

Internet Engineering Task Force (IETF)
Request for Comments: 7247
Category: Standards Track
ISSN: 2070-1721

P. Saint-Andre
&yet
A. Houri
IBM
J. Hildebrand
Cisco Systems, Inc.
May 2014

Interworking between the Session Initiation Protocol (SIP) and the
Extensible Messaging and Presence Protocol (XMPP):
Architecture, Addresses, and Error Handling

Abstract

As a foundation for the definition of bidirectional protocol mappings between the Session Initiation Protocol (SIP) and the Extensible Messaging and Presence Protocol (XMPP), this document specifies the architectural assumptions underlying such mappings as well as the mapping of addresses and error conditions.

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/rfc7247>.

Copyright Notice

Copyright (c) 2014 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	3
2. Intended Audience	3
3. Terminology	4
4. Architectural Assumptions	4
5. Interdomain Federation	6
6. Address Mapping	7
6.1. Overview	7
6.2. Local Part Mapping	9
6.3. Instance-Specific Mapping	11
6.4. SIP to XMPP	11
6.5. XMPP to SIP	12
7. Error Mapping	14
7.1. XMPP to SIP	15
7.2. SIP to XMPP	17
8. Security Considerations	20
9. References	21
9.1. Normative References	21
9.2. Informative References	22
Appendix A. Acknowledgements	24

1. Introduction

The IETF has worked on two signaling technologies that can be used for multimedia session negotiation, instant messaging, presence, capabilities discovery, notifications, and other functions served by real-time communication applications:

- o The Session Initiation Protocol [RFC3261], along with various SIP extensions developed within the SIP for Instant Messaging and Presence Leveraging Extensions (SIMPLE) Working Group.
- o The Extensible Messaging and Presence Protocol [RFC6120], along with various XMPP extensions developed by the IETF as well as by the XMPP Standards Foundation (XSF).

Because these technologies are widely deployed, it is important to clearly define mappings between them for the sake of interworking. This document provides the basis for a series of SIP-XMPP interworking specifications by defining architectural assumptions, address translation methods, and error condition mappings. Other documents in this series define mappings for presence, messaging, and media sessions.

The guidelines in this series are based on implementation and deployment experience, and they describe techniques that have worked well in the field with existing systems. In practice, interworking has been achieved through direct protocol mappings, not through mapping to an abstract model as described in, for example, [RFC3859] and [RFC3860]. Therefore, this series describes the direct mapping approach in enough detail for developers of new implementations to achieve practical interworking between SIP systems and XMPP systems.

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 would like to enable communication from that existing system to systems based on the other technology (e.g., XMPP). With regard to this document, we assume that readers are familiar with the core specifications for both SIP and XMPP; with regard to the other documents in this series, we assume that readers are familiar with this document as well as with the relevant SIP and XMPP extensions.

3. Terminology

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

Several examples use the "XML Notation" from the Internationalized Resource Identifier (IRI) specification [RFC3987] to represent Unicode characters outside the ASCII range (e.g., the string "ue" stands for the Unicode character [UNICODE] LATIN SMALL LETTER U WITH DIAERESIS, U+00FC).

In architectural diagrams, SIP traffic is shown using arrows such as "***>", whereas XMPP traffic is shown using arrows such as "...>".

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. Architectural Assumptions

Protocol translation between SIP and XMPP could occur in a number of different entities, depending on the architecture of real-time communication deployments. For example, protocol translation could occur within a multiprotocol server (which uses protocol-specific connection managers to initiate traffic to and accept traffic from clients or other servers natively using SIP/SIMPLE, XMPP, etc.), within a multiprotocol client (which enables a user to establish connections natively with various servers using SIP/SIMPLE, XMPP, etc.), or within a gateway that acts as a dedicated protocol translator (which takes one protocol as input and provides another protocol as output).

This document assumes that the protocol translation will occur within a gateway, specifically:

- o When information needs to flow from an XMPP-based system to a SIP-based system, protocol translation will occur within an "XMPP-to-SIP gateway" that translates XMPP syntax and semantics on behalf of an "XMPP server" (together, these two logical functions comprise an "XMPP service").
- o When information needs to flow from a SIP-based system to an XMPP-based system, protocol translation will occur within a "SIP-to-XMPP gateway" that translates SIP syntax and semantics on behalf of a "SIP server" (together, these two logical functions comprise a "SIP service").


Naturally, these logical functions could occur in one and the same actual entity; we differentiate between them mainly for explanatory purposes (although, in practice, such gateways are indeed fairly common).

Note: This assumption is not meant to discourage protocol translation within multiprotocol clients or servers; instead, this assumption is followed mainly to clarify the discussion and examples so that the protocol translation principles can be more easily understood and can be applied by client and server implementors with appropriate modifications to the examples and terminology.

This document assumes that a gateway will translate directly from one protocol to the other. For the sake of the examples, we further assume that protocol translation will occur within a gateway in the source domain, so that information generated by the user of an XMPP system will be translated by a gateway within the trust domain of that XMPP system, and information generated by the user of a SIP system will be translated by a gateway within the trust domain of that SIP system. However, nothing in this document ought to be taken as recommending against protocol translation at the destination domain.

An architectural diagram for a possible gateway deployment is shown below, where the entities have the following significance and the "#" character is used to show the boundary of a trust domain:

- o romeo@example.net -- a SIP user.
- o example.net -- a SIP server with an associated gateway ("S2X GW") to XMPP.
- o juliet@example.com -- an XMPP user.
- o example.com -- an XMPP server with an associated gateway ("X2S GW") to SIP.



Legend:

XMPP = ... or :

SIP = *

Figure 1: Possible Gateway Deployment Architecture

Note that bidirectional communication between the SIP server and the XMPP server can be established over either SIP or XMPP. If the XMPP user initiates the interaction, then the XMPP server would invoke its XMPP-to-SIP gateway; thus, the communication would occur over SIP. If the SIP user initiates the interaction, then the SIP server would invoke its SIP-to-XMPP gateway; thus, the communication would occur over XMPP.

5. Interdomain Federation

The architectural assumptions underlying this document imply that communication between a SIP system and an XMPP system will take place using interdomain federation: the SIP server invokes its associated SIP-to-XMPP gateway, which communicates with the XMPP server using native XMPP server-to-server methods; similarly, the XMPP server invokes its associated XMPP-to-SIP gateway, which communicates with the SIP server using native SIP server-to-server methods.

When an XMPP server receives an XMPP stanza whose 'to' address specifies or includes a domain other than the domain of the XMPP server, it needs to determine whether the destination domain communicates via XMPP or SIP. To do so, it performs one or more DNS SRV lookups [RFC2782] for "_xmpp-server" records as specified in [RFC6120]. If the response returns a hostname, the XMPP server can attempt XMPP communication. If not, it can determine the appropriate location for SIP communication at the target domain using the procedures specified in [RFC3263].

Similarly, when a SIP server receives a SIP message whose Request-URI or To header specifies or includes a domain other than the domain of the SIP server, it needs to determine whether the destination domain communicates via SIP or XMPP. To do so, it uses the procedures specified in [RFC3263]. If that response returns a hostname, the SIP server can attempt SIP communication. If not, it can perform one or more DNS SRV lookups [RFC2782] for "_xmpp-server" records as specified in [RFC6120].

In both cases, the server in question might have previously determined that the foreign domain communicates via SIP or XMPP, in which case it would not need to perform the relevant DNS lookups. The caching of such information is a matter of implementation and local service policy and is therefore out of scope for this document.

Existing SIP and XMPP server implementations do not typically include the ability to communicate using the other technology (XMPP for SIP implementations, SIP for XMPP implementations). One common architectural pattern is to associate a gateway with the core server implementation (e.g., in XMPP such a gateway might be called a "connection manager"). How exactly such a gateway interacts with the core server to complete tasks such as address lookups and communication with systems that use the other technology is a matter of implementation (e.g., the gateway might be an add-on module that is trusted by the core server to act as a fallback delivery mechanism if the remote domain does not support the server's native communication technology).

Because [RFC6120] specifies a binding of XMPP to TCP, a gateway from SIP to XMPP will need to support TCP as the underlying transport protocol. By contrast, as specified in [RFC3261], either TCP or UDP can be used as the underlying transport for SIP messages, and a given SIP deployment might support only UDP; therefore, a gateway from XMPP to SIP might need to communicate with a SIP server using either TCP or UDP.

6. Address Mapping

6.1. Overview

The basic SIP address format is a 'sip' or 'sips' URI as specified in [RFC3261]. When a SIP entity supports extensions for instant messaging it might be identified by an 'im' URI as specified in the Common Profile for Instant Messaging [RFC3860] (see [RFC3428]), and when a SIP entity supports extensions for presence it might be

identified by a 'pres' URI as specified in the Common Profile for Presence [RFC3859] (see [RFC3856]). SIP entities typically also support the 'tel' URI scheme [RFC3966] and might support other URI schemes as well.

The XMPP address format is specified in [RFC6122] (although note that XMPP URIs [RFC5122] are not used natively on the XMPP network); in addition, [RFC6121] encourages instant messaging and presence applications of XMPP to also support 'im' and 'pres' URIs as specified in [RFC3860] and [RFC3859], respectively, although such support might simply involve leaving resolution of such addresses up to an XMPP server.

In this document we primarily describe mappings for addresses of the form <user@domain>; however, we also provide guidelines for mapping the addresses of specific user agent instances, which take the form of Globally Routable User Agent URIs (GRUUs) in SIP and "resourceparts" in XMPP. Mapping of protocol-specific identifiers (such as telephone numbers) is out of scope for this specification. In addition, we have ruled the mapping of domain names as out of scope for now, since that is a matter for the Domain Name System; specifically, the issue for interworking between SIP and XMPP relates to the translation of fully internationalized domain names (IDNs) into non-internationalized domain names (IDNs are not allowed in the SIP address format but are allowed in the XMPP address via Internationalized Domain Names in Applications; see [RFC6122] and [XMPP-ADDRESS-FORMAT]). Therefore, in the following sections we focus primarily on the local part of an address (these are called variously "usernames", "instant inboxes", "presentities", and "localparts" in the protocols at issue), secondarily on the instance-specific part of an address, and not at all on the domain-name part of an address.

The 'sip'/'sips', 'im'/'pres', and XMPP address schemes allow different sets of characters (although all three allow alphanumeric characters and disallow both spaces and control characters). In some cases, characters allowed in one scheme are disallowed in others; these characters need to be mapped appropriately in order to ensure interworking across systems.

6.2. Local Part Mapping

The local part of a 'sip'/'sips' URI inherits from the "userinfo" rule in [RFC3986] with several changes; here we discuss the SIP "user" rule only (using ABNF as defined in [RFC5234]):

```

user          = 1*( unreserved / escaped / user-unreserved )
user-unreserved = "&" / "=" / "+" / "$" / "," / ";" / "?" / "/"
unreserved    = alphanum / mark
mark          = "-" / "_" / "." / "!" / "~" / "*" / "'"
              / "(" / ")"

```

Here we make the simplifying assumption that the local part of an 'im'/'pres' URI inherits from the "dot-atom-text" rule in [RFC5322] rather than the more complicated "local-part" rule:

```

dot-atom-text = 1*atext *("." 1*atext)
atext         = ALPHA / DIGIT /      ; Any character except
              "!" / "#" / "$" /      ; controls, SP, and
              "%" / "&" / "'" /      ; specials. Used for
              "*" / "+" / "-" /      ; atoms.
              "/" / "=" / "?" /
              "^" / "_" / "`" /
              "{" / "|" / "}" /
              "~"

```

The local part of an XMPP address allows any ASCII character except space, controls, and the " & ' / : < > @ characters.

To summarize the foregoing information, the following table lists the allowed and disallowed characters in the local part of identifiers for each protocol (aside from the alphanumeric, space, and control characters), in order by hexadecimal character number (where each "A" row shows the allowed characters and each "D" row shows the disallowed characters).

+-----+ SIP/SIPS CHARACTERS +-----+	
A	! \$ & ' () * + , - . / ; = ? [\] ^ _ ` { } ~
D	" # % : < > @ [\] ^ _ ` { }
+-----+	
IM/PRES CHARACTERS +-----+	
A	! # \$ % & ' * + - / = ? [\] ^ _ ` { } ~
D	" () , . : ; < > @ [\]
+-----+	
XMPP CHARACTERS +-----+	
A	! # \$ % () * + , - . ; = ? [\] ^ _ ` { } ~
D	" & ' / : < > @
+-----+	

Table 1: Allowed and Disallowed Characters

When transforming the local part of an address from one address format to another, an application SHOULD proceed as follows:

1. Unescape any escaped characters in the source address (e.g., from SIP to XMPP unescape "%23" to "#" per [RFC3986], and from XMPP to SIP unescape "%27" to "'" per [XEP-0106]).
2. Leave unmodified any characters that are allowed in the destination scheme.
3. Escape any characters that are allowed in the source scheme but reserved in the destination scheme, as escaping is defined for the destination scheme. In particular:
 - * Where the destination scheme is a URI (i.e., an 'im', 'pres', 'sip', or 'sips' URI), each reserved character MUST be percent-encoded to "%hexhex" as specified in Section 2.5 of [RFC3986] (e.g., when transforming from XMPP to SIP, encode "#" as "%23").

- * Where the destination scheme is a native XMPP address, each reserved character MUST be encoded to "\hexhex" as specified in [XEP-0106] (e.g., when transforming from SIP to XMPP, encode "'" as "\27").

6.3. Instance-Specific Mapping

The meaning of a resourcepart in XMPP (i.e., the portion of a Jabber ID (JID) after the slash character, such as "foo" in "user@example.com/foo") matches that of a Globally Routable User Agent URI (GRUU) in SIP [RFC5627]. In both cases, these constructs identify a particular device associated with the bare JID ("localpart@domainpart") of an XMPP entity or with the Address of Record (AOR) of a SIP entity. Therefore, it is reasonable to map the value of a "gr" URI parameter to an XMPP resourcepart and vice versa.

The mapping described here does not apply to temporary GRUUs -- only to GRUUs associated with an Address of Record.

The "gr" URI parameter in SIP can contain only characters from the ASCII range (although characters outside the ASCII range can be percent-encoded in accordance with [RFC3986]), whereas an XMPP resourcepart can contain nearly any Unicode character [UNICODE]. Therefore, Unicode characters outside the ASCII range need to be mapped to characters in the ASCII range, as described below.

6.4. SIP to XMPP

The following is a high-level algorithm for mapping a 'sip', 'sips', 'im', or 'pres' URI to an XMPP address:

1. Remove URI scheme.
2. Split at the first '@' character into local part and hostname (mapping the latter is out of scope).
3. Translate any percent-encoded strings ("%hexhex") to percent-decoded octets.
4. Treat result as a UTF-8 string.
5. Translate "&" to "\26", "'" to "\27", and "/" to "\2f", respectively in order to properly handle the characters disallowed in XMPP addresses but allowed in 'sip'/'sips' URIs and 'im'/'pres' URIs as shown in Table 1 above (this is consistent with [XEP-0106]).

6. Apply Nodeprep profile of stringprep [RFC3454] or its replacement (see [RFC6122] and [XMPP-ADDRESS-FORMAT]) for canonicalization (OPTIONAL).
7. Recombine local part with mapped hostname to form a bare JID ("localpart@domainpart").
8. If the (SIP) address contained a "gr" URI parameter, append a slash character "/" and the "gr" value to the bare JID to form a full JID ("localpart@domainpart/resourcepart").

Several examples follow, illustrating steps 3, 5, and 8 described above.

SIP URI	XMPP Address
sip:f%C3%BC@sip.example	fü@sip.example
sip:o'malley@sip.example	o\27malley@sip.example
sip:foo@sip.example;gr=bar	foo@sip.example/bar

In the first example, the string "%C3%BC" is a percent-encoded representation of the UTF-8-encoded Unicode character LATIN SMALL LETTER U WITH DIAERESIS (U+00FC), whereas the string "ü" is the same character shown for documentation purposes using the XML Notation defined in [RFC3987] (in XMPP, it would be sent directly as a UTF-8-encoded Unicode character and not percent-encoded as in a SIP URI to comply with the URI syntax defined in [RFC3986]).

6.5. XMPP to SIP

The following is a high-level algorithm for mapping an XMPP address to a 'sip', 'sips', 'im', or 'pres' URI:

1. Split XMPP address into localpart (mapping described in remaining steps), domainpart (hostname; mapping is out of scope), and resourcepart (specifier for particular device or connection, for which an OPTIONAL mapping is described below).
2. Apply Nodeprep profile of stringprep [RFC3454] or its replacement (see [RFC6122] and [XMPP-ADDRESS-FORMAT]) for canonicalization of the XMPP localpart (OPTIONAL).
3. Translate "\26" to "&", "\27" to "'", and "\2f" to "/", respectively (this is consistent with [XEP-0106]).

4. Determine if the foreign domain supports 'im' and 'pres' URIs (discovered via [RFC2782] lookup as specified in [RFC6121]), else assume that the foreign domain supports 'sip'/'sips' URIs.
5. If converting into 'im' or 'pres' URI, for each byte, if the byte is in the set (), . ; [\] or is a UTF-8 character outside the ASCII range, then percent-encode that byte to "%hexhex" format. If converting into 'sip' or 'sips' URI, for each byte, if the byte is in the set # % [\] ^ ` { | } or is a UTF-8 character outside the ASCII range, then percent-encode that byte to "%hexhex" format.
6. Combine resulting local part with mapped hostname to form local@domain address.
7. Prepend with the string 'im:' (for XMPP <message/> stanzas) or 'pres:' (for XMPP <presence/> stanzas) if foreign domain supports these, else prepend with the string 'sip:' or 'sips:' according to local service policy.
8. If the XMPP address included a resourcepart and the destination URI scheme is 'sip' or 'sips', optionally append the slash character '/' and then append the resourcepart (making sure to percent-encode any UTF-8 characters outside the ASCII range) as the "gr" URI parameter.

Several examples follow, illustrating steps 3, 5, and 8 described above.

XMPP Address	SIP URI
m\26m@xmpp.example	sip:m&m@xmpp.example
tschüss@xmpp.example	sip:tsch%C3%BCss@xmpp.example
baz@xmpp.example/qux	sip:baz@xmpp.example;gr=qux

As above, in the first example the string "ü" is the Unicode character LATIN SMALL LETTER U WITH DIAERESIS (U+00FC) shown for documentation purposes using the XML Notation defined in [RFC3987] (in XMPP, it would be sent directly as a UTF-8-encoded Unicode character and not percent-encoded, whereas the string "%C3%BC" is a percent-encoded representation of the same character).

7. Error Mapping

Various differences between XMPP error conditions and SIP response codes make it hard to provide a comprehensive and consistent mapping between the protocols:

- o Whereas the set of XMPP error conditions is fixed in the core XMPP specification (and supplemented where needed by application-specific extensions), the set of SIP response codes is more open to change, as evidenced by the IANA registry of SIP response codes.
- o XMPP has defined fewer error conditions related to stanza handling (22 are defined in [RFC6120]) than SIP has defined response codes related to message handling (at the date of this writing, 71 SIP response codes are registered with IANA as defined in [RFC3261] and numerous SIP extensions).
- o In many cases, the SIP response codes are more specific than the XMPP error conditions (e.g., from an XMPP perspective the SIP codes "413 Request Entity Too Large" and "414 Request-URI Too Long" are simply two different types of response to the same bad request, and the SIP codes "415 Unsupported Media Type" and "416 Unsupported URI Scheme" are two different responses to a request that is not acceptable).
- o SIP differentiates between responses about a particular endpoint or resource (the 4xx series) and responses about a user, i.e., all of a user's endpoints or resources (the 6xx series). There is no such distinction in XMPP, since the same error condition can be returned in relation to the "bare JID" (localpart@domainpart) of a user or the "full JID" (localpart@domainpart/resourcepart) of a particular endpoint or resource, depending on the 'to' address of the original request.

As a result of these and other factors, the mapping of error conditions and response codes is more of an art than a science. This document provides suggested mappings, but implementations are free to deviate from these mappings if needed. Also, because no XMPP error conditions are equivalent to the provisional (1xx) and successful (2xx) response codes in SIP, this document suggests mappings only for the SIP redirection (3xx), request failure (4xx), server failure (5xx), and global failure (6xx) response code families.

Supplementary information about SIP response codes can be expressed in the "Reason-Phrase" in the Status-Line header, and detailed information about XMPP error conditions can be expressed in the <text/> child of the <error/> element. Although the semantics of

these constructs are specified in a slightly different way, it is reasonable for a gateway to map these constructs to each other if they are found in a SIP response or XMPP error stanza.

7.1. XMPP to SIP

The mapping of specific XMPP error conditions to SIP response codes SHOULD be as described in the following table.

XMPP Error Condition	SIP Response Code
<bad-request/>	400
<conflict/>	400
<feature-not-implemented/>	405 or 501 (1)
<forbidden/>	403 or 603 (2)
<gone/>	301 or 410 (3)
<internal-server-error/>	500
<item-not-found/>	404 or 604 (2)
<jid-malformed/>	400
<not-acceptable/>	406 or 606 (2)
<not-allowed/>	403
<not-authorized/>	401
<policy-violation/>	403
<recipient-unavailable/>	480 or 600 (2)
<redirect/>	302
<registration-required/>	407
<remote-server-not-found/>	404 or 408 (4)
<remote-server-timeout/>	408

<resource-constraint/>	500	
<service-unavailable/>	see note (5) below	
<subscription-required/>	400	
<undefined-condition/>	400	
<unexpected-request/>	491 or 400	

Table 2: Mapping of XMPP Error Conditions to SIP Response Codes

- (1) If the error relates to a "full JID" (localpart@domainpart/resourcepart), the SIP 405 response code is RECOMMENDED. If the error relates to a "bare JID" (localpart@domainpart), the SIP 501 response code is RECOMMENDED.
- (2) If the error relates to a "full JID" (localpart@domainpart/resourcepart), the SIP response code from the 4xx series is RECOMMENDED. If the error relates to a "bare JID" (localpart@domainpart), the SIP response code from the 6xx series is RECOMMENDED.
- (3) If the <gone/> element includes XML character data specifying the new address, the error MUST be mapped to SIP 301; if not, it MUST be mapped to SIP 410.
- (4) The XMPP <remote-server-not-found/> error can mean that the remote server either (a) does not exist or (b) cannot be resolved. SIP has two different response codes here: 404 to cover (a) and 408 to cover (b).
- (5) The XMPP <service-unavailable/> error condition is widely used to inform the requesting entity that the intended recipient does not support the relevant feature, to signal that a server cannot perform the requested service either generally or in relation to a particular user, and to avoid disclosing whether a given account exists at all. This is quite different from the semantics of the SIP 503 Service Unavailable response code, which is used to signal that communication with a server is impossible (e.g., even if the XMPP <service-unavailable/> error condition is returned in relation to a specific user, the SIP 503 response code will be interpreted as applying to all future requests to this server, not just requests for the specific user). Therefore, mapping the XMPP <service-unavailable/> error condition to the SIP 503 Service Unavailable response code is

NOT RECOMMENDED. Although no precise mapping is available, the SIP 403 Forbidden and 405 Method Not Allowed response codes are closest in meaning to the XMPP <service-unavailable/> error condition.

7.2. SIP to XMPP

The mapping of SIP response codes to XMPP error conditions SHOULD be as described in the following table. If a gateway encounters a SIP response code that is not listed below, it SHOULD map a 3xx-series code to <redirect/>, a 4xx-series code to <bad-request/>, a 5xx-series code to <internal-server-error/>, and a 6xx-series code to <recipient-unavailable/>.

SIP Response Code	XMPP Error Condition
3xx	<redirect/>
300	<redirect/>
301	<gone/> (1)
302	<redirect/>
305	<redirect/>
380	<not-acceptable/>
4xx	<bad-request/>
400	<bad-request/>
401	<not-authorized/>
402	<bad-request/> (2)
403	<forbidden/> (3)
404	<item-not-found/> (4)
405	<feature-not-implemented/>
406	<not-acceptable/>
407	<registration-required/>

408	<remote-server-timeout/> (5)
410	<gone/> (1)
413	<policy-violation/>
414	<policy-violation/>
415	<not-acceptable/>
416	<not-acceptable/>
420	<feature-not-implemented/>
421	<not-acceptable/>
423	<resource-constraint/>
430	<recipient-unavailable/> (6)
439	<feature-not-implemented/> (6)
440	<policy-violation/> (7)
480	<recipient-unavailable/>
481	<item-not-found/>
482	<not-acceptable/>
483	<not-acceptable/>
484	<item-not-found/>
485	<item-not-found/>
486	<recipient-unavailable/>
487	<recipient-unavailable/>
488	<not-acceptable/>
489	<policy-violation/> (8)
491	<unexpected-request/>

493	<bad-request/>	
5xx	<internal-server-error/>	
500	<internal-server-error/>	
501	<feature-not-implemented/>	
502	<remote-server-not-found/>	
503	<internal-server-error/> (9)	
504	<remote-server-timeout/>	
505	<not-acceptable/>	
513	<policy-violation/>	
6xx	<recipient-unavailable/>	
600	<recipient-unavailable/>	
603	<recipient-unavailable/>	
604	<item-not-found/>	
606	<not-acceptable/>	

Table 3: Mapping of SIP Response Codes to XMPP Error Conditions

- (1) When mapping SIP 301 to XMPP <gone/>, the <gone/> element MUST include XML character data specifying the new address. When mapping SIP 410 to XMPP <gone/>, the <gone/> element MUST NOT include XML character data specifying a new address.
- (2) The XMPP <payment-required/> error condition was removed in [RFC6120]. Therefore, a mapping to XMPP <bad-request/> is suggested instead.
- (3) Depending on the scenario, other possible translations for SIP 403 are <not-allowed/> and <policy-violation/>.
- (4) Depending on the scenario, another possible translation for SIP 404 is <remote-server-not-found/>.

- (5) Depending on the scenario, another possible translation for SIP 408 is `<remote-server-not-found/>`.
- (6) Codes 430 and 439 are defined in [RFC5626].
- (7) Code 440 is defined in [RFC5393].
- (8) Code 489 is defined in [RFC6665].
- (9) Regarding the semantic mismatch between XMPP `<service-unavailable/>` and SIP code 503, see note (5) in Section 7.1 of this document.

8. Security Considerations

Detailed security considerations for SIP and XMPP are given in [RFC3261] and [RFC6120], respectively.

To protect information sent between SIP and XMPP systems, deployed gateways SHOULD use Transport Layer Security (TLS) [RFC5246] when communicating over TCP and Datagram Transport Layer Security (DTLS) [RFC6347] when communicating over UDP.

As specified in Section 26.4.4 of [RFC3261] and updated by [RFC5630], a To header or a Request-URI containing a Session Initiation Protocol Secure (SIPS) URI is used to indicate that all hops in a communication path need to be protected using TLS. Because XMPP lacks a way to signal that all hops need to be protected, if the To header or Request-URI of a SIP message is a SIPS URI then the SIP-to-XMPP gateway MUST NOT translate the SIP message into an XMPP stanza and MUST NOT route it to the destination XMPP server (there might be exceptions to such a policy, such as explicit agreement among two operators to enforce per-hop security, but currently they are quite rare).

A gateway between SIP and XMPP (in either direction) effectively acts as a SIP back-to-back user agent ("B2BUA"). The amplification vulnerability described in [RFC5393] can manifest itself with B2BUAs (see also [B2BUA-LOOP-DETECT]), and a gateway SHOULD implement the loop-detection methods defined in that specification to help mitigate the possibility of amplification attacks. Note that although it would be possible to signal the Max-Forwards and Max-Breadth SIP headers over XMPP using the Stanza Headers and Internet Metadata (SHIM) extension [XEP-0131], that extension is not widely implemented; therefore, defenses against excessive looping and amplification attacks when messages pass back and forth through SIP and XMPP networks are out of scope for this document. However, it

ought to be addressed in the future, and implementations are strongly encouraged to incorporate appropriate countermeasures wherever possible.

The ability to use a wide range of Unicode characters [UNICODE] can lead to security issues, especially the possibility of user confusion over identifiers containing visually similar characters (also called "confusable characters" or "confusables"). The PRECIS framework specification [PRECIS] describes some of these security issues, and additional guidance can be found in [UTS39].

9. References

9.1. Normative References

- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.
- [RFC3261] Rosenberg, J., Schulzrinne, H., Camarillo, G., Johnston, A., Peterson, J., Sparks, R., Handley, M., and E. Schooler, "SIP: Session Initiation Protocol", RFC 3261, June 2002.
- [RFC3263] Rosenberg, J. and H. Schulzrinne, "Session Initiation Protocol (SIP): Locating SIP Servers", RFC 3263, June 2002.
- [RFC3986] Berners-Lee, T., Fielding, R., and L. Masinter, "Uniform Resource Identifier (URI): Generic Syntax", STD 66, RFC 3986, January 2005.
- [RFC3987] Duerst, M. and M. Suignard, "Internationalized Resource Identifiers (IRIs)", RFC 3987, January 2005.
- [RFC5234] Crocker, D. and P. Overell, "Augmented BNF for Syntax Specifications: ABNF", STD 68, RFC 5234, January 2008.
- [RFC5246] Dierks, T. and E. Rescorla, "The Transport Layer Security (TLS) Protocol Version 1.2", RFC 5246, August 2008.
- [RFC5393] Sparks, R., Lawrence, S., Hawrylyshen, A., and B. Campen, "Addressing an Amplification Vulnerability in Session Initiation Protocol (SIP) Forking Proxies", RFC 5393, December 2008.
- [RFC5627] Rosenberg, J., "Obtaining and Using Globally Routable User Agent URIs (GRUUs) in the Session Initiation Protocol (SIP)", RFC 5627, October 2009.

- [RFC5630] Audet, F., "The Use of the SIPS URI Scheme in the Session Initiation Protocol (SIP)", RFC 5630, October 2009.
- [RFC6120] Saint-Andre, P., "Extensible Messaging and Presence Protocol (XMPP): Core", RFC 6120, March 2011.
- [RFC6122] Saint-Andre, P., "Extensible Messaging and Presence Protocol (XMPP): Address Format", RFC 6122, March 2011.
- [RFC6347] Rescorla, E. and N. Modadugu, "Datagram Transport Layer Security Version 1.2", RFC 6347, January 2012.
- [UNICODE] The Unicode Consortium, "The Unicode Standard, Version 6.3", 2013,
<<http://www.unicode.org/versions/Unicode6.3.0/>>.

9.2. Informative References

- [B2BUA-LOOP-DETECT] Kaplan, H. and V. Pascual, "Loop Detection Mechanisms for Session Initiation Protocol (SIP) Back-to-Back User Agents (B2BUAs)", Work in Progress, February 2014.
- [PRECIS] Saint-Andre, P. and M. Blanchet, "PRECIS Framework: Preparation and Comparison of Internationalized Strings in Application Protocols", Work in Progress, April 2014.
- [RFC2782] Gulbrandsen, A., Vixie, P., and L. Esibov, "A DNS RR for specifying the location of services (DNS SRV)", RFC 2782, February 2000.
- [RFC3428] Campbell, B., Rosenberg, J., Schulzrinne, H., Huitema, C., and D. Gurle, "Session Initiation Protocol (SIP) Extension for Instant Messaging", RFC 3428, December 2002.
- [RFC3454] Hoffman, P. and M. Blanchet, "Preparation of Internationalized Strings ("stringprep")", RFC 3454, December 2002.
- [RFC3856] Rosenberg, J., "A Presence Event Package for the Session Initiation Protocol (SIP)", RFC 3856, August 2004.
- [RFC3859] Peterson, J., "Common Profile for Presence (CPP)", RFC 3859, August 2004.
- [RFC3860] Peterson, J., "Common Profile for Instant Messaging (CPIM)", RFC 3860, August 2004.

- [RFC3966] Schulzrinne, H., "The tel URI for Telephone Numbers", RFC 3966, December 2004.
- [RFC5122] Saint-Andre, P., "Internationalized Resource Identifiers (IRIs) and Uniform Resource Identifiers (URIs) for the Extensible Messaging and Presence Protocol (XMPP)", RFC 5122, February 2008.
- [RFC5322] Resnick, P., Ed., "Internet Message Format", RFC 5322, October 2008.
- [RFC5626] Jennings, C., Mahy, R., and F. Audet, "Managing Client-Initiated Connections in the Session Initiation Protocol (SIP)", RFC 5626, October 2009.
- [RFC6121] Saint-Andre, P., "Extensible Messaging and Presence Protocol (XMPP): Instant Messaging and Presence", RFC 6121, March 2011.
- [RFC6665] Roach, A., "SIP-Specific Event Notification", RFC 6665, July 2012.
- [UTS39] The Unicode Consortium, "Unicode Technical Standard #39: Unicode Security Mechanisms", November 2013, <<http://unicode.org/reports/tr39/>>.
- [XEP-0106] Hildebrand, J. and P. Saint-Andre, "JID Escaping", XSF XEP 0106, June 2007, <<http://www.xmpp.org/extensions/xep-0106.html>>.
- [XEP-0131] Saint-Andre, P. and J. Hildebrand, "Stanza Headers and Internet Metadata", XSF XEP 0131, July 2006, <<http://xmpp.org/extensions/xep-0131.html>>.
- [XMPP-ADDRESS-FORMAT] Saint-Andre, P., "Extensible Messaging and Presence Protocol (XMPP): Address Format", Work in Progress, March 2014.

Appendix A. Acknowledgements

The authors wish to thank the following individuals for their feedback: Mary Barnes, Dave Cridland, Dave Crocker, Mike De Vries, Fabio Forno, Adrian Georgescu, Philipp Hancke, Saul Ibarra Corretge, Markus Isomaki, Olle Johansson, Paul Kyzivat, Salvatore Loreto, Daniel-Constantin Mierla, Tory Patnoe, and Robert Sparks.

Dan Romascanu reviewed the document on behalf of the General Area Review Team.

During IESG review, Stephen Farrell, Ted Lemon, Pete Resnick, and Sean Turner caught several issues that needed to be addressed.

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

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

Authors' Addresses

Peter Saint-Andre
&yet

EMail: ietf@stpeter.im

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
USA

EMail: jhildebr@cisco.com

