

Internet Engineering Task Force (IETF)
Request for Comments: 7572
Category: Standards Track
ISSN: 2070-1721

P. Saint-Andre
&yet
A. Hour
IBM
J. Hildebrand
Cisco Systems, Inc.
June 2015

Interworking between the Session Initiation Protocol (SIP) and the
Extensible Messaging and Presence Protocol (XMPP): Instant Messaging

Abstract

This document defines a bidirectional protocol mapping for the exchange of single instant messages between the Session Initiation Protocol (SIP) and the Extensible Messaging and Presence Protocol (XMPP).

Status of This Memo

This is an Internet Standards Track document.

This document is a product of the Internet Engineering Task Force (IETF). It represents the consensus of the IETF community. It has received public review and has been approved for publication by the Internet Engineering Steering Group (IESG). Further information on Internet Standards is available in Section 2 of RFC 5741.

Information about the current status of this document, any errata, and how to provide feedback on it may be obtained at <http://www.rfc-editor.org/info/rfc7572>.

Copyright Notice

Copyright (c) 2015 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to BCP 78 and the IETF Trust's Legal Provisions Relating to IETF Documents (<http://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document. Code Components extracted from this document must include Simplified BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

Table of Contents

| | |
|--|----|
| 1. Introduction | 2 |
| 2. Intended Audience | 3 |
| 3. Terminology | 3 |
| 4. XMPP to SIP | 4 |
| 5. SIP to XMPP | 6 |
| 6. Message Size | 8 |
| 7. Content Types | 8 |
| 8. Internationalization Considerations | 9 |
| 9. Security Considerations | 9 |
| 10. References | 10 |
| 10.1. Normative References | 10 |
| 10.2. Informative References | 11 |
| Acknowledgements | 13 |
| Authors' Addresses | 13 |

1. Introduction

In order to help ensure interworking between instant messaging (IM) systems that conform to the instant messaging / presence requirements [RFC2779], it is important to clearly define protocol mappings between such systems. Within the IETF, work has proceeded on two instant messaging technologies:

- o Various extensions to the Session Initiation Protocol ([RFC3261]) for instant messaging, in particular the MESSAGE method extension [RFC3428]; collectively the capabilities of SIP with these extensions are commonly called SIP for Instant Messaging and Presence Leveraging Extensions (SIMPLE).
- o The Extensible Messaging and Presence Protocol (XMPP), which consists of a formalization of the core XML streaming protocols developed originally by the Jabber open-source community; the relevant specifications are [RFC6120] for the XML streaming layer and [RFC6121] for basic presence and instant messaging extensions.

One approach to helping ensure interworking between these protocols is to map each protocol to the abstract semantics described in [RFC3860]; that is the approach taken by [SIMPLE-CPIM] and [RFC3922]. In contrast, the approach taken in this document is to directly map semantics from one protocol to another (i.e., from SIP / SIMPLE to XMPP and vice versa), since that is how existing systems solve the interworking problem.

Both XMPP systems and IM-capable SIP systems enable entities to exchange "instant messages". The term "instant message" usually refers to a message sent between two entities for delivery in close

to real time (rather than a message that is stored and forwarded to the intended recipient upon request). This document specifies mappings only for single messages (sometimes called "pager-mode" messaging), since they form the lowest common denominator for IM. Separate documents cover "session-mode" instant messaging in the form of one-to-one chat sessions [RFC7573] and multi-party chat sessions [GROUPCHAT]. In particular, session-mode instant messaging supports several features that are not part of pager-mode instant messaging, such as a higher level of assurance regarding end-to-end message delivery. As with all of these documents, the architectural assumptions underlying such direct mappings are provided in [RFC7247], including mapping of addresses and error conditions.

2. Intended Audience

The documents in this series are intended for use by software developers who have an existing system based on one of these technologies (e.g., SIP) and who would like to enable communication from that existing system to systems based on the other technology (e.g., XMPP). We assume that readers are familiar with the core specifications for both SIP [RFC3261] and XMPP [RFC6120], with the base document for this series [RFC7247], and with the following IM-related specifications:

- o "Session Initiation Protocol (SIP) Extension for Instant Messaging" [RFC3428]
- o "Extensible Messaging and Presence Protocol (XMPP): Instant Messaging and Presence" [RFC6121]

Note well that not all protocol-compliant messages are shown (such as SIP 100 TRYING messages), in order to focus the reader on the essential aspects of the protocol flows.

3. Terminology

A number of terms used here are explained in [RFC3261], [RFC3428], [RFC6120], and [RFC6121].

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "NOT RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in [RFC2119].

4. XMPP to SIP

As described in [RFC6121], a single instant message is an XML <message/> stanza of type "normal" sent over an XML stream (since "normal" is the default for the 'type' attribute of the <message/> stanza, the attribute is often omitted).

When the XMPP user Juliet with a Jabber Identifier (JID) of <juliet@example.com> wants to send an instant message to Romeo, she interacts with her XMPP client, which generates an XMPP <message/> stanza. The syntax of the <message/> stanza, including required and optional elements and attributes, is defined in [RFC6121] (for single instant messages, Section 5.1 of [RFC6121] recommends that the value of the 'to' address be a "bare JID" of the form "localpart@domainpart"). The following is an example of such a stanza:

Example 1: XMPP User Sends Message

```
| <message from='juliet@example.com/yn0cl4bnw0yr3vym'  
|   to='romeo@example.net'>  
|   <body>Art thou not Romeo, and a Montague?</body>  
| </message>
```

Upon receiving such a message stanza, the XMPP server needs to determine the identity of the domainpart in the 'to' address, which it does by following the procedures explained in Section 5 of [RFC7247]. If the domain is a SIP domain, the XMPP server will hand off the message stanza to an XMPP-to-SIP gateway or connection manager that natively communicates with IM-aware SIP servers.

The XMPP-to-SIP gateway is then responsible for translating the XMPP message stanza into a SIP MESSAGE request from the XMPP user to the SIP user:

Example 2: XMPP User Sends Message (SIP Transformation)

```
| MESSAGE sip:romeo@example.net SIP/2.0  
| Via: SIP/2.0/TCP x2s.example.com;branch=z9hG4bK776sgdkse  
| Max-Forwards: 70  
| To: sip:romeo@example.net  
| From: <sip:juliet@example.com;gr=yn0cl4bnw0yr3vym>;tag=12345  
| Call-ID: D9AA95FD-2BD5-46E2-AF0F-6CFAA96BDDFA  
| CSeq: 1 MESSAGE  
| Content-Type: text/plain  
| Content-Length: 35  
|  
| Art thou not Romeo, and a Montague?
```

The destination SIP server is responsible for delivering the message to the intended recipient, and the recipient is responsible for generating a response (e.g., 200 OK).

Example 3: SIP User Agent Indicates Receipt of Message

```
| SIP/2.0 200 OK
| Via: SIP/2.0/TCP x2s.example.com;branch=z9hG4bK776sgdkse
| From: sip:juliet@example.com;tag=12345
| To: sip:romeo@example.net;tag=vwxyz
| Call-ID: D9AA95FD-2BD5-46E2-AF0F-6CFAA96BDDFA
| CSeq: 1 MESSAGE
| Content-Length: 0
```

As described in [RFC3428], a downstream proxy could fork a MESSAGE request, but it would return only one 200 OK to the gateway.

Note: This document does not specify handling of the 200 OK by the XMPP-to-SIP gateway (e.g., to enable message acknowledgements). See [RFC7573] for a mapping of message acknowledgements in the context of one-to-one chat sessions.

The mapping of XMPP syntax to SIP syntax MUST be as shown in the following table.

Table 1: Message Syntax Mapping from XMPP to SIP

| XMPP Element or Attribute | SIP Header or Contents |
|---------------------------|------------------------|
| <body/> | body of MESSAGE |
| <subject/> | Subject |
| <thread/> | Call-ID |
| from | From (1) |
| id | transaction identifier |
| to | To or Request-URI |
| type | (no mapping) (2) |
| xml:lang | Content-Language |

- As shown in the foregoing example and described in [RFC7247], the XMPP-to-SIP gateway MUST map the bare JID ("localpart@domainpart") of the XMPP sender to the SIP From header and include the resourcepart of the "full JID" ("localpart@domainpart/resourcepart") as the Globally Routable User Agent URI (GRUU) portion [RFC5627] of the SIP URI.

2. Because there is no SIP header field that matches the meaning of the XMPP message 'type' values ("normal", "chat", "groupchat", "headline", "error"), no general mapping is possible here.

5. SIP to XMPP

As described in [RFC3428], a single instant message is a SIP MESSAGE request sent from a SIP user agent to an intended recipient who is most generally referenced by an Instant Messaging (IM) URI [RFC3861] of the form <im:user@domain> but who might be referenced by a SIP or SIPS URI of the form <sip:user@domain> or <sips:user@domain>.

When a SIP user Romeo with a SIP URI of <sip:romeo@example.net> wants to send an instant message to Juliet, he interacts with his SIP user agent, which generates a SIP MESSAGE request. The syntax of the MESSAGE request is defined in [RFC3428]. The following is an example of such a request:

Example 4: SIP User Sends Message

```
MESSAGE sip:juliet@example.com SIP/2.0
Via: SIP/2.0/TCP s2x.example.net;branch=z9hG4bKeskdg677
Max-Forwards: 70
To: sip:juliet@example.com
From: sip:romeo@example.net;tag=vwxyz
Call-ID: 9E97FB43-85F4-4A00-8751-1124FD4C7B2E
CSeq: 1 MESSAGE
Content-Type: text/plain
Content-Length: 44

Neither, fair saint, if either thee dislike.
```

Section 5 of [RFC3428] stipulates that a SIP user agent presented with an im: URI should resolve it to a sip: or sips: URI. Therefore, we assume that the Request-URI of a request received by an IM-capable SIP-to-XMPP gateway will contain a sip: or sips: URI. Upon receiving the MESSAGE, the SIP server needs to determine the identity of the domain portion of the Request-URI or To header, which it does by following the procedures explained in Section 5 of [RFC7247]. If the domain is an XMPP domain, the SIP server will hand off the MESSAGE to an associated SIP-to-XMPP gateway or connection manager that natively communicates with XMPP servers.

The SIP-to-XMPP gateway is then responsible for translating the request into an XMPP message stanza from the SIP user to the XMPP user and returning a SIP 200 OK message to the sender:

Example 5: SIP User Sends Message (XMPP Transformation)

```

| <message from='romeo@example.net/dr4hcr0st3lup4c'
|   to='juliet@example.com'>
|   <body>Neither, fair saint, if either thee dislike.</body>
| </message>

```

Note that the stanza-handling rules specified in [RFC6121] allow the receiving XMPP server to deliver a message stanza whose 'to' address is a bare JID ("localpart@domainpart") to multiple connected devices. This is similar to the "forking" of messages in SIP.

The mapping of SIP syntax to XMPP syntax MUST be as shown in the following table.

Table 2: Message Syntax Mapping from SIP to XMPP

| SIP Header or Contents | XMPP Element or Attribute |
|------------------------|---------------------------|
| Call-ID | <thread/> |
| Content-Language | xml:lang |
| CSeq | (no mapping) |
| From | from (1) |
| Subject | <subject/> |
| Request-URI or To | to |
| body of MESSAGE | <body/> |
| transaction identifier | id |

- As shown in the foregoing example and described in [RFC7247], if the IM-capable SIP-to-XMPP gateway has information about the GRUU [RFC5627] of the particular endpoint that sent the SIP message, then it MUST map the sender's address to a full JID ("localpart@domainpart/resourcepart") in the 'from' attribute of the XMPP stanza and include the GRUU as the resourcepart.

When transforming SIP pager-mode messages, an IM-capable SIP-to-XMPP gateway MUST specify no XMPP 'type' attribute or, equivalently, a 'type' attribute whose value is "normal" [RFC6121].

See Section 7 of this document about the handling of SIP message bodies that contain content types other than plain text.

6. Message Size

[RFC3428] specifies that (outside of a media session) the size of a MESSAGE request is not allowed to exceed 1300 bytes. Although, in practice, XMPP instant messages do not often exceed that size, neither [RFC6120] nor [RFC6121] sets an upper limit on the size of XMPP stanzas. However, XMPP server deployments usually do limit the size of stanzas in order to help prevent denial-of-service attacks, and [RFC6120] states that if a server sets a maximum stanza size, then the limit is not allowed to be less than 10,000 bytes. Because of this mismatch, an XMPP-to-SIP gateway SHOULD return a <policy-violation/> stanza error if an XMPP user attempts to send an XMPP message stanza that would result in a SIP MESSAGE greater than 1300 bytes. Although such a gateway might decide to "upgrade" from page mode to session mode using the Message Session Relay Protocol (MSRP) -- thus treating the instant message as part of a chat session as described in [RFC7573] -- such behavior is application-specific and this document provides no guidelines for how to complete such an upgrade.

7. Content Types

SIP requests of type "MESSAGE" are allowed to contain essentially any content type. The recommended procedures for SIP-to-XMPP gateways to use in handling these content types are as follows.

An IM-aware SIP-to-XMPP gateway MUST process SIP messages that contain message bodies of type "text/plain" and MUST encapsulate such message bodies as the XML character data of the XMPP <body/> element.

An IM-aware SIP-to-XMPP gateway SHOULD process SIP messages that contain message bodies of type "text/html"; if so, a gateway MUST transform the "text/html" content into XHTML content that conforms to the XHTML-IM Integration Set specified in [XEP-0071].

Although an IM-aware SIP-to-XMPP gateway MAY process SIP messages that contain message bodies of types other than "text/plain" and "text/html", the handling of such content types is a matter of implementation.

8. Internationalization Considerations

Both XMPP and SIP support the UTF-8 encoding [RFC3629] of Unicode characters [UNICODE] within messages, along with tagging of the language for a particular message (in XMPP via the 'xml:lang' attribute and in SIP via the Content-Language header). Gateways MUST map these language tagging mechanisms if they are present in the original message. Several examples follow, using the "XML Notation" [RFC3987] for Unicode characters outside the ASCII range.

Example 6: SIP User Sends Message

```
MESSAGE sip:juliet@example.com SIP/2.0
Via: SIP/2.0/TCP s2x.example.net;branch=z9hG4bKeskdg677
Max-Forwards: 70
To: sip:juliet@example.com
From: sip:romeo@example.net;tag=vwxyz
Call-ID: 5A37A65D-304B-470A-B718-3F3E6770ACAF
CSeq: 1 MESSAGE
Content-Type: text/plain
Content-Length: 45
Content-Language: cs

Nic z ob&#xC3A9;ho, m&#xC3A1; d&#xC49B;vo spanil&#xC3A1;,
nenavid&#xC3AD;&#xC5A1;-li jedno nebo druh&#xC3A9;.
```

Example 7: SIP User Sends Message (XMPP Transformation)

```
<message from='romeo@example.net'
  to='juliet@example.com'
  xml:lang='cs'>
  <body>
    Nic z ob&#xC3A9;ho, m&#xC3A1; d&#xC49B;vo spanil&#xC3A1;,
    nenavid&#xC3AD;&#xC5A1;-li jedno nebo druh&#xC3A9;.
  </body>
</message>
```

9. Security Considerations

Detailed security considerations are given in the following documents:

- o For instant messaging protocols in general, see [RFC2779]
- o For SIP-based instant messaging, see [RFC3428] and also [RFC3261]
- o For XMPP-based instant messaging, see [RFC6121] and also [RFC6120]

- o For SIP-XMPP interworking in general, see [RFC7247]

This document specifies methods for exchanging "pager-mode" instant messages through a gateway that translates between SIP and XMPP, and [RFC7573] specifies such methods for "session-mode" instant messaging between MSRP and XMPP. Such a gateway MUST be compliant with the minimum security requirements of the textual chat protocols for which it translates (i.e., SIP or MSRP and XMPP).

The addition of gateways to the security model of instant messaging specified in [RFC2779] introduces some new risks. In particular, end-to-end security properties (especially confidentiality and integrity) between instant messaging clients that interface through a gateway can be provided only if common formats are supported. Specification of those common formats is out of scope for this document. For instant messages, it is possible to use the methods described in [RFC3862] and [RFC3923], but those methods are not widely implemented. A more widely implemented, albeit nonstandardized, method for interoperable end-to-end encryption would be Off-the-Record Messaging [OTR].

10. References

10.1. Normative References

- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, DOI 10.17487/RFC2119, March 1997, <<http://www.rfc-editor.org/info/rfc2119>>.
- [RFC3261] Rosenberg, J., Schulzrinne, H., Camarillo, G., Johnston, A., Peterson, J., Sparks, R., Handley, M., and E. Schooler, "SIP: Session Initiation Protocol", RFC 3261, DOI 10.17487/RFC3261, June 2002, <<http://www.rfc-editor.org/info/rfc3261>>.
- [RFC3428] Campbell, B., Ed., Rosenberg, J., Schulzrinne, H., Huitema, C., and D. Gurle, "Session Initiation Protocol (SIP) Extension for Instant Messaging", RFC 3428, DOI 10.17487/RFC3428, December 2002, <<http://www.rfc-editor.org/info/rfc3428>>.
- [RFC3861] Peterson, J., "Address Resolution for Instant Messaging and Presence", RFC 3861, DOI 10.17487/RFC3861, August 2004, <<http://www.rfc-editor.org/info/rfc3861>>.

- [RFC5627] Rosenberg, J., "Obtaining and Using Globally Routable User Agent URIs (GRUUs) in the Session Initiation Protocol (SIP)", RFC 5627, DOI 10.17487/RFC5627, October 2009, <<http://www.rfc-editor.org/info/rfc5627>>.
- [RFC6120] Saint-Andre, P., "Extensible Messaging and Presence Protocol (XMPP): Core", RFC 6120, DOI 10.17487/RFC6120, March 2011, <<http://www.rfc-editor.org/info/rfc6120>>.
- [RFC6121] Saint-Andre, P., "Extensible Messaging and Presence Protocol (XMPP): Instant Messaging and Presence", RFC 6121, DOI 10.17487/RFC6121, March 2011, <<http://www.rfc-editor.org/info/rfc6121>>.
- [RFC7247] Saint-Andre, P., Hourii, A., and J. Hildebrand, "Interworking between the Session Initiation Protocol (SIP) and the Extensible Messaging and Presence Protocol (XMPP): Architecture, Addresses, and Error Handling", RFC 7247, DOI 10.17487/RFC7247, May 2014, <<http://www.rfc-editor.org/info/rfc7247>>.
- [XEP-0071] Saint-Andre, P., "XHTML-IM", XSF XEP 0071, November 2012.

10.2. Informative References

- [GROUPCHAT] Saint-Andre, P., Corretge, S., and S. Loreto, "Interworking between the Session Initiation Protocol (SIP) and the Extensible Messaging and Presence Protocol (XMPP): Groupchat", Work in Progress, draft-ietf-stox-groupchat-11, March 2015.
- [OTR] Goldberg, I., "Off-the-Record Messaging", <<https://otr.cypherpunks.ca/>>.
- [RFC2779] Day, M., Aggarwal, S., Mohr, G., and J. Vincent, "Instant Messaging / Presence Protocol Requirements", RFC 2779, DOI 10.17487/RFC2779, February 2000, <<http://www.rfc-editor.org/info/rfc2779>>.
- [RFC3629] Yergeau, F., "UTF-8, a transformation format of ISO 10646", STD 63, RFC 3629, DOI 10.17487/RFC3629, November 2003, <<http://www.rfc-editor.org/info/rfc3629>>.
- [RFC3860] Peterson, J., "Common Profile for Instant Messaging (CPIM)", RFC 3860, DOI 10.17487/RFC3860, August 2004, <<http://www.rfc-editor.org/info/rfc3860>>.

- [RFC3862] Klyne, G. and D. Atkins, "Common Presence and Instant Messaging (CPIM): Message Format", RFC 3862, DOI 10.17487/RFC3862, August 2004, <<http://www.rfc-editor.org/info/rfc3862>>.
- [RFC3922] Saint-Andre, P., "Mapping the Extensible Messaging and Presence Protocol (XMPP) to Common Presence and Instant Messaging (CPIM)", RFC 3922, DOI 10.17487/RFC3922, October 2004, <<http://www.rfc-editor.org/info/rfc3922>>.
- [RFC3923] Saint-Andre, P., "End-to-End Signing and Object Encryption for the Extensible Messaging and Presence Protocol (XMPP)", RFC 3923, DOI 10.17487/RFC3923, October 2004, <<http://www.rfc-editor.org/info/rfc3923>>.
- [RFC3987] Duerst, M. and M. Suignard, "Internationalized Resource Identifiers (IRIs)", RFC 3987, DOI 10.17487/RFC3987, January 2005, <<http://www.rfc-editor.org/info/rfc3987>>.
- [RFC7573] Saint-Andre, P. and S. Loreto, "Interworking between the Session Initiation Protocol (SIP) and the Extensible Messaging and Presence Protocol (XMPP): One-to-One Text Chat Sessions", RFC 7573, DOI 10.17487/RFC7573, June 2015, <<http://www.rfc-editor.org/info/rfc7573>>.
- [SIMPLE-CPIM] Campbell, B. and J. Rosenberg, "CPIM Mapping of SIMPLE Presence and Instant Messaging", Work in Progress, draft-ietf-simple-cpim-mapping-01, June 2002.
- [UNICODE] The Unicode Consortium, "The Unicode Standard", <<http://www.unicode.org/versions/latest/>>.

Acknowledgements

The authors wish to thank the following individuals for their feedback: Mary Barnes, Dave Cridland, Dave Crocker, Adrian Georgescu, Christer Holmberg, Saul Ibarra Corretge, Olle Johansson, Paul Kyzivat, Salvatore Loreto, Daniel-Constantin Mierla, and Tory Patnoe.

Special thanks to Ben Campbell for his detailed and insightful reviews.

Francis Dupont reviewed the document on behalf of the General Area Review Team.

Spencer Dawkins, Stephen Farrell, and Barry Leiba provided helpful input during IESG review.

The authors gratefully acknowledge the assistance of Markus Isomaki and Yana Stamcheva as the working group chairs and Gonzalo Camarillo and Alissa Cooper as the sponsoring Area Directors.

Peter Saint-Andre wishes to acknowledge Cisco Systems, Inc., for employing him during his work on earlier draft versions of this document.

Authors' Addresses

Peter Saint-Andre
&yet
EMail: peter@andyet.com
URI: <https://andyet.com/>

Avshalom Houri
IBM
Rorberg Building, Pekris 3
Rehovot 76123
Israel
EMail: avshalom@il.ibm.com

Joe Hildebrand
Cisco Systems, Inc.
1899 Wynkoop Street, Suite 600
Denver, CO 80202
United States
EMail: jhildebr@cisco.com

