Network Working Group
Internet-Draft
Intended status: Informational
Expires: April 25, 2013

E. Ivov
Jitsi
E. Marocco
Telecom Italia
P. Saint-Andre
Cisco Systems, Inc.
October 22, 2012

Combined Use of the Session Initiation Protocol (SIP) and the Extensible Messaging and Presence Protocol (CUSAX) draft-ivov-xmpp-cusax-02

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 complementary subsets of features from each of the protocols. Typically such subsets would include telephony capabilities from SIP and instant messaging and presence capabilities from XMPP. This specification does not define any new protocols or syntax for either SIP or 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 use 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.

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 instant messaging and presence [RFC6121], as well as related services such as groupchat rooms [XEP-0045].

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.

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-based software. XMPP implementations, on the other hand, greatly outnumber and outperform those available for instant messaging and presence extensions developed in the SIMPLE WG, such as MSRP [RFC4975] and XCAP [RFC4825].

For these reasons, in a number of cases adopters have found 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 the combined use of SIP and XMPP ("CUSAX") would employ SIP exclusively for audio, video, and telephony services and rely on XMPP for anything else varying from chat, contact list management, and presence to whiteboarding and exchanging files.

This document explains how such hybrid offerings can be achieved with

a minimum of modifications to 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 determining an XMPP Jabber Identifier ("JID") from incoming SIP requests. Most of the text here pertains to client behavior but it also recommends certain server-side configurations.

Note that this document is focused on coexistence of SIP and XMPP functionality in end-user-oriented clients. By intent it does not define methods for protocol-level mapping between SIP and XMPP, as might be used within a server-side gateway between a SIP network and an XMPP network. A separate series of documents has been produced that defines such mappings.

2. Client Bootstrap

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

While it should be possible for CUSAX users to manually configure their separate SIP and XMPP accounts, dual-stack SIP/XMPP clients ought to provide means of online provisioning. While the specifics of such mechanisms are outside the scope of this specification, they should make it possible for a service provider to remotely configure the clients based on minimal user input (e.g., only a user ID and password).

Because many of the features that a CUSAX client 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. In particular, it is suggested that clients allow for audio/video calling features to be disabled for XMPP accounts. Additionally, instant messaging and presence features should 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 need to log in independently to the SIP and XMPP accounts that make up the CUSAX "service" and then 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 implementers. The point is that, in cases where SIP and XMPP are components of a service offered by a single provider, such presentation could help users better understand why they are being shown two different connections for what they perceive as a single service. It could alleviate especially situations where one of these connections is disrupted while the other one is successfully maintained.

3. 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 [RFC6350] "tel" fields for each contact is properly populated with a SIP URI or a phone number when an XMPP protocol for vCard storage (e.g., [XEP-0054] or [XEP-0292]) is used. 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). However, it ought to be reachable through the SIP aspect of this CUSAX service.

To ensure that the foregoing approach is always respected, service providers might consider (1) preventing clients (and hence users) from modifying the vCard "tel" fields or (2) applying some form of validation before recording changes. Of course such validation would be feasible mostly in cases where one single provider controls both the XMPP and the SIP service since such providers would "know" (e.g., based on use of a common user database for both services) what SIP AOR corresponds to a given XMPP user.

When rendering the XMPP roster 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 ought to 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 have a better user experience. If on the other hand, the CUSAX client is aware of multiple telephony-capable accounts, it ought to 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.

In addition to discovering phone numbers from vCards, clients may also check presence broadcasts and the appropriate Personal Eventing Protocol nodes as described in XEP-0152: Reachability Addresses [XEP-0152].

The client should use XMPP for all other forms of communication with the contacts from its roster, which will occur naturally because they were retrieved through XMPP and only voice/video features were disabled in the XMPP stack.

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 the XMPP URI corresponding to their JID as per [RFC5122].

4. Federation

An alternate mechanism would be for CUSAX clients to add to

5. Security Considerations

Use of the same user agent with two different accounts providing complementary features introduces the possibility of mismatches between the security profiles of those accounts or features. For example, the SIP aspect and XMPP aspect of the CUSAX service might offer different authentication options (e.g., digest authentication for SIP as specified in [RFC3261] and SCRAM authentication [RFC5802] for XMPP as specified in [RFC6120]). Similarly, a CUSAX client might

successfully negotiate Transport Layer Security (TLS) [RFC5246] when connecting to the XMPP aspect of the service but not when connecting to the SIP aspect. Such mismatches could introduce the possibility of downgrade attacks. User agent developers and service providers ought to ensure that such mismatches are avoided as much as possible.

Refer to the specifications for the relevant SIP and XMPP features for detailed security considerations applying to each "stack" in a CUSAX client.

6. Acknowledgements

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

7. 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.

[RFC4825] Rosenberg, J., "The Extensible Markup Language (XML) Configuration Access Protocol (XCAP)", RFC 4825, May 2007.

[RFC4975] Campbell, B., Mahy, R., and C. Jennings, "The Message Session Relay Protocol (MSRP)", RFC 4975, September 2007.

[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.

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

[RFC5246] Dierks, T. and E. Rescorla, "The Transport Layer Security (TLS) Protocol Version 1.2", RFC 5246, August 2008.

[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.

[RFC5802] Newman, C., Menon-Sen, A., Melnikov, A., and N. Williams, "Salted Challenge Response Authentication Mechanism (SCRAM) SASL and GSS-API Mechanisms", RFC 5802, July 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.

[RFC6121] Saint-Andre, P., "Extensible Messaging and Presence Protocol (XMPP): Instant Messaging and Presence", RFC 6121, March 2011.

[RFC6189] Zimmermann, P., Johnston, A., and J. Callas, "ZRTP: Media

Path Key Agreement for Unicast Secure RTP", RFC 6189, April 2011.

[RFC6350] Perreault, S., "vCard Format Specification", RFC 6350, August 2011.

[XEP-0045] Saint-Andre, P., "Multi-User Chat", XSF XEP 0045, February 2012.

[XEP-0054] Saint-Andre, P., "vcard-temp", XSF XEP 0054, July 2008.

[XEP-0152] Hildebrand, J. and P. Saint-Andre, "XEP-0152: Reachability Addresses", XEP XEP-0152, October 2008.

[XEP-0292] Saint-Andre, P. and S. Mizzi, "vCard4 Over XMPP", XSF XEP 0292, October 2011.

Authors' Addresses

Emil Ivov
Jitsi
Strasbourg 67000
France

Phone: +33-672-811-555
Email: emcho@jitsi.org

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

Email: enrico.marocco@telecomitalia.it

Peter Saint-Andre
Cisco Systems, Inc.
1899 Wynkoop Street, Suite 600
Denver, CO 80202
USA

Phone: +1-303-308-3282
Email: psaintan@cisco.com