

Common Profile for Instant Messaging (CPIM)

Status of this Memo

This document specifies an Internet standards track protocol for the Internet community, and requests discussion and suggestions for improvements. Please refer to the current edition of the "Internet Official Protocol Standards" (STD 1) for the standardization state and status of this protocol. Distribution of this memo is unlimited.

Copyright Notice

Copyright (C) The Internet Society (2004).

Abstract

At the time this document was written, numerous instant messaging protocols were in use, and little interoperability between services based on these protocols has been achieved. This specification defines common semantics and data formats for instant messaging to facilitate the creation of gateways between instant messaging services.

Table of Contents

1.	Introduction	2
2.	Terminology	3
3.	Abstract Instant Messaging Service	4
3.1.	Overview of Instant Messaging Service	4
3.2.	Identification of INSTANT INBOXes	5
3.2.1.	Address Resolution	5
3.3.	Format of Instant Messages	5
3.4.	The Messaging Service	5
3.4.1.	The Message Operation	5
3.4.2.	Looping	6
4.	Security Considerations	7
5.	IANA Considerations	8
5.1.	The IM URI Scheme.	8
6.	Contributors	8
7.	References	9
7.1.	Normative References	9
7.2.	Informative References	9

A.	IM URI IANA Registration Template	10
A.1.	URI Scheme Name	10
A.2.	URI Scheme Syntax	10
A.3.	Character Encoding Considerations	10
A.4.	Intended Usage	10
A.5.	Applications and/or Protocols which use this URI Scheme Name	10
A.6.	Security Considerations	10
A.7.	Relevant Publications	11
A.8.	Person & Email Address to Contact for Further Information.	11
A.9.	Author/Change Controller	11
A.10.	Applications and/or Protocols which use this URI Scheme Name	11
B.	Issues of Interest	11
B.1.	Address Mapping	11
B.2.	Source-Route Mapping	11
C.	Acknowledgments	12
	Author's Address	12
	Full Copyright Statement	13

1. Introduction

Instant messaging is defined in RFC2778 [5]. At the time this document was written, numerous instant messaging protocols are in use, and little interoperability between services based on these protocols has been achieved. This specification defines semantics and data formats for common services of instant messaging to facilitate the creation of gateways between instant messaging services: a common profile for instant messaging (CPIM).

Service behavior is described abstractly in terms of operations invoked between the consumer and provider of a service. Accordingly, each IM service must specify how this behavior is mapped onto its own protocol interactions. The choice of strategy is a local matter, providing that there is a clear relation between the abstract behaviors of the service (as specified in this memo) and how it is faithfully realized by a particular instant messaging service. For example, one strategy might transmit an instant message as textual key/value pairs, another might use a compact binary representation, and a third might use nested containers.

The attributes for each operation are defined using an abstract syntax. Although the syntax specifies the range of possible data values, each IM service must specify how well-formed instances of the abstract representation are encoded as a concrete series of bits.

In order to provide a means for the preservation of end-to-end features (especially security) to pass through instant messaging interoperability gateways, this specification also provides recommendations for instant messaging document formats that could be employed by instant messaging protocols.

2. Terminology

In this document, the key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "NOT RECOMMENDED", "MAY", and "OPTIONAL" are to be interpreted as described in RFC 2119 [1] and indicate requirement levels for compliant implementations.

This memo makes use of the vocabulary defined in RFC 2778 [5]. Terms such as CLOSED, INSTANT INBOX, INSTANT MESSAGE, and OPEN are used in the same meaning as defined therein.

The term 'gateway' used in this document denotes a network element responsible for interworking between diverse instant messaging protocols. Although the instant messaging protocols themselves are diverse, under the model used in this document these protocols can carry a common payload that is relayed by the gateway. Whether these interworking intermediaries should be called 'gateways' or 'relays' is therefore somewhat debatable; for the purposes of this document, they are called 'CPIM gateways'.

The term 'instant messaging service' also derives from RFC 2778, but its meaning changes slightly due to the existence of gateways in the CPIM model. When a client sends an operation to an instant messaging service, that service might either be an endpoint or an intermediary such as a CPIM gateway - in fact, the client should not have to be aware which it is addressing, as responses from either will appear the same.

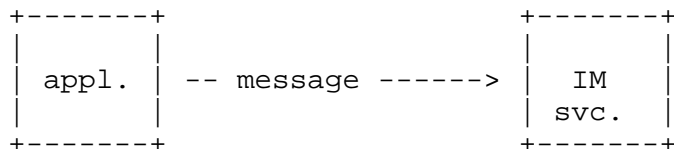
This document defines operations and attributes of an abstract instant messaging protocol. In order for a compliant protocol to interface with an instant messaging gateway, it must support all of the operations described in this document (i.e., the instant messaging protocol must have some message or capability that provides the function described by each of the given operations). Similarly, the attributes defined for these operations must correspond to information available in the instant messaging protocol in order for the protocol to interface with gateways defined by this specification. Note that these attributes provide only the minimum possible information that needs to be specified for interoperability

- the functions in an instant messaging protocol that correspond to the operations described in this document can contain additional information that will not be mapped by CPIM.

3. Abstract Instant Messaging Service

3.1. Overview of Instant Messaging Service

When an application wants to send a message to an INSTANT INBOX, it invokes the message operation, e.g.,

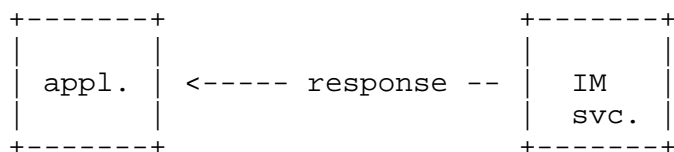


The message operation has the following attributes: source, destination, MaxForwards and TransID. 'source' and 'destination' identify the originator and recipient of an instant message, respectively, and consist of an INSTANT INBOX identifier (as described in Section 3.2). The MaxForwards is a hop counter to avoid loops through gateways, with usage detailed defined in Section 3.4.2; its initial value is set by the originator. The TransID is a unique identifier used to correlate message operations to response operations; gateways should be capable of handling TransIDs up to 40 bytes in length.

The message operation also has some content, the instant message itself, which may be textual, or which may consist of other data. Content details are specified in Section 3.3.

Note that this specification assumes that instant messaging protocols provide reliable message delivery; there are no application-layer message delivery assurance provisions in this specification.

Upon receiving a message operation, the service immediately responds by invoking the response operation containing the same transaction-identifier, e.g.,



The response operation contains the following attributes: TransID and status. The TransID is used to correlate the response to a particular instant message. Status indicates whether the delivery of the message succeeded or failed. Valid status values are described in Section 3.4.1.

3.2. Identification of INSTANT INBOXes

An INSTANT INBOX is specified using an instant messaging URI with the 'im:' URI scheme. The full syntax of the IM URI scheme is given in Appendix A. An example would be: "im:fred@example.com"

3.2.1. Address Resolution

An IM service client determines the next hop to forward the IM to by resolving the domain name portion of the service destination. Compliant implementations SHOULD follow the guidelines for dereferencing URIs given in [2].

3.3. Format of Instant Messages

This specification defines an abstract interoperability mechanism for instant messaging protocols; the message content definition given here pertains to semantics rather than syntax. However, some important properties for interoperability can only be provided if a common end-to-end format for instant messaging is employed by the interoperating instant messaging protocols, especially with respect to security. In order to maintain end-to-end security properties, applications that send message operations to a CPIM gateway MUST implement the format defined in MSGFMT [4]. Applications MAY support other content formats.

CPIM gateways MUST be capable of relaying the content of a message operation between supported instant messaging protocols without needing to modify or inspect the content.

3.4. The Messaging Service

3.4.1. The Message Operation

When an application wants to send an INSTANT MESSAGE, it invokes the message operation.

When an instant messaging service receives the message operation, it performs the following preliminary checks:

1. If the source or destination does not refer to a syntactically valid INSTANT INBOX, a response operation having status "failure" is invoked.
2. If the destination of the operation cannot be resolved by the recipient, and the recipient is not the final recipient, a response operation with the status "failure" is invoked.
3. If access control does not permit the application to request this operation, a response operation having status "failure" is invoked.
4. Provided these checks are successful:

If the instant messaging service is able to successfully deliver the message, a response operation having status "success" is invoked.

If the service is unable to successfully deliver the message, a response operation having status "failure" is invoked.

If the service must delegate responsibility for delivery (i.e., if it is acting as a gateway or proxying the operation), and if the delegation will not result in a future authoritative indication to the service, a response operation having status "indeterminant" is invoked.

If the service must delegate responsibility for delivery, and if the delegation will result in a future authoritative indication to the service, then a response operation is invoked immediately after the indication is received.

When the service invokes the response operation, the transID parameter is identical to the value found in the message operation invoked by the application.

3.4.2. Looping

The dynamic routing of instant messages can result in looping of a message through a relay. Detection of loops is not always obvious, since aliasing and group list expansions can legitimately cause a message to pass through a relay more than one time.

This document assumes that instant messaging protocols that can be gatewayed by CPIM support some semantic equivalent to an integer value that indicates the maximum number of hops through which a message can pass. When that number of hops has been reached, the message is assumed to have looped.

When a CPIM gateway relays an instant message, it decrements the value of the MaxForwards attribute. This document does not mandate any particular initial setting for the MaxForwards element in instant messaging protocols, but it is recommended that the value be reasonably large (over one hundred).

If a CPIM gateway receives an instant message operation that has a MaxForwards attribute of 0, it discards the message and invokes a failure operation.

4. Security Considerations

Detailed security considerations for instant messaging protocols are given in RFC 2779 [6] (in particular, requirements are given in section 5.4 and some motivating discussion with 8.1).

CPIM defines an interoperability function that is employed by gateways between instant messaging protocols. CPIM gateways **MUST** be compliant with the minimum security requirements of the instant messaging protocols with which they interface.

The introduction of gateways to the security model of instant messaging in RFC 2779 also introduces some new risks. End-to-end security properties (especially confidentiality and integrity) between instant messaging user agents that interface through a CPIM gateway can only be provided if a common instant message format (such as the format described in MSGFMT [4]) is supported by the protocols interfacing with the CPIM gateway.

When end-to-end security is required, the message operation **MUST** use MSGFMT, and **MUST** secure the MSGFMT MIME body with S/MIME [8], with encryption (CMS EnvelopeData) and/or S/MIME signatures (CMS SignedData).

The S/MIME algorithms are set by CMS [9]. The AES [11] algorithm should be preferred, as it is expected that AES best suits the capabilities of many platforms. Implementations **MAY** use AES as an encryption algorithm, but are **REQUIRED** to support only the baseline algorithms mandated by S/MIME and CMS.

When IM URIs are placed in instant messaging protocols, they convey the identity of the sender and/or the recipient. Certificates that are used for S/MIME IM operations SHOULD, for the purposes of reference integrity, contain a subjectAltName field containing the IM URI of their subject. Note that such certificates may also contain other identifiers, including those specific to particular instant messaging protocols. In order to further facilitate interoperability of secure messaging through CPIM gateways, users and service providers are encouraged to employ trust anchors for certificates that are widely accepted rather than trust anchors specific to any particular instant messaging service or provider.

In some cases, anonymous messaging may be desired. Such a capability is beyond the scope of this specification.

5. IANA Considerations

The IANA has assigned the "im" scheme.

5.1. The IM URI Scheme

The Instant Messaging (IM) URI scheme designates an Internet resource, namely an INSTANT INBOX.

The syntax of an IM URI is given in Appendix A.

6. Contributors

Dave Crocker edited earlier versions of this document.

The following individuals made substantial textual contributions to this document:

Athanassios Diacakis (thanos.diacakis@openwave.com)

Florencio Mazzoldi (flo@networkprojects.com)

Christian Huitema (huitema@microsoft.com)

Graham Klyne (gk@ninebynine.org)

Jonathan Rosenberg (jdrosen@dynamicsoft.com)

Robert Sparks (rsparks@dynamicsoft.com)

Hiroyasu Sugano (suga@flab.fujitsu.co.jp)

7. References

7.1. Normative References

- [1] Bradner, S., "Key words for use in RFCs to indicate requirement levels", BCP 14, RFC 2119, March 1997.
- [2] Peterson, J., "Address Resolution for Instant Messaging and Presence", RFC 3861, August 2004.
- [3] Resnick, P., "Internet Message Format", STD 11, RFC 2822, April 2001.
- [4] Atkins, D. and G. Klyne, "Common Presence and Instant Messaging: Message Format", RFC 3862, August 2004.
- [5] Day, M., Rosenberg, J., and H. Sugano, "A Model for Presence and Instant Messaging", RFC 2778, February 2000.
- [6] Day, M., Aggarwal, S., and J. Vincent, "Instant Messaging / Presence Protocol Requirements", RFC 2779, February 2000.
- [7] Allocchio, C., "GSTN Address Element Extensions in Email Services", RFC 2846, June 2000.
- [8] Ramsdell, B., "Secure/Multipurpose Internet Mail Extensions (S/MIME) Version 3.1 Message Specification", RFC 3851, July 2004.
- [9] Housley, R., "Cryptographic Message Syntax (CMS)", RFC 3852, July 2004.
- [10] Berners-Lee, T., Fielding, R., and L. Masinter, "Uniform Resource Identifiers (URI): Generic Syntax", RFC 2396, August 1998.

7.2. Informative References

- [11] Schaad, J., "Use of the Advanced Encryption Standard (AES) Encryption Algorithm and in Cryptographic Message Syntax (CMS)", RFC 3565, August 2003.

Appendix A. IM URI IANA Registration Template

This section provides the information to register the im: instant messaging URI.

A.1. URI Scheme Name

im

A.2. URI Scheme Syntax

The syntax follows the existing mailto: URI syntax specified in RFC 2368. The ABNF is:

```
IM-URI      = "im:" [ to ] [ headers ]
to           = mailbox
headers      = "?" header *( "&" header )
header       = hname "=" hvalue
hname        = *uric
hvalue       = *uric
```

Here the symbol "mailbox" represents an encoded mailbox name as defined in RFC 2822 [3], and the symbol "uric" denotes any character that is valid in a URL (defined in RFC 2396 [10]).

A.3. Character Encoding Considerations

Representation of non-ASCII character sets in local-part strings is limited to the standard methods provided as extensions to RFC 2822 [3].

A.4. Intended Usage

Use of the im: URI follows closely usage of the mailto: URI. That is, invocation of an IM URI will cause the user's instant messaging application to start, with destination address and message headers fill-in according to the information supplied in the URI.

A.5. Applications and/or Protocols which use this URI Scheme Name

It is anticipated that protocols compliant with RFC 2779, and meeting the interoperability requirements specified here, will make use of this URI scheme name.

A.6. Security Considerations

See Section 4.

A.7. Relevant Publications

RFC 2779, RFC 2778

A.8. Person & Email Address to Contact for Further Information

Jon Peterson [mailto:jon.peterson@neustar.biz]

A.9. Author/Change Controller

This scheme is registered under the IETF tree. As such, IETF maintains change control.

A.10. Applications and/or Protocols which use this URI Scheme Name

Instant messaging service

Appendix B. Issues of Interest

This appendix briefly discusses issues that may be of interest when designing an interoperation gateway.

B.1. Address Mapping

When mapping the service described in this memo, mappings that place special information into the im: address local-part MUST use the meta-syntax defined in RFC 2846 [7].

B.2. Source-Route Mapping

The easiest mapping technique is a form of source-routing and usually is the least friendly to humans having to type the string. Source-routing also has a history of operational problems.

Use of source-routing for exchanges between different services is by a transformation that places the entire, original address string into the im: address local part and names the gateway in the domain part.

For example, if the destination INSTANT INBOX is "pepp://example.com/fred", then, after performing the necessary character conversions, the resulting mapping is:

im:pepp=example.com/fred@relay-domain

where "relay-domain" is derived from local configuration information.

Experience shows that it is vastly preferable to hide this mapping from end-users - if possible, the underlying software should perform the mapping automatically.

Appendix C. Acknowledgments

The author would like to acknowledge John Ramsdell for his comments, suggestions and enthusiasm. Thanks to Derek Atkins for editorial fixes.

Author's Address

Jon Peterson
NeuStar, Inc.
1800 Sutter St
Suite 570
Concord, CA 94520
US

Phone: +1 925/363-8720
EMail: jon.peterson@neustar.biz

Full Copyright Statement

Copyright (C) The Internet Society (2004). This document is subject to the rights, licenses and restrictions contained in BCP 78, and except as set forth therein, the authors retain all their rights.

This document and the information contained herein are provided on an "AS IS" basis and THE CONTRIBUTOR, THE ORGANIZATION HE/SHE REPRESENTS OR IS SPONSORED BY (IF ANY), THE INTERNET SOCIETY AND THE INTERNET ENGINEERING TASK FORCE DISCLAIM ALL WARRANTIES, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO ANY WARRANTY THAT THE USE OF THE INFORMATION HEREIN WILL NOT INFRINGE ANY RIGHTS OR ANY IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.

Intellectual Property

The IETF takes no position regarding the validity or scope of any Intellectual Property Rights 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; nor does it represent that it has made any independent effort to identify any such rights. Information on the procedures with respect to rights in RFC documents can be found in BCP 78 and BCP 79.

Copies of IPR disclosures made to the IETF Secretariat 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 specification can be obtained from the IETF on-line IPR repository at <http://www.ietf.org/ipr>.

The IETF invites any interested party to bring to its attention any copyrights, patents or patent applications, or other proprietary rights that may cover technology that may be required to implement this standard. Please address the information to the IETF at ietf-ipr@ietf.org.

Acknowledgement

Funding for the RFC Editor function is currently provided by the Internet Society.

