Network Working Group
Internet-Draft
Intended status: BCP
Expires: September 6, 2012
E. Ivov
Jitsi
E. Marocco, Ed.
Telecom Italia
March 5, 2012

Combined Use of the Session Initiation Protocol (SIP) and the eXtensible Messaging and Presence Protocol (CUSAX)
draft-ivov-xmpp-cusax-00

Abstract

This document describes current practices for combined use of the Session Initiation Protocol (SIP) and the eXtensible Messaging and Presence Protocol (XMPP). Such practices aim to provide a single fully featured real-time communication service by using complimenting subsets of features from each of the protocols. Typically such subsets would include telephony oriented from SIP and instant messaging and presence capabilities from XMPP. This specification does not define any new protocols or syntax for neither SIP nor XMPP. However, implementing it may require modifying or at least reconfiguring existing client and server-side software. Also, it is not the purpose of this document to make recommendations as to whether or not such combined us should be preferred to the mechanisms provided natively by each protocol like for example SIP's SIMPLE or XMPP's Jingle. It merely aims to provide guidance to those who are interested in such a combined use.

Status of this Memo

Copyright Notice

Table of Contents

1. Introduction

Historically SIP [RFC3261] and XMPP [RFC6120] have often been implemented and deployed with different purposes: from its very start SIP's primary goal has been to provide a means of conducting "Internet telephone calls". XMPP on the other hand, has from its Jabber days been mostly used for its instant messaging and presence capabilities.

For various reasons, these trends have continued through the years even after each of the protocols had been equipped to provide the features it was initially lacking.

Today, in the context of the SIMPLE working group, the IETF has defined a number of protocols and protocol extensions that not only allow for SIP to be used for regular instant messaging and presence but that also provide mechanisms for elaborated features such as multi-user chats, server-stored contact lists, file transfer and others.

Similarly, the XMPP community and the XMPP Standards Foundation have worked on defining a number of XMPP Extension Protocols (XEPs) that provide XMPP implementations with the means of establishing end-to-end sessions. These extensions are often jointly referred to as Jingle and their arguably most popular use case are audio and video calls.

Yet, despite these advances SIP remains the protocol of choice for telephony-like services, especially in enterprises where users are accustomed to features such as voice mail, call park, call queues, conference bridges and many others that are rarely (if at all) available in Jingle servers. XMPP implementations on the other hand, greatly outnumber and outperform those available for protocols recommended by SIMPLE, such as [MSRP] and [XCAP].

For these reasons in a number of cases, adopters may find themselves needing a set of features that are not offered by any single-protocol solution but that separately exist in SIP and XMPP products. The idea of seamlessly using both protocols together would hence often appeal to service providers.

Most often such combined use would employ SIP exclusively for audio, video and telephony services and it would rely on XMPP for anything else varying from chat, roster management and presence to exchanging files.

This document explains how the above could be achieved with a minimum amount of modifications on existing software while providing an

optimal user experience. It tries to cover points such as server discovery, determining a SIP AOR while using XMPP and an XMPP JID from incoming SIP requests. Most of the text here pertains to client behavior but it also recommends certain server-side configurations.

2. Terminology

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

3. Client Bootstrap

One of the main problems of using two distinct protocols when providing one service is how it affects usability. E-mail services for example have long been affected by the mixed use of SMTP on for outgoing mail and POP3 and IMAP for incoming, making it rather complicated for inexperienced users to configure a mail client and start using it with a new service. As a result mailing list services often need to provide configuration instructions for various mail clients. Client developers and communications device manufacturers on the other hand often ship with a number of wizards that allow to easily set up a new account for a number of popular e-mail services. While this may improve the situation to some extent, user experience is still clearly sub-optimal.

While it should be possible for CUSAX users to manually configure their separate SIP and XMPP accounts, it is RECOMMENDED that dual stack SIP/XMPP clients provide means of online provisioning. While the specifics of such mechanisms are not in the scope of this specification, they should make it possible for service providers to remotely configure the clients based on minimal user input (e.g. user id and password).

Given that many of the features that CUSAX would privilege in one protocol would also be available in the other, clients should make it possible for such features to be disabled for a specific account. Specifically it is RECOMMENDED that clients allow for audio/video calling features to be disabled for XMPP accounts. Additionally instant messaging and presence features MAY also be made optional for SIP accounts.

The main advantage of the above would be that clients would be able to continue to function properly and use the complete feature set of stand-alone SIP and XMPP accounts.

Once client bootstrap has completed, clients SHOULD log independently to the SIP and XMPP accounts that make up the CUSAX service and should maintain both these connections. In order to improve user experience, when reporting connection status clients may also wish to present the CUSAX XMPP connection as an "instant messaging" or a "chat" account. Similarly they could also depict the SIP CUSAX connection as a "Voice and Video" or a "Telephony" connection. The exact naming is of course entirely up to implementors. The point is that such presentation could help users better understand why they are being shown two different connections for a single service. It could even alleviate especially situations where one of these connections is disrupted while the other one is successfully maintained.

4. Operation

Once a CUSAX client has been provisioned/configured to connect to the corresponding SIP and XMPP services it would proceed by retrieving its XMPP roster. In order for CUSAX to function properly, XMPP service administrators should make sure that at least one of the [VCARD] "tel" fields for each contact is properly populated with a SIP URI or a phone number. There are no limitations as to the form of that number (e.g. it does not need to respect any equivalence with the XMPP JID). It SHOULD however be reachable through the SIP counterpart of this CUSAX service.

In order to make sure that the above is always respected, service maintainers MAY prevent clients (and hence users) from modifying the VCARD "tel" fields or they MAY apply some form of validation before recording changes.

When rendering the XMPP roaster CUSAX clients should make sure that users are presented with a "Call" option for each roster entry that has a properly set "tel" field even if calling has been disabled for that particular XMPP account. The usefulness of such a feature is not limited to CUSAX. After all, numbers are entered in VCARDs in order to be dialed and called. Hence, as long as an XMPP client is equipped with accounts that have calling features it may wish to present the user with the option of using these accounts to reach numbers from an XMPP VCARD. In order to improve usability, in cases where clients are provisioned with only a single telephony capable account they SHOULD do so immediately upon user request without asking for confirmation. This way CUSAX users whose only account with calling capabilities would often be the SIP part of their service would be having better user experience. If on the other hand, the CUSAX client is aware of multiple telephony-capable accounts, it SHOULD present the user with the choice of reaching the

phone number through any of them (including the source XMPP account where the VCARD was obtained) in order to guarantee proper operation for XMPP accounts that are not part of a CUSAX deployment.

The client should use XMPP for all other forms of communication with the contacts from its roster so it should and this should occur naturally given that they were retrieved through XMPP.

When receiving SIP calls, clients may wish to determine the identity of the caller and bind it to a roster entry so that users could revert to chatting or other forms of communication that require XMPP. To do so clients could search their roster for an entry whose VCARD has a "tel" field matching the originator of the call.

An alternate mechanism would be for CUSAX clients to add to their SIP invite requests a contact header containing their XMPP JID, but at this point we are not really sure if that's ' such a good idea. (After all Contact headers carry URIs and JIDs are not URIs).

5. Security Considerations

TBD

6. Acknowledgements

This draft is inspired by work from Markus Isomaki and Simo Veikkolainen.

7. References

7.1. Normative References

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.

7.2. Informative References

[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.

[RFC3264] Rosenberg, J. and H. Schulzrinne, "An Offer/Answer Model with Session Description Protocol (SDP)", RFC 3264, June 2002.

[RFC3489] Rosenberg, J., Weinberger, J., Huitema, C., and R. Mahy, "STUN - Simple Traversal of User Datagram Protocol (UDP) Through Network Address Translators (NATs)", RFC 3489, March 2003.

[RFC3711] Baugher, M., McGrew, D., Naslund, M., Carrara, E., and K. Norrman, "The Secure Real-time Transport Protocol (SRTP)", RFC 3711, March 2004.

[RFC4474] Peterson, J. and C. Jennings, "Enhancements for Authenticated Identity Management in the Session Initiation Protocol (SIP)", RFC 4474, August 2006.

[RFC4566] Handley, M., Jacobson, V., and C. Perkins, "SDP: Session Description Protocol", RFC 4566, July 2006.

[RFC4787] Audet, F. and C. Jennings, "Network Address Translation (NAT) Behavioral Requirements for Unicast UDP", BCP 127, RFC 4787, January 2007.

[RFC5245] Rosenberg, J., "Interactive Connectivity Establishment (ICE): A Protocol for Network Address Translator (NAT) Traversal for Offer/Answer Protocols", RFC 5245, April 2010.

[RFC5389] Rosenberg, J., Mahy, R., Matthews, P., and D. Wing, "Session Traversal Utilities for NAT (STUN)", RFC 5389, October 2008.

[RFC5751] Ramsdell, B. and S. Turner, "Secure/Multipurpose Internet Mail Extensions (S/MIME) Version 3.2 Message Specification", RFC 5751, January 2010.

[RFC5766] Mahy, R., Matthews, P., and J. Rosenberg, "Traversal Using Relays around NAT (TURN): Relay Extensions to Session Traversal Utilities for NAT (STUN)", RFC 5766, April 2010.

[RFC5853] Hautakorpi, J., Camarillo, G., Penfield, R., Hawrylyshen, A., and M. Bhatia, "Requirements from Session Initiation Protocol (SIP) Session Border Control (SBC) Deployments", RFC 5853, April 2010.

[RFC6120] Saint-Andre, P., "Extensible Messaging and Presence Protocol (XMPP): Core", RFC 6120, March 2011.

[RFC6189] Zimmermann, P., Johnston, A., and J. Callas, "ZRTP: Media Path Key Agreement for Unicast Secure RTP", RFC 6189, April 2011.

[XEP-0177]
Beda, J., Saint-Andre, P., Hildebrand, J., and S. Egan, "XEP-0177: Jingle Raw UDP Transport Method", XEP XEP-0177, December 2009.

Authors' Addresses

Emil Ivov
Jitsi
Strasbourg 67000
France

Email: emcho@jitsi.org

Enrico Marocco (editor)
Telecom Italia
Via G. Reiss Romoli, 274
Turin 10148
Italy

Email: enrico.marocco@telecomitalia.it