000074"/>
</type>
```

The application-properties section is a part of the bare message used for structured application data. Intermediaries can use the data within this structure for the purposes of filtering or routing.

The keys of this map are restricted to be of type `string` (which excludes the possibility of a null key) and the values are restricted to be of simple types only, that is, excluding `map`, `list`, and array types.

3.2.6 Data

```
<type name="data" class="restricted" source="binary" provides="section">
  <descriptor name="amqp:data:binary" code="0x00000000:0x00000075"/>
</type>
```

A data section contains opaque binary data.

3.2.7 AMQP Sequence

```
<type name="amqp-sequence" class="restricted" source="list" provides="section">
  <descriptor name="amqp:amqp-sequence:list" code="0x00000000:0x00000076"/>
</type>
```

A sequence section contains an arbitrary number of structured data elements.

3.2.8 AMQP Value

```
<type name="amqp-value" class="restricted" source="*" provides="section">
  <descriptor name="amqp:amqp-value:*" code="0x00000000:0x00000077"/>
</type>
```

An amqp-value section contains a single AMQP value.

3.2.9 Footer

Transport footers for a message.

```
<type name="footer" class="restricted" source="annotations" provides="section">
  <descriptor name="amqp:footer:map" code="0x00000000:0x00000078"/>
</type>
```

The footer section is used for details about the message or delivery which can only be calculated or evaluated once the whole bare message has been constructed or seen (for example message hashes, HMACs, signatures and encryption details).

A registry of defined footers and their meanings is maintained [AMQPFOOTER].

3.2.10 Annotations

```
<type name="annotations" class="restricted" source="map"/>
```

The *annotations* type is a map where the keys are restricted to be of type `symbol` or of type `ulong`. All `ulong` keys, and all symbolic keys except those beginning with “x-” are reserved. Keys beginning with “x-opt-” MUST be ignored if not understood. On receiving an annotation key which is not understood, and which does not begin with “x-opt”, the receiving AMQP container MUST detach the link with a `not-implemented` error.

3.2.11 Message ID ULong

```
<type name="message-id-ulong" class="restricted" source="ulong" provides="message-id"/>
```

3.2.12 Message ID UUID

```
<type name="message-id-uuid" class="restricted" source="uuid" provides="message-id"/>
```

3.2.13 Message ID Binary

```
<type name="message-id-binary" class="restricted" source="binary" provides="message-id"/>
```

3.2.14 Message ID String

```
<type name="message-id-string" class="restricted" source="string" provides="message-id"/>
```

3.2.15 Address String

Address of a node.

```
<type name="address-string" class="restricted" source="string" provides="address"/>
```

3.2.16 Constant Definitions

MESSAGE-FORMAT	0	the format + revision for the messages defined by this document. This value goes into the message-format field of the transfer frame when transferring messages of the format defined herein.
----------------	---	--

3.3 Distribution Nodes

The messaging layer defines a set of states for a *distribution node*, defined as a node that stores messages for distribution. Not all nodes are distribution nodes; however, these definitions permit some standardized interaction with those nodes that do. The transitions between these states are controlled by the transfer of messages to/from a distribution node and the resulting terminal delivery state. Note that the state of a message at one distribution node does not affect the state of the same message at a separate node.

By default a message will begin in the AVAILABLE state. Prior to initiating an *acquiring* transfer, the message will transition to the ACQUIRED state. Once in the ACQUIRED state, a message is ineligible for *acquiring* transfers to any other links.

A message will remain ACQUIRED at the distribution node until the transfer is settled. The delivery state at the receiver determines how the message transitions when the transfer is settled. If the delivery state at the receiver is not yet known, (e.g., the link endpoint is destroyed before recovery occurs) the *default-outcome* of the source is used (see *source*).

State transitions can also occur spontaneously at the distribution node. For example if a message with a *ttl* expires, the effect of expiry might be (depending on specific type and configuration of the distribution node) to move spontaneously from the AVAILABLE state into the ARCHIVED state. In this case any transfers of the message are transitioned to a terminal outcome at the distribution node regardless of receiver state.




Figure 3.1: Message State Transitions

3.4 Delivery State

The messaging layer defines a concrete set of delivery states which can be used (via the *disposition* frame) to indicate the state of the message at the receiver. Delivery states can be either terminal or non-terminal. Once a delivery reaches a terminal delivery state, the state for that delivery will no longer change. A terminal delivery state is referred to as an *outcome*.

The following outcomes are formally defined by the messaging layer to indicate the result of processing at the receiver:

- **accepted**: indicates successful processing at the receiver.
- **rejected**: indicates an invalid and unprocessable message.
- **released**: indicates that the message was not (and will not be) processed.

- `modified`: indicates that the message was modified, but not processed.

The following non-terminal delivery-state is formally defined by the messaging layer for use during link recovery to allow the sender to resume the transfer of a large message without retransmitting all the message data:

- `received`: indicates partial message data seen by the receiver as well as the starting point for a resumed transfer.

3.4.1 Received

```
<type name="received" class="composite" source="list" provides="delivery-state">
  <descriptor name="amqp:received:list" code="0x00000000:0x00000023"/>
  <field name="section-number" type="uint" mandatory="true"/>
  <field name="section-offset" type="ulong" mandatory="true"/>
</type>
```

At the target the `received` state indicates the furthest point in the payload of the message which the target will not need to have resent if the link is resumed. At the source the `received` state represents the earliest point in the payload which the sender is able to resume transferring at in the case of link resumption. When resuming a delivery, if this state is set on the first `transfer` performative it indicates the offset in the payload at which the first resumed delivery is starting. The sender **MUST NOT** send the `received` state on `transfer` or `disposition` performatives except on the first `transfer` performative on a resumed delivery.

Field Details

`section-number`

When sent by the sender this indicates the first section of the message (with `section-number` 0 being the first section) for which data can be resent. Data from sections prior to the given section cannot be retransmitted for this delivery.

When sent by the receiver this indicates the first section of the message for which all data might not yet have been received.

`section-offset`

When sent by the sender this indicates the first byte of the encoded section data of the section given by `section-number` for which data can be resent (with `section-offset` 0 being the first byte). Bytes from the same section prior to the given offset section cannot be retransmitted for this delivery.

When sent by the receiver this indicates the first byte of the given section which has not yet been received. Note that if a receiver has received all of section number `X` (which contains `N` bytes of data), but none of section number `X + 1`, then it can indicate this by sending either `Received(section-number=X, section-offset=N)` or `Received(section-number=X+1, section-offset=0)`. The state `Received(section-number=0, section-offset=0)` indicates that no message data at all has been transferred.

3.4.2 Accepted

The accepted outcome.

```
<type name="accepted" class="composite" source="list" provides="delivery-state, outcome">
  <descriptor name="amqp:accepted:list" code="0x00000000:0x00000024"/>
</type>
```

At the source the `accepted` state means that the message has been retired from the node, and transfer of payload data will not be able to be resumed if the link becomes suspended. A delivery can become accepted at the source

even before all transfer frames have been sent, this does not imply that the remaining transfers for the delivery will not be sent - only the aborted flag on the `transfer` performative can be used to indicate a premature termination of the transfer.

At the target, the accepted outcome is used to indicate that an incoming message has been successfully processed, and that the receiver of the message is expecting the sender to transition the delivery to the accepted state at the source.

The accepted outcome does not increment the *delivery-count* in the header of the accepted message.

3.4.3 Rejected

The rejected outcome.

```
<type name="rejected" class="composite" source="list" provides="delivery-state, outcome">
  <descriptor name="amqp:rejected:list" code="0x00000000:0x00000025"/>
  <field name="error" type="error"/>
</type>
```

At the target, the rejected outcome is used to indicate that an incoming message is invalid and therefore unprocessable. The rejected outcome when applied to a message will cause the *delivery-count* to be incremented in the header of the rejected message.

At the source, the rejected outcome means that the target has informed the source that the message was rejected, and the source has taken the necessary action. The delivery SHOULD NOT ever spontaneously attain the rejected state at the source.

Field Details

`error` *error that caused the message to be rejected*

This field contains diagnostic information about the cause of the message rejection.

3.4.4 Released

The released outcome.

```
<type name="released" class="composite" source="list" provides="delivery-state, outcome">
  <descriptor name="amqp:released:list" code="0x00000000:0x00000026"/>
</type>
```

At the source the released outcome means that the message is no longer acquired by the receiver, and has been made available for (re-)delivery to the same or other targets receiving from the node. The message is unchanged at the node (i.e., the *delivery-count* of the header of the released message MUST NOT be incremented). As released is a terminal outcome, transfer of payload data will not be able to be resumed if the link becomes suspended. A delivery can become released at the source even before all transfer frames have been sent. This does not imply that the remaining transfers for the delivery will not be sent. The source MAY spontaneously attain the released outcome for a message (for example the source might implement some sort of time-bound acquisition lock, after which the acquisition of a message at a node is revoked to allow for delivery to an alternative consumer).

At the target, the released outcome is used to indicate that a given transfer was not and will not be acted upon.

3.4.5 Modified

The modified outcome.

```
<type name="modified" class="composite" source="list" provides="delivery-state, outcome">
  <descriptor name="amqp:modified:list" code="0x00000000:0x00000027"/>
  <field name="delivery-failed" type="boolean"/>
  <field name="undeliverable-here" type="boolean"/>
  <field name="message-annotations" type="fields"/>
</type>
```

At the source the modified outcome means that the message is no longer acquired by the receiver, and has been made available for (re-)delivery to the same or other targets receiving from the node. The message has been changed at the node in the ways indicated by the fields of the outcome. As modified is a terminal outcome, transfer of payload data will not be able to be resumed if the link becomes suspended. A delivery can become modified at the source even before all transfer frames have been sent. This does not imply that the remaining transfers for the delivery will not be sent. The source MAY spontaneously attain the modified outcome for a message (for example the source might implement some sort of time-bound acquisition lock, after which the acquisition of a message at a node is revoked to allow for delivery to an alternative consumer with the message modified in some way to denote the previous failed, e.g., with `delivery-failed` set to true).

At the target, the modified outcome is used to indicate that a given transfer was not and will not be acted upon, and that the message SHOULD be modified in the specified ways at the node.

Field Details

`delivery-failed` *count the transfer as an unsuccessful delivery attempt*

If the `delivery-failed` flag is set, any messages modified MUST have their `delivery-count` incremented.

`undeliverable-here` *prevent redelivery*

If the `undeliverable-here` is set, then any messages released MUST NOT be redelivered to the modifying link endpoint.

`message-annotations` *message attributes*

Map containing attributes to combine with the existing *message-annotations* held in the message's header section. Where the existing *message-annotations* of the message contain an entry with the same key as an entry in this field, the value in this field associated with that key replaces the one in the existing headers; where the existing *message-annotations* has no such value, the value in this map is added.

3.4.6 Resuming Deliveries Using Delivery States

Part 2: 2.6.13 Resuming Deliveries provides the general scheme for how two endpoints can re-establish state after link resumption was provided. The concrete delivery states defined above allow for a more comprehensive set of examples of link resumption.

```

Peer                                     Partner
=====
ATTACH(name=N, handle=1,                --+   +-- ATTACH(name=N, handle=2,
role=sender,                             \   /    role=receiver,
source=X,                                 x   x    source=X,
target=Y,                                  /   \    target=Y,
unsettled=                                <--+ +-- unsettled=
{ 1 -> null,                               { 2 -> Received(3,0),
  2 -> null,                               3 -> Accepted,
  3 -> null,                               4 -> null,
  4 -> null,                               6 -> Received(2,0),
  5 -> Received(0,200),                    7 -> Received(0,100),
  6 -> Received(1,150),                    8 -> Accepted,
  7 -> Received(0,500),                    9 -> null,
  8 -> Received(3,5),                      11 -> Received(1,2000),
  9 -> Received(2,0),                      12 -> Accepted,
  10 -> Accepted,                          13 -> Released,
  11 -> Accepted,                          14 -> null }
  12 -> Accepted,
  13 -> Accepted,
  14 -> Accepted }
-----
Key:
Received(x,y) means Received(section-number=x, section-offset=y)

```

In this example, for delivery-tags 1 to 4 inclusive the sender indicates that it can resume sending from the start of the message.

For delivery-tag 1, the receiver has no record of the delivery. To preserve “exactly once” or “at least once” delivery guarantees, the sender **MUST** resend the message; however, the delivery is not being resumed (since the receiver does not remember the delivery tag) so transfers **MUST NOT** have the resume flag set to true. If the sender were to mark the transfers as resumes then they would be ignored at the receiver.

For delivery-tag 2, the receiver has retained some of the data making up the message, but not the whole. In order to complete the delivery the sender needs to resume sending from some point before or at the next position which the receiver is expecting.

```

TRANSFER(delivery-id=1, -----> ** Append message data not **
  delivery-tag=2,                ** seen previously to delivery **
  (1) state=Received(3,0),       ** state. **
      resume=true)
{ ** payload ** }

(1) state could be
a) null, meaning that the transfer is being resumed from the first
   byte of section number 0 (and the receiver MUST ignore all data
   up to the first position it has not previously received).
b) Received with section number 0, 1 or 2 and an offset, indicating
   that the payload data on the first frame of the resumed delivery
   starts at the given point, and that the receiver MUST ignore all
   data up to the first position it has not previously received.
c) Received(3,0) indicating that the resumption will start at the
   first point which the receiver has not previously received.

```

For delivery-tag 3, the receiver indicates that it has processed the delivery to the point where it desires a terminal outcome (in this case accepted). In this case the sender will either apply that outcome at the source, or in the rare case that it cannot apply that outcome, indicate the terminal outcome that has been applied. To do this the sender **MUST** send a resuming transfer to associate delivery-tag 3 with a delivery-id. On this transfer the sender **SHOULD** set the delivery-state at the source. If this is the same outcome as at the receiver then the sender **MAY** also send the resuming transfer as settled.

```
TRANSFER(delivery-id=2, -----> ** Processes confirmation that **
          delivery-tag=3,          ** was accepted, and settles. **
          settled=true,
          more=false,
          state=Accepted,
          resume=true)
```

For delivery-tag 4, the receiver indicates that it is aware that the delivery had begun, but does not provide any indication that it has retained any data about that delivery except the fact of its existence. To preserve “exactly once” or “at least once” delivery guarantees, the sender **MUST** resend the whole message. Unlike the case with delivery-tag 1 the resent delivery **MUST** be sent with the resume flag set to true and the delivery-tag set to 4. (While this use of null in the receivers map is valid, it is discouraged. It is **RECOMMENDED** that receiver **SHOULD NOT** retain such an entry in its map, in which case the situation would be as for delivery-tag 1 in this example).

```
TRANSFER(delivery-id=3, -----> ** Processes in the same way **
          delivery-tag=4,          ** as we be done for a non- **
          (1) state=null,          ** resumed delivery.      **
          resume=true)
{ ** payload ** }
```

(1) Alternatively (and equivalently) state could be
 Received(section-number=0, section-offset=0)

For delivery-tags 5 to 9 inclusive the sender indicates that it can resume at some point beyond start of the message data. This is usually indicative of the fact that the receiver had previously confirmed reception of message data to the given point, removing responsibility from the sender to retain the ability to resend that data after resuming the link. The sender **MAY** still retain the ability to resend the message as a new delivery (i.e. it **MAY** not have completely discarded the data from which the original delivery was generated).

For delivery-tag 5, the receiver has no record of the delivery. To preserve “exactly once” or “at least once” delivery guarantees, the sender **MUST** resend the message; however, the delivery is not being resumed (since the receiver does not remember the delivery tag) so transfers **MUST NOT** have the resume flag set to true. If the sender does not have enough data to resend the message, then the sender **MAY** take some action to indicate that it believes there is a possibility that there has been message loss, for example, notify the application.

For delivery-tag 6, the receiver has retained some of the data making up the message, but not the whole. The first position within the message which the receiver has not received is after the first position at which the sender can resume sending. In order to complete the delivery the sender needs to resume sending from some point before or at the next position which the receiver is expecting.

```
TRANSFER(delivery-id=4, -----> ** Append message data not **
          delivery-tag=6,          ** seen previously to delivery **
          (1) state=Received(2,0), ** state.                      **
          resume=true)
{ ** payload ** }
```

(1) state could be any point between Received(1,150) and Received(2,0) inclusive. The receiver **MUST** ignore all data up to the first position it has not previously received (i.e. section 2 offset 0).

For delivery-tag 7, the receiver has retained some of the data making up the message, but not the whole. The first position within the message which the receiver has not received is before the first position at which the sender can resume sending. It is thus not possible for the sender to resume sending the message to completion. The only option available to the sender is to abort the transfer and to then (if possible) resend as a new delivery or else to report the possible message loss in some way if it cannot.

```
TRANSFER(delivery-id=5, -----> ** Discard any state relating **
          delivery-tag=7,          ** to the message delivery.   **
          resume=true,
          aborted=true)
```

For delivery-tag 8, the receiver indicates that it has processed the delivery to the point where it desires a terminal outcome (in this case `accepted`). This is the same case as for delivery-tag 3.

```
TRANSFER(delivery-id=6, -----> ** Processes confirmation that **
delivery-tag=8,          ** was accepted, and settles. **
settled=true,
more=false,
state=Accepted,
resume=true)
```

For delivery-tag 9, the receiver indicates that it is aware that the delivery had begun, but does not provide any indication that it has retained any data about that delivery except the fact of its existence. This is the same case as for delivery-tag 7.

```
TRANSFER(delivery-id=7, -----> ** Discard any state relating **
delivery-tag=9,          ** to the message delivery.   **
resume=true,
aborted=true)
```

For delivery-tags 10 to 14 inclusive the sender indicates that it has reached a terminal outcome, namely `accepted`. Once the sender has arrived at a terminal outcome it **MUST NOT** change. As such, if a sender is capable of resuming a delivery (even if the only possible outcome of the delivery is a pre-defined terminal outcome such as `accepted`) it **MUST NOT** use this state as the value of the state in its unsettled map until it is sure that the receiver will not require the resending of the message data.

For delivery-tag 10 the receiver has no record of the delivery. However, in contrast to the cases of delivery-tag 1 and delivery-tag 5, since it is known that the sender can only have arrived at this state through knowing that the receiver has received the whole message (or that the sender had spontaneously reached a terminal outcome with no possibility of resumption) there is no need to resend the message.

For delivery-tag 11 it **MUST** be assumed that the sender spontaneously attained the terminal outcome (and is unable to resume). In this case the sender can simply abort the delivery as it cannot be resumed.

```
TRANSFER(delivery-id=8, -----> ** Discard any state relating **
delivery-tag=11,         ** to the message delivery.   **
resume=true,
aborted=true)
```

For delivery-tag 12 both the sender and receiver have attained the same view of the terminal outcome, but neither has settled. In this case the sender **SHOULD** simply settle the delivery.

```
TRANSFER(delivery-id=9, -----> ** Locally settle the delivery **
delivery-tag=12,
settled=true,
resume=true)
```

For delivery-tag 13 the sender and receiver have both attained terminal outcomes, but the outcomes differ. In this case, since the outcome actually takes effect at the sender, it is the sender's view that is definitive. The sender thus **MUST** restate this as the terminal outcome, and the receiver **SHOULD** then echo this and settle.

```
TRANSFER(delivery-id=10 -----> ** Update any state affected **
delivery-tag=13,          ** by the actual outcome, then **
settled=false,           ** settle the delivery         **
state=Accepted,
resume=true)
<----- DISPOSITION(first=10, last=10,
state=Accepted,
settled=true)
```

For delivery-tag 14 the case is essentially the same as for delivery-tag 11, as the null state at the receiver is essentially identical to having the state `Received` with `section-number=0` and `section-offset=0`.

```
TRANSFER(delivery-id=11, -----> ** Discard any state relating **
          delivery-tag=14,                ** to the message delivery.   **
          resume=true,
          aborted=true)
```

3.5 Sources and Targets

The messaging layer defines two concrete types (*source* and *target*) to be used as the *source* and *target* of a link. These types are supplied in the *source* and *target* fields of the *attach* frame when establishing or resuming link. The *source* is comprised of an address (which the container of the outgoing link endpoint will resolve to a node within that container) coupled with properties which determine:

- which messages from the sending node will be sent on the link,
- how sending the message affects the state of that message at the sending node,
- the behavior of messages which have been transferred on the link, but have not yet reached a terminal state at the receiver, when the source is destroyed.

3.5.1 Filtering Messages

A source can restrict the messages transferred from a source by specifying a filter. A filter can be thought of as a function which takes a message as input and returns a boolean value: true if the message will be accepted by the source, false otherwise. A *filter* MUST NOT change its return value for a message unless the state or annotations on the message at the node change (e.g., through an updated delivery state).

3.5.2 Distribution Modes

The source optionally defines a distribution-mode that informs and/or indicates how distribution nodes are to behave with respect to the link. The distribution-mode of a source determines how messages from a distribution node are distributed among its associated links. There are two defined distribution-modes: *move* and *copy*. When specified, the distribution-mode has two related effects on the behavior of a distribution node with respect to the link associated with the source.

The *move* distribution-mode causes messages transferred from the distribution node to transition to the ACQUIRED state prior to transfer over the link, and subsequently to the ARCHIVED state when the transfer is settled with a successful outcome. The *copy* distribution-mode leaves the state of the message unchanged at the distribution node.

A source MUST NOT resend a message which has previously been successfully transferred from the source, i.e., reached an ACCEPTED delivery state at the receiver. For a *move* link with a default configuration this is trivially achieved as such an end result will lead to the message in the ARCHIVED state on the node, and thus ineligible for transfer. For a *copy* link, state MUST be retained at the source to ensure compliance. In practice, for nodes which maintain a strict order on messages at the node, the state might simply be a record of the most recent message transferred.

A registry of commonly defined non-standard distribution-modes and their meanings is maintained [AMQPDIST-MODE].

3.5.3 Source

```
<type name="source" class="composite" source="list" provides="source">
  <descriptor name="amqp:source:list" code="0x00000000:0x00000028"/>
  <field name="address" type="*" requires="address"/>
  <field name="durable" type="terminus-durability" default="none"/>
  <field name="expiry-policy" type="terminus-expiry-policy" default="session-end"/>
  <field name="timeout" type="seconds" default="0"/>
  <field name="dynamic" type="boolean" default="false"/>
  <field name="dynamic-node-properties" type="node-properties"/>
  <field name="distribution-mode" type="symbol" requires="distribution-mode"/>
  <field name="filter" type="filter-set"/>
  <field name="default-outcome" type="*" requires="outcome"/>
  <field name="outcomes" type="symbol" multiple="true"/>
  <field name="capabilities" type="symbol" multiple="true"/>
</type>
```

For containers which do not implement address resolution (and do not admit spontaneous link attachment from their partners) but are instead only used as producers of messages, it is unnecessary to provide spurious detail on the source. For this purpose it is possible to use a “minimal” source in which all the fields are left unset.

Field Details

`address` *the address of the source*

The address of the source MUST NOT be set when sent on a `attach` frame sent by the receiving link endpoint where the dynamic flag is set to true (that is where the receiver is requesting the sender to create an addressable node).

The address of the source MUST be set when sent on a `attach` frame sent by the sending link endpoint where the dynamic flag is set to true (that is where the sender has created an addressable node at the request of the receiver and is now communicating the address of that created node). The generated name of the address SHOULD include the link name and the container-id of the remote container to allow for ease of identification.

`durable` *indicates the durability of the terminus*

Indicates what state of the terminus will be retained durably: the state of durable messages, only existence and configuration of the terminus, or no state at all.

`expiry-policy` *the expiry policy of the source*

See subsection 3.5.6 Terminus Expiry Policy.

`timeout` *duration that an expiring source will be retained*

The source starts expiring as indicated by the expiry-policy.

`dynamic` *request dynamic creation of a remote node*

When set to true by the receiving link endpoint, this field constitutes a request for the sending peer to dynamically create a node at the source. In this case the address field MUST NOT be set. When set to true by the sending link endpoint this field indicates creation of a dynamically created node. In this case the address field will contain the address of the created node. The generated address SHOULD include the link name and other available information on the initiator of the request (such as the remote container-id) in some recognizable form for ease of traceability.

`dynamic-node-properties` *properties of the dynamically created node*

If the dynamic field is not set to true this field **MUST** be left unset.

When set by the receiving link endpoint, this field contains the desired properties of the node the receiver wishes to be created. When set by the sending link endpoint this field contains the actual properties of the dynamically created node. See subsection 3.5.9 Node Properties for standard node properties. A registry of other commonly used node-properties and their meanings is maintained [AMQPNODEPROP].

`distribution-mode` *the distribution mode of the link*

This field **MUST** be set by the sending end of the link if the endpoint supports more than one distribution-mode. This field **MAY** be set by the receiving end of the link to indicate a preference when a node supports multiple distribution modes.

`filter` *a set of predicates to filter the messages admitted onto the link*

See subsection 3.5.8 Filter Set. The receiving endpoint sets its desired filter, the sending endpoint sets the filter actually in place (including any filters defaulted at the node). The receiving endpoint **MUST** check that the filter in place meets its needs and take responsibility for detaching if it does not.

`default-outcome` *default outcome for unsettled transfers*

Indicates the outcome to be used for transfers that have not reached a terminal state at the receiver when the transfer is settled, including when the source is destroyed. The value **MUST** be a valid outcome (e.g., `released` or `rejected`).

`outcomes` *descriptors for the outcomes that can be chosen on this link*

The values in this field are the symbolic descriptors of the outcomes that can be chosen on this link. This field **MAY** be empty, indicating that the *default-outcome* will be assumed for all message transfers (if the default-outcome is not set, and no outcomes are provided, then the `accepted` outcome **MUST** be supported by the source).

When present, the values **MUST** be a symbolic descriptor of a valid outcome, e.g., `"amqp:accepted:list"`.

`capabilities` *the extension capabilities the sender supports/desires*

A registry of commonly defined source capabilities and their meanings is maintained [AMQP-SOURCECAP].

3.5.4 Target

```
<type name="target" class="composite" source="list" provides="target">
  <descriptor name="amqp:target:list" code="0x00000000:0x00000029"/>
  <field name="address" type="*" requires="address"/>
  <field name="durable" type="terminus-durability" default="none"/>
  <field name="expiry-policy" type="terminus-expiry-policy" default="session-end"/>
  <field name="timeout" type="seconds" default="0"/>
  <field name="dynamic" type="boolean" default="false"/>
  <field name="dynamic-node-properties" type="node-properties"/>
  <field name="capabilities" type="symbol" multiple="true"/>
</type>
```

For containers which do not implement address resolution (and do not admit spontaneous link attachment from their partners) but are instead only used as consumers of messages, it is unnecessary to provide spurious detail on the source. For this purpose it is possible to use a “minimal” target in which all the fields are left unset.

Field Details

address	<i>The address of the target.</i>
	The address of the target MUST NOT be set when sent on a <code>attach</code> frame sent by the sending link endpoint where the dynamic flag is set to true (that is where the sender is requesting the receiver to create an addressable node). The address of the source MUST be set when sent on a <code>attach</code> frame sent by the receiving link endpoint where the dynamic flag is set to true (that is where the receiver has created an addressable node at the request of the sender and is now communicating the address of that created node). The generated name of the address SHOULD include the link name and the container-id of the remote container to allow for ease of identification.
durable	<i>indicates the durability of the terminus</i>
	Indicates what state of the terminus will be retained durably: the state of durable messages, only existence and configuration of the terminus, or no state at all.
expiry-policy	<i>the expiry policy of the target</i>
	See subsection 3.5.6 Terminus Expiry Policy.
timeout	<i>duration that an expiring target will be retained</i>
	The target starts expiring as indicated by the expiry-policy.
dynamic	<i>request dynamic creation of a remote node</i>
	When set to true by the sending link endpoint, this field constitutes a request for the receiving peer to dynamically create a node at the target. In this case the address field MUST NOT be set. When set to true by the receiving link endpoint this field indicates creation of a dynamically created node. In this case the address field will contain the address of the created node. The generated address SHOULD include the link name and other available information on the initiator of the request (such as the remote container-id) in some recognizable form for ease of traceability.
dynamic-node-properties	<i>properties of the dynamically created node</i>
	If the dynamic field is not set to true this field MUST be left unset. When set by the sending link endpoint, this field contains the desired properties of the node the sender wishes to be created. When set by the receiving link endpoint this field contains the actual properties of the dynamically created node. See subsection 3.5.9 Node Properties for standard node properties. A registry of other commonly used node-properties and their meanings is maintained [AMQPNODEPROP].
capabilities	<i>the extension capabilities the sender supports/desires</i>
	A registry of commonly defined target capabilities and their meanings is maintained [AMQPTARGETCAP].

3.5.5 Terminus Durability

Durability policy for a terminus.

```
<type name="terminus-durability" class="restricted" source="uint">
  <choice name="none" value="0"/>
  <choice name="configuration" value="1"/>
  <choice name="unsettled-state" value="2"/>
</type>
```

Determines which state of the terminus is held durably.

Valid Values

- 0 No terminus state is retained durably.
- 1 Only the existence and configuration of the terminus is retained durably.
- 2 In addition to the existence and configuration of the terminus, the unsettled state for durable messages is retained durably.

3.5.6 Terminus Expiry Policy

Expiry policy for a terminus.

```
<type name="terminus-expiry-policy" class="restricted" source="symbol">
  <choice name="link-detach" value="link-detach"/>
  <choice name="session-end" value="session-end"/>
  <choice name="connection-close" value="connection-close"/>
  <choice name="never" value="never"/>
</type>
```

Determines when the expiry timer of a terminus starts counting down from the timeout value. If the link is subsequently re-attached before the terminus is expired, then the count down is aborted. If the conditions for the terminus-expiry-policy are subsequently re-met, the expiry timer restarts from its originally configured timeout value.

Valid Values

- link-detach** The expiry timer starts when terminus is detached.
- session-end** The expiry timer starts when the most recently associated session is ended.
- connection-close** The expiry timer starts when most recently associated connection is closed.
- never** The terminus never expires.

3.5.7 Standard Distribution Mode

Link distribution policy.

```
<type name="std-dist-mode" class="restricted" source="symbol" provides="distribution-mode">
  <choice name="move" value="move"/>
  <choice name="copy" value="copy"/>
</type>
```

Policies for distributing messages when multiple links are connected to the same node.

Valid Values

- move** once successfully transferred over the link, the message will no longer be available to other links from the same node
- copy** once successfully transferred over the link, the message is still available for other links from the same node

3.5.8 Filter Set

```
<type name="filter-set" class="restricted" source="map"/>
```

A set of named filters. Every key in the map **MUST** be of type `symbol`, every value **MUST** be either `null` or of a described type which provides the archetype *filter*. A filter acts as a function on a message which returns a boolean result indicating whether the message can pass through that filter or not. A message will pass through a filter-set if and only if it passes through each of the named filters. If the value for a given key is `null`, this acts as if there were no such key present (i.e., all messages pass through the null filter).

Filter types are a defined extension point. The filter types that a given `source` supports will be indicated by the capabilities of the `source`. A registry of commonly defined filter types and their capabilities is maintained [AMQPFILTERS].

3.5.9 Node Properties

Properties of a node.

```
<type name="node-properties" class="restricted" source="fields"/>
```

A symbol-keyed map containing properties of a node used when requesting creation or reporting the creation of a dynamic node.

The following common properties are defined:

lifetime-policy The lifetime of a dynamically generated node.

Definitionally, the lifetime will never be less than the lifetime of the link which caused its creation, however it is possible to extend the lifetime of dynamically created node using a lifetime policy. The value of this entry **MUST** be of a type which provides the *lifetime-policy* archetype. The following standard *lifetime-policies* are defined below: `delete-on-close`, `delete-on-no-links`, `delete-on-no-messages` Or `delete-on-no-links-or-messages`.

supported-dist-modes

The distribution modes that the node supports.

The value of this entry **MUST** be one or more symbols which are valid *distribution-modes*. That is, the value **MUST** be of the same type as would be valid in a field defined with the following attributes:

```
type="symbol" multiple="true" requires="distribution-mode"
```

3.5.10 Delete On Close

Lifetime of dynamic node scoped to lifetime of link which caused creation.

```
<type name="delete-on-close" class="composite" source="list" provides="lifetime-policy">
  <descriptor name="amqp:delete-on-close:list" code="0x00000000:0x0000002b"/>
</type>
```

A node dynamically created with this lifetime policy will be deleted at the point that the link which caused its creation ceases to exist.

3.5.11 Delete On No Links

Lifetime of dynamic node scoped to existence of links to the node.

```
<type name="delete-on-no-links" class="composite" source="list" provides="lifetime-policy">
  <descriptor name="amqp:delete-on-no-links:list" code="0x00000000:0x0000002c"/>
</type>
```

A node dynamically created with this lifetime policy will be deleted at the point that there remain no links for which the node is either the source or target.

3.5.12 Delete On No Messages

Lifetime of dynamic node scoped to existence of messages on the node.

```
<type name="delete-on-no-messages" class="composite" source="list" provides="lifetime-policy">
  <descriptor name="amqp:delete-on-no-messages:list" code="0x00000000:0x0000002d"/>
</type>
```

A node dynamically created with this lifetime policy will be deleted at the point that the link which caused its creation no longer exists and there remain no messages at the node.

3.5.13 Delete On No Links Or Messages

Lifetime of node scoped to existence of messages on or links to the node.

```
<type name="delete-on-no-links-or-messages" class="composite" source="list" provides="lifetime-policy">
  <descriptor name="amqp:delete-on-no-links-or-messages:list" code="0x00000000:0x0000002e"/>
</type>
```

A node dynamically created with this lifetime policy will be deleted at the point that there are no links which have this node as their source or target, and there remain no messages at the node.

Part 4: Transactions

4.1 Transactional Messaging

Transactional messaging allows for the coordinated outcome of otherwise independent transfers. This extends to an arbitrary number of transfers spread across any number of distinct links in either direction.

For every transactional interaction, one container acts as the *transactional resource*, and the other container acts as the *transaction controller*. The *transactional resource* performs *transactional work* as requested by the *transaction controller*.

The *transactional controller* and *transactional resource* communicate over a *control link* which is established by the *transactional controller*. The `declare` and `dispatch` messages are sent by the *transactional controller* over the *control link* to allocate and complete transactions respectively (they do not represent the demarcation of transactional work). No transactional work is allowed on the *control link*. Each transactional operation requested is explicitly identified with the desired transaction-id and therefore can occur on any link within the controlling session, or, if permitted by the capabilities of the controller, any link on the controlling connection. If the *control link* is closed while there exist non-discharged transactions it created, then all such transactions are immediately rolled back, and attempts to perform further transactional work on them will lead to failure.

4.2 Declaring a Transaction

The container acting as the transactional resource defines a special target that functions as a coordinator. The *transaction controller* establishes a control link to this target. Note that links to the coordinator cannot be resumed.

To begin transactional work, the transaction controller needs to obtain a transaction identifier from the resource. It does this by sending a message to the coordinator whose body consists of the `declare` type in a single `amqp-value` section. Other standard message sections such as the header section SHOULD be ignored. This message MUST NOT be sent settled as the sender is REQUIRED to receive and interpret the outcome of the `declare` from the receiver. If the coordinator receives a `transfer` that has been settled by the sender, it SHOULD `detach` with an `amqp:illegal-state` error.

If the declaration is successful, the coordinator responds with a disposition outcome of `declared` which carries the assigned identifier for the transaction.

If the coordinator was unable to perform the `declare` as specified by the transaction controller, the transaction coordinator MUST convey the error to the controller as a `transaction-error`. If the source for the link to the coordinator supports the `rejected` outcome, then the message MUST be `rejected` with this outcome carrying the `transaction-error`. If the source does not support the `rejected` outcome, the *transactional resource* MUST `detach` the link to the coordinator, with the `detach` performative carrying the `transaction-error`.

Transaction controllers SHOULD establish a control link that allows the `rejected` outcome.

```

Transaction Controller          Transactional Resource
=====
ATTACH(name=txn-ctl,          ----->
      ...,
      target=
      Coordinator(
      capabilities=
      "amqp:local-transactions")
    )

      <----- ATTACH(name=txn-ctl,
      ...,
      target=
      Coordinator(
      capabilities=
      ["amqp:local-transactions",
      "amqp:multi-txns-per-ssn"]
      )
    )

      <----- FLOW(...,handle=1, link-credit=1)

TRANSFER(delivery-id=0, ...) ----->
{ AmqpValue( Declare() ) }

      <----- DISPOSITION(first=0, last=0,
      state=Declared(txn-id=0) )
=====

```

Figure 4.1: Declaring a Transaction

4.3 Discharging a Transaction

The controller will conclude the transactional work by sending a discharge message (encoded in a single `amqp-value` section) to the coordinator. The controller indicates that it wishes to commit or rollback the transactional work by setting the *fail* flag on the discharge body. As with the `declare` message, it is an error if the sender sends the transfer pre-settled.

If the coordinator is unable to complete the discharge, the coordinator **MUST** convey the error to the controller as a `transaction-error`. If the source for the link to the coordinator supports the `rejected` outcome, then the message **MUST** be rejected with this outcome carrying the `transaction-error`. If the source does not support the `rejected` outcome, the *transactional resource* **MUST** detach the link to the coordinator, with the `detach` performative carrying the `transaction-error`. Note that the coordinator **MUST** always be able to complete a discharge where the `fail` flag is set to true (since coordinator failure leads to rollback, which is what the controller is asking for).

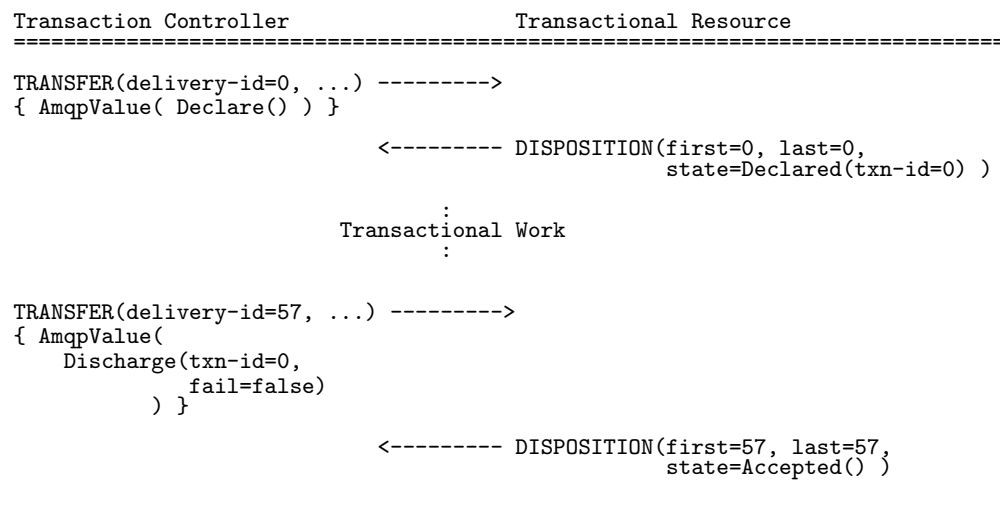


Figure 4.2: Discharging a Transaction

4.4 Transactional Work

Transactional work is described in terms of the message states defined in Part 3: section 3.3. Transactional work is formally defined to be composed of the following operations:

- *posting* a message at a target, i.e., making it *available*
- *acquiring* a message at a source, i.e., transitioning it to *acquired*
- *retiring* a message at a source, i.e., applying the terminal outcome

The transactional resource performs these operations when triggered by the transaction controller:

- *posting* messages is initiated by incoming `transfer` frames
- *acquiring* messages is initiated by incoming `flow` frames
- *retiring* messages is initiated by incoming `disposition` frames

In each case, it is the responsibility of the transaction controller to identify the transaction with which the requested work is to be associated. This is done with the transactional delivery state `transactional-state` that combines a `txn-id` together with one of the terminal delivery states defined in Part 3: section 3.4. The `transactional-state` is carried by both the `transfer` and the `disposition` frames allowing both the *posting* and *retiring* of messages to be associated with a transaction.

The `transfer`, `disposition`, and `flow` frames can travel in either direction, i.e., from the controller to the resource and from the resource to the controller. When these frames travel from the controller to the resource, any embedded `txn-ids` are requesting that the resource assigns transactional work to the indicated transaction. When traveling in the other direction, from resource to controller, the `transfer` and `disposition` frames indicate work performed, and the `txn-ids` included **MUST** correctly indicate with which (if any) transaction this work is associated. In the case of the `flow` frame traveling from resource to controller, the `txn-id` does not indicate work that has been performed, but indicates with which transaction future transfers from that link will be performed.

4.4.1 Transactional Posting

If the transaction controller wishes to associate an outgoing `transfer` with a transaction, it **MUST** set the state of the transfer with a `transactional-state` carrying the appropriate transaction identifier. Note that if delivery is

split across several `transfer` frames then all frames MUST be explicitly associated with the same transaction. It is an error for the controller to attempt to discharge a transaction against which a partial delivery has been posted. If this happens, the control link MUST be terminated with the `transaction-rollback` error.

The effect of transactional posting is that the message does not become available at the destination node within the transactional resource until after the transaction has been successfully discharged.

```

Transaction Controller          Transactional Resource
=====
TRANSFER(handle=0,             ----->
          delivery-id=0, ...)
{ AmqpValue( Declare() ) }

                                <----- DISPOSITION(first=0, last=0,
                                                state=Declared(txn-id=0) )

TRANSFER(handle=1,             ----->
          delivery-id=1,
          state=
            TransactionalState(
              txn-id=0) )
{ ... payload ... }

                                <----- DISPOSITION(first=1, last=1,
                                                state=TransactionalState(
                                                    txn-id=0,
                                                    outcome=Accepted()
                                                )
)
-----

```

Figure 4.3: Transactional Publish

On receiving a non-settled delivery associated with a live transaction, the transactional resource MUST inform the controller of the presumptive terminal outcome before it can successfully discharge the transaction. That is, the resource MUST send a `disposition` performative which covers the posted transfer with the state of the delivery being a `transactional-state` with the correct transaction identified, and a terminal outcome. This informs the controller of the outcome that will be in effect at the point that the transaction is successfully discharged.

4.4.2 Transactional Retirement

The transaction controller might wish to associate the outcome of a delivery with a transaction. The delivery itself need not be associated with the same transaction as the outcome, or indeed with any transaction at all. However, the delivery MUST NOT be associated with a different non-discharged transaction than the outcome. If this happens then the control link MUST be terminated with a `transaction-rollback` error.

To associate an outcome with a transaction the controller sends a `disposition` performative which sets the state of the delivery to a `transactional-state` with the desired transaction identifier and the outcome to be applied upon a successful discharge.

```

Transaction Controller                               Transactional Resource
=====
TRANSFER(handle=0,                                ----->
          delivery-id=0, ...)
{ AmqpValue( Declare() ) }

<----- DISPOSITION(first=0, last=0,
                    state=Declared(txn-id=0) )

FLOW(handle=2,                                     ----->
     link-credit=10)

<----- TRANSFER(handle=2,
                  delivery-id=11,
                  state=null,
                  { ... payload ... }

          :
          :

<----- TRANSFER(handle=2,
                  delivery-id=20,
                  state=null,
                  { ... payload ... }

DISPOSITION(first=11,                               ----->
            last=20,
            state=TransactionalState(
              txn-id=0,
              outcome=Accepted()
            )
-----

```

Figure 4.4: Transactional Receive

On a successful *discharge*, the resource will apply the given outcome and can immediately settle the transfers. In the event of a controller-initiated rollback (a *discharge* where the fail flag is set to true) or a resource-initiated rollback (the *discharge* message being rejected, or the link to the coordinator being detached with an error), the outcome will not be applied, and the deliveries will still be “live” and will remain acquired by the controller, i.e., the resource can expect the controller to request a disposition for the delivery (either transactionally on a new transaction, or non-transactionally).

4.4.3 Transactional Acquisition

In the case of the *flow* frame, the transactional work is not necessarily directly initiated or entirely determined when the *flow* frame arrives at the resource, but can in fact occur at some later point and in ways not necessarily anticipated by the controller. To accommodate this, the resource associates an additional piece of state with outgoing link endpoints, a *txn-id* that identifies the transaction with which *acquired* messages will be associated. This state is determined by the controller by specifying a *txn-id* entry in the *properties* map of the flow frame. When a transaction is discharged, the *txn-id* of any link endpoints will be cleared.

If the link endpoint does not support transactional acquisition, the link **MUST** be terminated with a *not-implemented* error.

While the *txn-id* is cleared when the transaction is discharged, this does not affect the level of outstanding credit. To prevent the sending link endpoint from acquiring outside of any transaction, the *controller* **SHOULD** ensure there is no outstanding credit at the sender before it discharges the transaction. The *controller* can do this by either setting the drain mode of the sending link endpoint to *true* before discharging the transaction, or by reducing the *link-credit* to zero, and waiting to hear back that the sender has seen this state change.

If a transaction is discharged at a point where a message has been transactionally acquired, but has not been fully sent (i.e., the delivery of the message will require more than one *transfer* frame and at least one, but not all, such frames have been sent), then the resource **MUST** interpret this to mean that the fate of the acquisition is

fully decided by the discharge. If the `discharge` indicates the failure of the transaction the resource **MUST** abort or complete sending the remainder of the message before completing the discharge.

```

Transaction Controller          Transactional Resource
=====
TRANSFER(handle=0,             ----->
          delivery-id=0, ...)
{ AmqpValue( Declare() ) }

                                <----- DISPOSITION(first=0, last=0,
                                                state=Declared(txn-id=0) )

FLOW(handle=2,                  ----->
     link-credit=10,
     drain=true,
     properties={
       txn-id=0
     })

                                <----- TRANSFER(handle=2,
                                                delivery-id=11,
                                                state=
                                                  TransactionalState(txn-id=0),
                                                { ... payload ... }

                                :
                                :

                                <----- TRANSFER(handle=2,
                                                delivery-id=20,
                                                state=
                                                  TransactionalState(txn-id=0),
                                                { ... payload ... }

DISPOSITION(first=11,          ----->
            last=20,
            state=TransactionalState(
              txn-id=0,
              outcome=Accepted())
          )
-----

```

Figure 4.5: Transactional Acquisition

4.4.4 Interaction Of Settlement With Transactions

The transport layer defines a notion of *settlement* which refers to the point at which the association of a delivery-tag to a delivery attempt is forgotten. Transactions do not in themselves change this notion, however the fact that transactional work can be rolled back does have implications for deliveries which the endpoint has marked as settled (and for which it can therefore no longer exchange state information using the previously allocated transport level identifiers).

4.4.4.1 Transactional Posting

Delivery Sent Settled By Controller

The delivered message will not be made available at the node until the transaction has been successfully discharged. If the transaction is rolled back then the delivery is not made available. If the resource is unable to process the delivery it **MUST NOT** allow the successful discharge of the associated transaction. This can be communicated by immediately destroying the controlling link on which the transaction was declared, or by rejecting any attempt to discharge the transaction where the fail flag is not set to true.

Delivery Sent Unsettled By Controller; Resource Settles

The resource **MUST** determine the outcome of the delivery before committing the transaction, and this **MUST** be communicated to the controller before the acceptance of a successful discharge. The outcome

communicated by the resource **MUST** be associated with the same transaction with which the `transfer` from controller to resource was associated.

If the transaction is rolled back then the delivery is not made available at the target. If the resource can no longer apply the outcome that it originally indicated would be the result of a successful discharge then it **MUST NOT** allow the successful discharge of the associated transaction. This can be communicated by immediately destroying the controlling link on which the transaction was declared, or by rejecting any attempt to discharge the transaction where the fail flag is not set to true.

Delivery Sent Unsettled By Controller; Resource Does Not Settle

Behavior prior to discharge is the same as the previous case.

After a successful discharge, the state of unsettled deliveries at the resource **MUST** reflect the outcome that was applied. If the discharge was unsuccessful then an outcome **SHOULD NOT** be associated with the unsettled deliveries. The controller **SHOULD** settle any outstanding unsettled deliveries in a timely fashion after the transaction has discharged.

4.4.4.2 Transactional Retirement

This section considers the cases where the resource has sent messages to the controller in a non-transactional fashion. For the cases where the resource sends the messages transactionally, see subsection 4.4.4.3.

Delivery Sent Unsettled By Resource; Controller Settles

Upon a successful discharge the outcome specified by the controller is applied at the source. If the controller requests a rollback or the discharge attempt be unsuccessful, then the outcome is not applied. At this point the controller can no longer influence the state of the delivery as it is settled, and the resource **MUST** apply the default outcome.

Delivery Sent Unsettled By Resource; Controller Does Not Settle

The resource might or might not settle the delivery before the transaction is discharged. If the resource settles the delivery before the discharge then the behavior after discharge is the same as the case above.

Upon a successful discharge the outcome is applied. Otherwise the state reverts to that which occurred before the controller sent its (transactional) disposition. The controller is free to update the state using subsequent transactional or non-transactional updates.

4.4.4.3 Transactional Acquisition

Delivery Sent Settled By Resource

In the event of a successful discharge the outcome applies at the resource, otherwise the acquisition and outcome are rolled back.

Delivery Sent Unsettled By Resource; Controller Sends Outcome

An outcome sent by the controller before the transaction has discharged **MUST** be associated with the same transaction. In the even of a successful discharge the outcome is applied at the source, otherwise both the acquisition and outcome are rolled back.

4.5 Coordination

4.5.1 Coordinator

Target for communicating with a transaction coordinator.

```
<type name="coordinator" class="composite" source="list" provides="target">
  <descriptor name="amqp:coordinator:list" code="0x00000000:0x00000030"/>
  <field name="capabilities" type="symbol" requires="txn-capability" multiple="true"/>
</type>
```

The coordinator type defines a special target used for establishing a link with a transaction coordinator.

Field Details

`capabilities` *the capabilities supported at the coordinator*

When sent by the transaction controller (the sending endpoint), indicates the desired capabilities of the coordinator. When sent by the resource (the receiving endpoint), defined the actual capabilities of the coordinator. Note that it is the responsibility of the transaction controller to verify that the capabilities of the controller meet its requirements. See `txn-capability`.

4.5.2 Declare

Message body for declaring a transaction id.

```
<type name="declare" class="composite" source="list">
  <descriptor name="amqp:declare:list" code="0x00000000:0x00000031"/>
  <field name="global-id" type="*" requires="global-tx-id"/>
</type>
```

The declare type defines the message body sent to the coordinator to declare a transaction. The `txn-id` allocated for this transaction is chosen by the transaction controller and identified in the `declared` resulting outcome.

Field Details

`global-id` *global transaction id*

Specifies that the `txn-id` allocated by this declare MUST be associated with the indicated global transaction. If not set, the allocated `txn-id` will be associated with a local transaction. This field MUST NOT be set if the coordinator does not have the `distributed-transactions` capability. Note that the details of distributed transactions within AMQP 1.0 will be provided in a separate specification.

4.5.3 Discharge

Message body for discharging a transaction.

```
<type name="discharge" class="composite" source="list">
  <descriptor name="amqp:discharge:list" code="0x00000000:0x00000032"/>
  <field name="txn-id" type="*" requires="txn-id" mandatory="true"/>
  <field name="fail" type="boolean"/>
</type>
```

The discharge type defines the message body sent to the coordinator to indicate that the `txn-id` is no longer in use. If the transaction is not associated with a `global-id`, then this also indicates the disposition of the local transaction.

Field Details

`txn-id` *identifies the transaction to be discharged*

`fail` *indicates the transaction has failed*

If set, this flag indicates that the work associated with this transaction has failed, and the controller wishes the transaction to be rolled back. If the transaction is associated with a global-id this will render the global transaction rollback-only. If the transaction is a local transaction, then this flag controls whether the transaction is committed or aborted when it is discharged. (Note that the specification for distributed transactions within AMQP 1.0 will be provided separately in Part 6 Distributed Transactions).

4.5.4 Transaction ID

```
<type name="transaction-id" class="restricted" source="binary" provides="txn-id"/>
```

A transaction-id can be up to 32 octets of binary data.

4.5.5 Declared

```
<type name="declared" class="composite" source="list" provides="delivery-state, outcome">
  <descriptor name="amqp:declared:list" code="0x00000000:0x00000033"/>
  <field name="txn-id" type="*" requires="txn-id" mandatory="true"/>
</type>
```

Indicates that a transaction identifier has successfully been allocated in response to a declare message sent to a transaction coordinator.

Field Details

`txn-id` *the allocated transaction id*

4.5.6 Transactional State

The state of a transactional message transfer.

```
<type name="transactional-state" class="composite" source="list" provides="delivery-state">
  <descriptor name="amqp:transactional-state:list" code="0x00000000:0x00000034"/>
  <field name="txn-id" type="*" mandatory="true" requires="txn-id"/>
  <field name="outcome" type="*" requires="outcome"/>
</type>
```

The transactional-state type defines a delivery-state that is used to associate a delivery with a transaction as well as to indicate which outcome is to be applied if the transaction commits.

Field Details

`txn-id` *identifies the transaction with which the state is associated*

`outcome` *provisional outcome*

This field indicates the provisional outcome to be applied if the transaction commits.

4.5.7 Transaction Capability

Symbols indicating (desired/available) capabilities of a transaction coordinator.

```
<type name="txn-capability" class="restricted" source="symbol" provides="txn-capability">
  <choice name="local-transactions" value="amqp:local-transactions"/>
  <choice name="distributed-transactions" value="amqp:distributed-transactions"/>
  <choice name="promotable-transactions" value="amqp:promotable-transactions"/>
  <choice name="multi-txns-per-ssn" value="amqp:multi-txns-per-ssn"/>
  <choice name="multi-ssns-per-txn" value="amqp:multi-ssns-per-txn"/>
</type>
```

Valid Values

amqp:local-transactions

Support local transactions.

amqp:distributed-transactions

Support AMQP Distributed Transactions.

amqp:promotable-transactions

Support AMQP Promotable Transactions.

amqp:multi-txns-per-ssn

Support multiple active transactions on a single session.

amqp:multi-ssns-per-txn

Support transactions whose txn-id is used across sessions on one connection.

4.5.8 Transaction Error

Symbols used to indicate transaction errors.

```
<type name="transaction-error" class="restricted" source="symbol" provides="error-condition">
  <choice name="unknown-id" value="amqp:transaction:unknown-id"/>
  <choice name="transaction-rollback" value="amqp:transaction:rollback"/>
  <choice name="transaction-timeout" value="amqp:transaction:timeout"/>
</type>
```

Valid Values

amqp:transaction:unknown-id

The specified txn-id does not exist.

amqp:transaction:rollback

The transaction was rolled back for an unspecified reason.

amqp:transaction:timeout

The work represented by this transaction took too long.

Part 5: Security

5.1 Security Layers

Security layers are used to establish an authenticated and/or encrypted transport over which regular AMQP traffic can be tunneled. Security layers can be tunneled through one another (for instance a security layer used by the peers to do authentication might be tunneled through a security layer established for encryption purposes).

The framing and protocol definitions for security layers are expected to be defined externally to the AMQP specification as in the case of TLS [RFC5246]. An exception to this is the SASL [RFC4422] security layer which depends on its host protocol to provide framing. Therefore section 5.3 defines the frames necessary for SASL to function. When a security layer terminates (either before or after a secure tunnel is established), the TCP connection MUST be closed by first shutting down the outgoing stream and then reading the incoming stream until it is terminated.

5.2 TLS

To establish a TLS session, each peer MUST start by sending a protocol header before commencing with TLS negotiation. The protocol header consists of the upper case ASCII letters "AMQP" followed by a protocol id of two, followed by three unsigned bytes representing the major, minor, and revision of the specification version (currently 1 (TLS-MAJOR), 0 (TLS-MINOR), 0 (TLS-REVISION)). In total this is an 8-octet sequence:




Figure 5.1: Protocol Header for TLS Security Layer

Other than using a protocol id of two, the exchange of TLS protocol headers follows the same rules specified in the version negotiation section of the transport specification (See Part 2: section 2.2).

The following diagram illustrates the interaction involved in creating a TLS security layer:




Figure 5.2: Establishing a TLS Security Layer

When the use of the TLS security layer is negotiated, the following rules apply:

- The TLS client peer and TLS server peer are determined by the TCP client peer and TCP server peer respectively.

- The TLS client peer SHOULD use the server name indication extension as described in RFC-4366 [RFC4366]. If it does so, then it is undefined what happens if this differs from hostname in the `sasl-init` and `open` frame frames.

This field can be used by AMQP proxies to determine the correct back-end service to connect the client to, and to determine the domain to validate the client's credentials against if TLS client certificates are being used.

- The TLS client MUST validate the certificate presented by the TLS server.
- Implementations MAY choose to use TLS with unidirectional shutdown, i.e., an application initiating shutdown using `close.notify` is not obliged to wait for the peer to respond, and MAY close the write half of the TCP socket.

5.2.1 Alternative Establishment

In certain situations, such as connecting through firewalls, it might not be possible to establish a TLS security layer using the above procedure (for example, because a deep packet inspecting firewall sees the first few bytes of the connection 'as not being TLS').

As an alternative, implementations MAY run a pure TLS server, i.e., one that does not expect the initial TLS-invoking handshake. The IANA service name for this is `amqps` and the port is `SECURE-PORT` (5671). Implementations MAY also choose to run this pure TLS server on other ports, if this is operationally necessary (e.g., to tunnel through a legacy firewall that only expects TLS traffic on port 443).

5.2.2 Constant Definitions

TLS-MAJOR	1	major protocol version.
TLS-MINOR	0	minor protocol version.
TLS-REVISION	0	protocol revision.

5.3 SASL

To establish a SASL layer, each peer MUST start by sending a protocol header. The protocol header consists of the upper case ASCII letters "AMQP" followed by a protocol id of three, followed by three unsigned bytes representing the major, minor, and revision of the specification version (currently 1 (SASL-MAJOR), 0 (SASL-MINOR), 0 (SASL-REVISION)). In total this is an 8-octet sequence:




Figure 5.3: Protocol Header for SASL Security Layer

Other than using a protocol id of three, the exchange of SASL layer headers follows the same rules specified in the version negotiation section of the transport specification (See Part 2: section 2.2).

The following diagram illustrates the interaction involved in creating a SASL security layer:




Figure 5.4: Establishing a SASL Security Layer

5.3.1 SASL Frames

SASL performatives are framed as per Part 2: section 2.3. A SASL frame has a type code of 0x01. Bytes 6 and 7 of the header are ignored. Implementations SHOULD set these to 0x00. The extended header is ignored. Implementations SHOULD therefore set DOFF to 0x02.

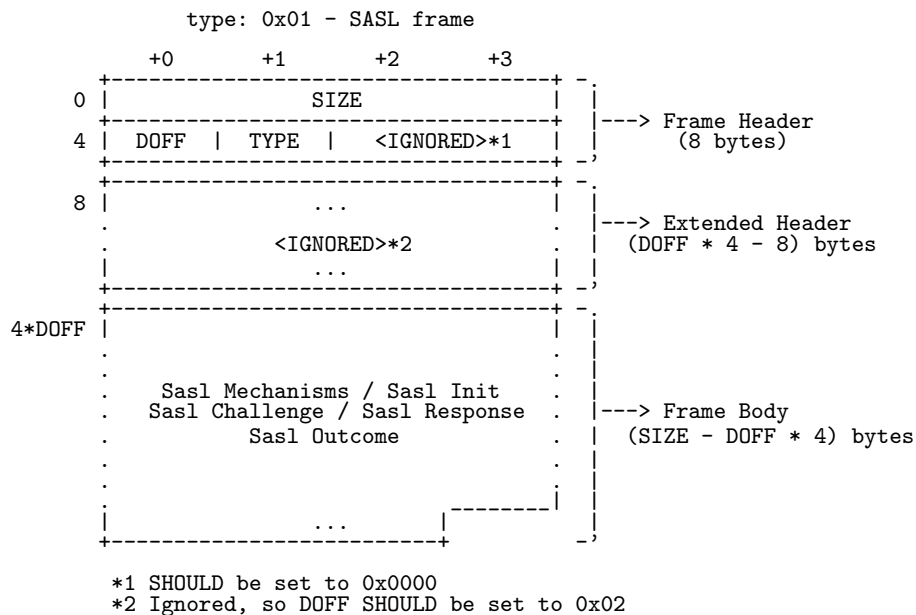


Figure 5.5: SASL Frame

The maximum size of a SASL frame is defined by MIN-MAX-FRAME-SIZE. There is no mechanism within the SASL negotiation to negotiate a different size. The frame body of a SASL frame MUST contain exactly one AMQP type, whose type encoding MUST have provides="sasl-frame". Receipt of an empty frame is an irrecoverable error.

5.3.2 SASL Negotiation

The peer acting as the SASL server MUST announce supported authentication mechanisms using the `sasl-mechanisms` frame. The partner MUST then choose one of the supported mechanisms and initiate a sasl exchange.

```

SASL Client      SASL Server
=====
SASL-INIT        <-- SASL-MECHANISMS
                  -->
                  ...
SASL-RESPONSE    <-- SASL-CHALLENGE *
                  -->
                  ...
                  <-- SASL-OUTCOME
-----
* Note that the SASL
  challenge/response step can
  occur zero or more times
  depending on the details of
  the SASL mechanism chosen.

```

Figure 5.6: SASL Exchange

The peer playing the role of the SASL client and the peer playing the role of the SASL server MUST correspond to the TCP client and server respectively.

5.3.3 Security Frame Bodies

5.3.3.1 SASL Mechanisms

Advertise available sasl mechanisms.

```

<type name="sasl-mechanisms" class="composite" source="list" provides="sasl-frame">
  <descriptor name="amqp:sasl-mechanisms:list" code="0x00000000:0x00000040"/>
  <field name="sasl-server-mechanisms" type="symbol" multiple="true" mandatory="true"/>
</type>

```

Advertises the available SASL mechanisms that can be used for authentication.

Field Details

`sasl-server-mechanisms` *supported sasl mechanisms*

A list of the sasl security mechanisms supported by the sending peer. It is invalid for this list to be null or empty. If the sending peer does not require its partner to authenticate with it, then it SHOULD send a list of one element with its value as the SASL mechanism *ANONYMOUS*. The server mechanisms are ordered in decreasing level of preference.

5.3.3.2 SASL Init

Initiate sasl exchange.

```

<type name="sasl-init" class="composite" source="list" provides="sasl-frame">
  <descriptor name="amqp:sasl-init:list" code="0x00000000:0x00000041"/>
  <field name="mechanism" type="symbol" mandatory="true"/>
  <field name="initial-response" type="binary"/>
  <field name="hostname" type="string"/>
</type>

```

Selects the sasl mechanism and provides the initial response if needed.

Field Details

`mechanism` *selected security mechanism*

The name of the SASL mechanism used for the SASL exchange. If the selected mechanism is not supported by the receiving peer, it MUST close the connection with the authentication-failure close-code. Each peer MUST authenticate using the highest-level security profile it can handle from the list provided by the partner.

`initial-response` *security response data*

A block of opaque data passed to the security mechanism. The contents of this data are defined by the SASL security mechanism.

`hostname` *the name of the target host*

The DNS name of the host (either fully qualified or relative) to which the sending peer is connecting. It is not mandatory to provide the hostname. If no hostname is provided the receiving peer SHOULD select a default based on its own configuration.

This field can be used by AMQP proxies to determine the correct back-end service to connect the client to, and to determine the domain to validate the client's credentials against.

This field might already have been specified by the server name indication extension as described in RFC-4366 [RFC4366], if a TLS layer is used, in which case this field SHOULD either be null or contain the same value. It is undefined what a different value to those already specified means.

5.3.3.3 SASL Challenge

Security mechanism challenge.

```
<type name="sasl-challenge" class="composite" source="list" provides="sasl-frame">
  <descriptor name="amqp:sasl-challenge:list" code="0x00000000:0x00000042"/>
  <field name="challenge" type="binary" mandatory="true"/>
</type>
```

Send the SASL challenge data as defined by the SASL specification.

Field Details

`challenge` *security challenge data*

Challenge information, a block of opaque binary data passed to the security mechanism.

5.3.3.4 SASL Response

Security mechanism response.

```
<type name="sasl-response" class="composite" source="list" provides="sasl-frame">
  <descriptor name="amqp:sasl-response:list" code="0x00000000:0x00000043"/>
  <field name="response" type="binary" mandatory="true"/>
</type>
```

Send the SASL response data as defined by the SASL specification.

Field Details

`response` *security response data*

A block of opaque data passed to the security mechanism. The contents of this data are defined by the SASL security mechanism.

5.3.3.5 SASL Outcome

Indicates the outcome of the sasl dialog.

```
<type name="sasl-outcome" class="composite" source="list" provides="sasl-frame">
  <descriptor name="amqp:sasl-outcome:list" code="0x00000000:0x00000044"/>
  <field name="code" type="sasl-code" mandatory="true"/>
  <field name="additional-data" type="binary"/>
</type>
```

This frame indicates the outcome of the SASL dialog. Upon successful completion of the SASL dialog the security layer has been established, and the peers **MUST** exchange protocol headers to either start a nested security layer, or to establish the AMQP connection.

Field Details

`code` *indicates the outcome of the sasl dialog*
A reply-code indicating the outcome of the SASL dialog.

`additional-data` *additional data as specified in RFC-4422*
The additional-data field carries additional data on successful authentication outcome as specified by the SASL specification [RFC4422]. If the authentication is unsuccessful, this field is not set.

5.3.3.6 SASL Code

Codes to indicate the outcome of the sasl dialog.

```
<type name="sasl-code" class="restricted" source="ubyte">
  <choice name="ok" value="0"/>
  <choice name="auth" value="1"/>
  <choice name="sys" value="2"/>
  <choice name="sys-perm" value="3"/>
  <choice name="sys-temp" value="4"/>
</type>
```

Valid Values

- | | |
|----------|---|
| 0 | Connection authentication succeeded. |
| 1 | Connection authentication failed due to an unspecified problem with the supplied credentials. |
| 2 | Connection authentication failed due to a system error. |
| 3 | Connection authentication failed due to a system error that is unlikely to be corrected without intervention. |
| 4 | Connection authentication failed due to a transient system error. |

5.3.4 Constant Definitions

SASL-MAJOR 1 major protocol version.

SASL-MINOR 0 minor protocol version.

SASL-REVISION 0 protocol revision.

XML Document Type Definition (DTD)

```
1 <?xml version="1.0" encoding="UTF-8"?>
2
3 <!--
4 Copyright OASIS Open 2012. All Rights Reserved.
5
6
7 All capitalized terms in the following text have the meanings assigned to them in the OASIS
8 Intellectual Property Rights Policy (the "OASIS IPR Policy"). The full Policy may be found at the
9 OASIS website.
10
11
12 This document and translations of it may be copied and furnished to others, and derivative works
13 that comment on or otherwise explain it or assist in its implementation may be prepared, copied,
14 published, and distributed, in whole or in part, without restriction of any kind, provided that
15 the above copyright notice and this section are included on all such copies and derivative works.
16 However, this document itself may not be modified in any way, including by removing the copyright
17 notice or references to OASIS, except as needed for the purpose of developing any document or
18 deliverable produced by an OASIS Technical Committee (in which case the rules applicable to
19 copyrights, as set forth in the OASIS IPR Policy, must be followed) or as required to translate
20 it into languages other than English.
21
22
23 The limited permissions granted above are perpetual and will not be revoked by OASIS or its
24 successors or assigns.
25
26
27 This document and the information contained herein is provided on an "AS IS" basis and OASIS
28 DISCLAIMS ALL WARRANTIES, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO ANY WARRANTY THAT THE
29 USE OF THE INFORMATION HEREIN WILL NOT INFRINGE ANY OWNERSHIP RIGHTS OR ANY IMPLIED WARRANTIES OF
30 MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.
31
32
33 OASIS requests that any OASIS Party or any other party that believes it has patent claims that
34 would necessarily be infringed by implementations of this OASIS Committee Specification or OASIS
35 Standard, to notify OASIS TC Administrator and provide an indication of its willingness to grant
36 patent licenses to such patent claims in a manner consistent with the IPR Mode of the OASIS
37 Technical Committee that produced this specification.
38
39
40 OASIS invites any party to contact the OASIS TC Administrator if it is aware of a claim of
41 ownership of any patent claims that would necessarily be infringed by implementations of this
42 specification by a patent holder that is not willing to provide a license to such patent claims in
43 a manner consistent with the IPR Mode of the OASIS Technical Committee that produced this
44 specification. OASIS may include such claims on its website, but disclaims any obligation to do
45 so.
46
47
48 OASIS takes no position regarding the validity or scope of any intellectual property or other
49 rights that might be claimed to pertain to the implementation or use of the technology described
50 in this document or the extent to which any license under such rights might or might not be
51 available; neither does it represent that it has made any effort to identify any such rights.
52 Information on OASIS' procedures with respect to rights in any document or deliverable produced
53 by an OASIS Technical Committee can be found on the OASIS website. Copies of claims of rights
```


54 made available for publication and any assurances of licenses to be made available, or the result
55 of an attempt made to obtain a general license or permission for the use of such proprietary
56 rights by implementers or users of this OASIS Committee Specification or OASIS Standard, can be
57 obtained from the OASIS TC Administrator. OASIS makes no representation that any information or
58 list of intellectual property rights will at any time be complete, or that any claims in such list
59 are, in fact, Essential Claims.
60
61
62 The name "OASIS" is a trademark of OASIS, the owner and developer of this specification, and
63 should be used only to refer to the organization and its official outputs. OASIS welcomes
64 reference to, and implementation and use of, specifications, while reserving the right to enforce
65 its marks against misleading uses. Please see <http://www.oasis-open.org/who/trademark.php> for above guidance.
66 -->
67
68 <!-- The AMQP specification is defined by a collection of XML source
69 files each conforming to the DTD defined below. Each file
70 specifies a distinct "part" of the specification.
71 -->
72
73 <!-- The amqp element is the root element for each source file and
74 thus identifies a distinct part of the AMQP specification. The
75 amqp element may contain any number of doc or section elements
76 and has the following attributes:
77
78 name - The name of the part of the specification defined within
79 this element.
80 label - A short label for the part of the specification defined
81 within this element
82 -->
83 <!ELEMENT amqp (doc|section)*>
84 <!ATTLIST amqp
85 xmlns CDATA #IMPLIED
86 name CDATA #REQUIRED
87 label CDATA #IMPLIED
88 >
89
90 <!-- The section element identifies a section within a part of the
91 AMQP specification. Sections provide a logical grouping of
92 content and definitions within a part of the specification. The
93 section element may contain any number of doc, definition, or
94 type elements and has the following attributes:
95
96 name - The name of the section.
97 title - A title for the section.
98 label - A short label for the section.
99 -->
100 <!ELEMENT section (doc|definition|type)*>
101 <!ATTLIST section
102 name CDATA #REQUIRED
103 title CDATA #IMPLIED
104 label CDATA #IMPLIED
105 >
106
107 <!-- The definition element formally defines a named constant. The
108 definition element may contain any number of doc elements and has
109 the following attributes:
110
111 name - The name of the constant.
112 value - The value of the constant.
113 label - A short label for the constant.

```
114 -->
115 <!ELEMENT definition (doc)*>
116 <!ATTLIST definition
117     name CDATA #REQUIRED
118     value CDATA #REQUIRED
119     label CDATA #IMPLIED
120 >
121
122 <!-- The type element formally defines a semantic type used to
123     exchange data on the wire. Every type definition may have the
124     following attributes:
125
126     name      - The name of the type.
127     label     - A short description of the type.
128     class     - A string identifying what class of type is being
129                 defined.
130     provides  - A comma separated list of archetypes (see field
131                 element).
132
133     There are four different classes of types identified by the
134     "class" attribute: primitive, composite, restricted, and union.
135     All classes of types may contain doc elements.
136
137     A "primitive" type will contain one or more encoding elements
138     that describe how the data is formatted on the wire. Primitive
139     types do not contain descriptor, field, or choice elements.
140
141     A "composite" type definition specifies a new kind of record
142     type, i.e. a type composed of a fixed number of fields whose
143     values are each of a specific type. A composite type does not
144     have a new encoding, but is sent on the wire as a list annotated
145     by a specific descriptor value that identifies it as a
146     representation of the defined composite type. A composite type
147     definition will contain a descriptor and zero or more field
148     elements. Composite types do not contain encoding or choice
149     elements. The source attribute of a composite type will always be
150     "list", other values are reserved for future use.
151
152     A "restricted" type definition specifies a new kind of type whose
153     values are restricted to a subset of the values that may be
154     represented with another type. The source attribute identifies
155     the base type from which a restricted type is derived. A
156     restricted type may have a descriptor element in which case it is
157     identified by a descriptor on the wire. The values permitted by a
158     restricted type may be specified either via documentation, or
159     directly limited to a fixed number of values by the choice
160     elements contained within the definition.
161
162     The "union" class of types is currently reserved.
163 -->
164 <!ELEMENT type (encoding|descriptor|field|choice|doc)*>
165 <!ATTLIST type
166     name CDATA #REQUIRED
167     class (primitive|composite|restricted|union) #REQUIRED
168     source CDATA #IMPLIED
169     provides CDATA #IMPLIED
170     label CDATA #IMPLIED
171 >
172
173 <!-- The encoding element defines how a primitive type is encoded. The
```

```
174 specification defines 4 general categories of encodings: fixed,
175 variable, compound, or array. A specific encoding provides
176 further details of how data is formatted within its general
177 category.
178
179 name - The name of the encoding.
180 label - A short description of the encoding.
181 code - The numeric value that prefixes the encoded data on
182 the wire.
183 category - The category of the encoding: "fixed", "variable",
184 "compound", or "array".
185 width - The width of the encoding or the size/count
186 prefixes depending on the category.
187 -->
188 <!ELEMENT encoding (doc)*>
189 <!ATTLIST encoding
190 name CDATA #IMPLIED
191 label CDATA #IMPLIED
192 code CDATA #REQUIRED
193 category (fixed|variable|compound|array) #REQUIRED
194 width CDATA #IMPLIED
195 >
196
197 <!-- The descriptor element specifies what annotation is used to
198 identify encoded values as representations of a described type.
199
200 name - A symbolic name for the representation.
201 code - The numeric value.
202 -->
203 <!ELEMENT descriptor (doc)*>
204 <!ATTLIST descriptor
205 name CDATA #IMPLIED
206 code CDATA #IMPLIED
207 >
208
209 <!-- The field element identifies a field within a composite type.
210 Every field has the following attributes:
211
212 name - A symbolic name that uniquely identifies the field
213 within the type.
214 type - The type of the field. This attribute defines the
215 range of values that are permitted to appear in
216 this field. It may name a specific type, in which
217 case the values are restricted to that type, or it
218 may be the special character "*", in which case a
219 value of any type is permitted. In the latter case
220 the range of values may be further restricted by
221 the requires attribute.
222 requires - A comma separated list of archetypes. Field values
223 are restricted to types providing at least one of
224 the specified archetypes.
225 default - A default value for the field if no value is encoded.
226 label - A short description of the field.
227 mandatory - "true" iff a non null value for the field is always encoded.
228 multiple - "true" iff the field may have multiple values of its specified type.
229 -->
230 <!ELEMENT field (doc)*>
231 <!ATTLIST field
232 name CDATA #REQUIRED
233 type CDATA #IMPLIED
```

```
234         requires CDATA #IMPLIED
235         default CDATA #IMPLIED
236         label CDATA #IMPLIED
237         mandatory CDATA #IMPLIED
238         multiple CDATA #IMPLIED
239 >
240
241 <!-- The choice element identifies a legal value for a restricted
242      type. The choice element has the following attributes:
243
244         name - A symbolic name for the value.
245         value - The permitted value.
246 -->
247 <!ELEMENT choice (doc)*>
248 <!ATTLIST choice
249         name CDATA #REQUIRED
250         value CDATA #REQUIRED
251 >
252
253 <!-- The doc element identifies the basic unit of documentation that
254      may appear at nearly any point within the specification. A doc
255      element may optionally have a symbolic name and a title for cross
256      reference:
257
258         name - The symbolic name of the doc element.
259         title - The title of the doc element.
260
261      A doc element may contain any number of the following
262      presentational sub elements:
263
264         doc - nested doc elements
265         p - paragraphs
266         ul - unordered lists
267         ol - ordered lists
268         dl - definition lists
269         picture - preformatted ascii art diagrams
270 -->
271
272 <!ELEMENT doc (doc|p|ul|ol|dl|picture)*>
273 <!ATTLIST doc
274         name CDATA #IMPLIED
275         title CDATA #IMPLIED
276 >
277
278 <!-- A paragraph element may be optionally titled and contains
279      freeform text with the following markup elements:
280
281         anchor - a reference point
282         xref - a cross reference
283         b - bold text
284         i - italic text
285         term - a formal term
286         sup - superscript
287         sub - subscript
288         br - line break
289 -->
290 <!ELEMENT p (#PCDATA|anchor|xref|b|i|term|sup|sub|br)*>
291 <!ATTLIST p
292         title CDATA #IMPLIED
293 >
```

```
294
295 <!-- A cross reference. -->
296 <!ELEMENT xref (#PCDATA)>
297 <!ATTLIST xref
298     type CDATA #IMPLIED
299     name CDATA #REQUIRED
300     choice CDATA #IMPLIED
301 >
302
303 <!-- A reference point. -->
304 <!ELEMENT anchor (#PCDATA)>
305 <!ATTLIST anchor
306     name CDATA #REQUIRED
307 >
308
309 <!-- A line break. -->
310 <!ELEMENT br EMPTY>
311
312 <!-- A span of boldface markup. -->
313 <!ELEMENT b (#PCDATA|sub|sup|i|br|anchor)*>
314 <!-- A span of italic markup. -->
315 <!ELEMENT i (#PCDATA|sub|sup|b|br|anchor)*>
316 <!-- A formal term. -->
317 <!ELEMENT term (#PCDATA)>
318 <!-- A span of superscript markup. -->
319 <!ELEMENT sup (#PCDATA|sup|sub|b|i)*>
320 <!-- A span of subscript markup. -->
321 <!ELEMENT sub (#PCDATA|sup|sub|b|i)*>
322
323 <!-- An unordered list. -->
324 <!ELEMENT ul (li)*>
325 <!ATTLIST ul
326     title CDATA #IMPLIED
327 >
328 <!-- An ordered list. -->
329 <!ELEMENT ol (li)*>
330 <!ATTLIST ol
331     title CDATA #IMPLIED
332 >
333
334 <!-- An item in an ordered or unordered list. -->
335 <!ELEMENT li (#PCDATA|p|ul|dl)*>
336
337 <!-- A definition list. -->
338 <!ELEMENT dl (dt, dd)*>
339 <!ATTLIST dl
340     title CDATA #IMPLIED
341 >
342 <!-- A definition term. -->
343 <!ELEMENT dt (#PCDATA)>
344 <!-- A definition description -->
345 <!ELEMENT dd (p|picture|ul|ol)*>
346
347 <!-- A preformatted ascii art diagram. -->
348 <!ELEMENT picture (#PCDATA)>
349 <!ATTLIST picture
350     title CDATA #IMPLIED
351 >
```