

MMUSIC WG
Internet-Draft
Intended status: Informational
Expires: November 26, 2011

J. Elwell
A. Hutton
T. Stach
Siemens Enterprise
Communications
May 25, 2011

ICE Updated Offer Problematic
draft-elwell-mmusic-ice-updated-offer-01.txt

Abstract

Interactive Connectivity Establishment (ICE) requires an updated offer-answer cycle under some circumstances, to satisfy the needs of some devices on the signalling path. When used with SIP, this additional offer-answer cycle interacts with other things, such as fax and third party call control (3PCC). This document describes the problems and discusses possible remedies.

This work is being discussed on the mmusic@ietf.org mailing list.

Status of this Memo

This Internet-Draft is submitted in full conformance with the provisions of BCP 78 and BCP 79.

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at <http://datatracker.ietf.org/drafts/current/>.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

This Internet-Draft will expire on November 26, 2011.

Copyright Notice

Copyright (c) 2011 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. Fax calls	3
2.1. Problem statement	3
2.2. Possible remedies	5
2.2.1. Delay the ICE updated offer	5
2.2.2. Delay the fax updated offer	5
2.2.3. Use reliable provisional responses and pre-conditions	5
3. Third party call control (3PCC)	6
3.1. Problem statement	6
3.2. Possible remedies	7
3.2.1. Avoid 3PCC	7
3.2.2. Delay the updated offer	8
3.2.3. Delay ICE until UA2 answers	8
3.2.4. Issue an early 200 response to the INVITE request to UA2	8
3.2.5. Use reliable provisional responses	9
4. Do we really need the updated offer?	9
4.1. Types of devices that rely on the updated offer	9
4.2. Types of environment in which ICE is deployed	9
4.3. Relaxing the requirement	10
5. Conclusions	10
6. IANA considerations	11
7. Security considerations	11
8. Informative References	11
Authors' Addresses	12

1. Introduction

Interactive Connectivity Establishment (ICE) [RFC5245] specifies a mechanism for NAT traversal for multimedia sessions established using the Session Description Protocol (SDP) [RFC4566] offer-answer model [RFC3264]. It allows a pair of endpoints to exchange candidate IP addresses and ports, perform checks to see which pairs of candidates work, and agree which pairs to use for a given component of a given medium (e.g., RTP stream, RTCP stream). The mechanism can also be used for IPv6 transition, for determining whether to use IPv4 or IPv6. A particular application of ICE is with the Session Initiation Protocol (SIP) [RFC3261].

Connectivity checks are performed on the media path between candidate pairs. Based on the results of connectivity checks and certain rules, the two endpoints each determine which pair of candidates to use for a given component and can then start exchanging data (e.g., RTP packets) on the agreed path. Further exchanges on the signalling path (i.e., the path on which the offer-answer exchange is performed) are not necessary for the endpoints to agree which candidates to use.

However, certain SIP/SDP-aware devices on the signalling path need to know which candidates have been selected (e.g., to prioritize that traffic or to remove the resources for non-selected candidates). For this reason ICE mandates a further offer-answer exchange in some circumstances, i.e., an updated SDP offer followed by an updated SDP answer. In some situations with SIP, this updated offer-answer exchange can be problematic. This document examines these problems.

2. Fax calls

2.1. Problem statement

Except where dedicated fax devices are involved, fax calls typically start as audio. Detection of CNG tone (calling tone) from the initiating fax machine and CED (called) tone from the receiving fax machine initiates a switch to T.38, i.e., a switch from audio to image. Where the audio call uses a compressed codec (e.g., G.729), if one tone is detected there may first be a switch to G.711, for more reliable tone detection or in case the call turns out to be a non-fax modem call. Thus there can be:

a switch from a compressed codec to G.711; or

a switch from audio to image; or

both in sequence.

Switching codec or switching from audio to image requires an SDP offer-answer cycle. ICE also requires an updated offer-answer cycle where the selected candidates are not those in the m/c-lines of the original offer-answer. If the UA that detects the need to switch because of fax is also the controlling agent from the ICE perspective, updated offer-answer for fax can follow the updated offer-answer for ICE and probably won't lead to problems.

However, if the UA that detects the need to switch because of fax is not the controlling agent from the ICE perspective, there is a significant danger of the two re-INVITE or UPDATE [RFC3311] requests colliding, resulting in a 491 response to each. According to [RFC3261] and [RFC3311], one UA (the one that owns the Call-ID) backs off for between 2.1 and 4 seconds, and the other UA backs off for between 0 and 2 seconds, before trying again. This can result in a delay of up to 4 seconds before the switch to fax, long enough in practice to cause fax calls to fail. It can also result in a delay of up to 4 seconds before the post-ICE updated offer-answer. SIP/SDP-aware devices that need the post-ICE updated offer-answer might not permit the flow of RTP packets throughout this period, which might also lead to failure of the fax call. An example flow is shown below:

```

UA1 (Call-ID owner)                                UA2 (fax gateway)
-----INVITE / SDP offer audio----->
<-----183 / SDP answer audio-----
<=====ICE negotiation=====>
<-----200 / SDP answer audio -----
-----ACK----->
<=====RTP=====>
(ICE requires updated offer)
-----UPDATE / SDP offer audio---
<-----UPDATE / SDP offer image-----
                                     \
                                     (Fax detected)
                                     /
                                     ----->
-----491-----
                                     \
                                     -----491-----
                                     /
                                     ----->
                                     (back-off 0-2s)
<-----UPDATE / SDP offer image-----
-----200 / SDP answer image----->
(back-off 2.1-4s)
----UPDATE / SDP offer image (selected candidates)----->

```

<----200 / SDP answer image (selected candidates)-----

In this example UA1 is the ICE controlling agent and issues an updated offer on completion of ICE, and UA2 is a fax gateway that detects fax and attempts to change to image. UPDATE is supported by both and used for the updated offers. UA1 owns the Call-ID and has the longer back-off. In this example, the switch to image will probably be accomplished fast enough (back-off does not exceed 2 seconds), but the post-ICE updated offer can be delayed up to 4 seconds, perhaps leading to undesirable behaviour by SIP/SDP-aware devices on the signalling path, which might disrupt the flow of RTP and cause the fax call to fail.

Of course, collision of UPDATE or re-INVITE requests will not always occur - it is matter of timing. However, the probability of collision is not insignificant and, if that occurs, the probability of the fax call being adversely affected to the extent that it fails is not insignificant.

2.2. Possible remedies

2.2.1. Delay the ICE updated offer

UA1, as the ICE controlling agent, will be unaware that UA2 will detect fax. Therefore any delay in sending the ICE updated offer will need to apply to all calls and will need to be long enough to allow for differing amounts of time for UA2 to detect fax (perhaps several seconds). The question then is whether this would be long enough to introduce a risk of undesirable behaviour by SIP/SDP-aware devices on the signalling path, which could impact all calls, not just fax calls.

2.2.2. Delay the fax updated offer

UA2 will know that ICE has been used, and therefore can expect an updated offer from UA1, the ICE controlling agent. Normally this should arrive quite quickly (e.g., well under 100 ms), although it depends on the number of SIP intermediaries on the path and whether any retransmissions are needed because of packet loss. Therefore a delay of a 100 ms., say, would probably not impact the fax call and would help avoid collisions but would not be a guarantee.

2.2.3. Use reliable provisional responses and pre-conditions

By using a reliable 183 in accordance with [RFC3262] to send the SDP answer, when ICE completes the updated offer can be sent in an UPDATE request, rather than waiting for the 200 response to the INVITE request and then sending the updated offer. However, the fax machine

might auto-answer and send the 200 response to the INVITE request as soon as ICE procedures complete, so the updated offer might collide with the 200 response, again leading to further signalling delays before things are resolved. This in turn could be avoided by using pre-conditions [RFC3312] to delay answering of the call until the updated offer has occurred.

This might work, although it is unclear how pre-conditions are intended to interact with ICE, i.e., whether ICE procedures can continue without waiting for pre-conditions to be satisfied. Perhaps an extension to pre-conditions would be required. Also this might introduce further adverse interactions with SIP/SDP-aware devices on the signalling path.

Even if it could be made to work, this approach would require the entities involved to support [RFC3262] and [RFC3312]. [RFC3262] is known to be rather complicated to implement (hence the reason the ICE mechanism was specifically designed to allow SDP answer to be sent in an unreliable provisional response (ICE provides acknowledgement on the media path, rather than requiring the use of PRACK). Pre-conditions are a further complication and not widely implemented. Therefore ICE implementations should not be expected to support [RFC3262] and [RFC3312].

3. Third party call control (3PCC)

3.1. Problem statement

3PCC [RFC3725] is a common technique used with SIP where calls are controlled from an application at a SIP B2BUA. In particular, calls can be established by 3PCC, whereby the application first makes a call to the first party (typically the device of a user requesting the call) and then makes a second call to the second party, the two calls being joined together such that media flow directly between the two devices. A typical implementation is in accordance with Flow I in [RFC3725]: the controlling B2BUA sends an offerless INVITE request to UA1, which alerts the first user. When the user answers, UA1 sends an offer in a 200 response to the INVITE request, and this offer is used by the B2BUA in a second INVITE request, this time to UA2.

The problem with using ICE with 3PCC is that 3PCC signalling does not permit the updated offer-answer for ICE to occur in a timely manner. UA2 will often take some time (seconds or tens of seconds) before sending the 200 response to its INVITE request. Yet if UA2 has already sent an SDP answer (e.g., in a 183 response), ICE can complete on the media paths and UA1, as the ICE controlling agent,

can attempt an updated offer. This updated offer cannot be forwarded to UA2 until the INVITE transaction on that leg of the call has completed.

More specifically, consider the following example flow:

```

UA1 (Call-ID owner)      B2BUA      UA2
<----INVITE (no SDP)-----
-----200 / SDP offer----->      ----INVITE / SDP offer---->
<----ACK / SDP answer-----      <-----183 / SDP answer-----
<=====ICE negotiation=====>
(ICE requires updated offer)
-----UPDATE / SDP offer----> What next?

```

In this case, UA2 sends a 183 provisional response to its INVITE request. This contains an SDP answer, which is passed to UA1 through the ACK request. Thus UA1 and UA2 are able to conduct ICE negotiation on the media paths. UA2 will probably not alert its user until ICE negotiation is complete, but anyway, there will often be a significant delay before the user answers and UA2 sends a 200 response to its INVITE request. Meanwhile, UA1, as the ICE controlling agent, attempts to send an updated offer. In this case it chooses to use an UPDATE request, but similar considerations apply if it uses a re-INVITE request. The B2BUA cannot pass that request on until the INVITE transaction with UA2 has completed. Either the UPDATE request has to be delayed somehow or rejected, in either case leading to the possibility of undesirable behaviour by SIP/SDP-aware devices that require a timely updated offer. For example, UA2 might be transmitting early media, which might fail to be passed through correctly, or clipping might occur when the user answers.

It should be noted that the issue of sending an updated offer in a 3PCC situation before UA2 has answered is not solely an ICE issue. However, ICE substantially increases the need for such an updated offer.

3.2. Possible remedies

3.2.1. Avoid 3PCC

There are alternatives to this form of 3PCC. For example, UA1 could be instructed to issue a conventional INVITE request by sending a SIP REFER request to UA1, or by some non-SIP means. However, using a REFER request is not an option for some types of UA, for example PSTN gateways. If user 1 is a PSTN user, it is necessary to make a PSTN call to the user, and this can be achieved by sending an INVITE request to UA1, but not by sending a REFER request to UA1. Non-SIP means are either not standardized or little deployed.

A particular example of an application that uses 3PCC is one where the user uses a web page to make the call, having selected in advance the device he/she wishes to use to make the call. The application causes the B2BUA to send an INVITE request to that selected device, followed by an INVITE request to the called destination. If the selected device is, for example, a cellular device reachable via PSTN, that initial INVITE request will be sent to a PSTN gateway.

Because of the difficulties supporting such applications by other means, 3PCC is a commonly deployed technique. It is not possible to scrap 3PCC in order to introduce ICE.

3.2.2. Delay the updated offer

UA1 will typically not be aware of the state of the INVITE transaction to UA2, and will issue the updated offer in an UPDATE or re-INVITE request without knowing whether the B2BUA can pass it on. Therefore the onus is on the B2BUA to handle the situation when it receives the UPDATE or re-INVITE request. As a non-INVITE transaction, an UPDATE request has a relatively short timeout, but one possibility would be for the B2BUA to reject it with a 500 response and a Retry-After header field, relying on UA1 to try again later. In the case of re-INVITE, the B2BUA could delay forwarding the request to UA2 until the original transaction is complete. However, in either case, SIP/SDP-aware devices between the B2BUA and UA2 will not see the updated offer in a timely manner, and therefore might take action that prevents the correct handling of early media or clips media for a short time after the call is answered.

3.2.3. Delay ICE until UA2 answers

UA2 could delay ICE until UA2 answers, which means UA2 would not need to send SDP answer in a provisional response but could wait for the 200 response. This would mean the user would answer and experience a delay (clipping) before ICE completes and media start to flow. Since UA2 would not be aware of the 3PCC situation, this would impact all calls to UA2, not just those that use 3PCC.

3.2.4. Issue an early 200 response to the INVITE request to UA2

UA2 could issue a 200 response instead of a 183 response, even though the user has not yet been alerted and answered. This would be different from normal practice and might adversely impact behaviour at other SIP entities, e.g., charging, logging, forking, call forwarding. Again UA 2 would not be aware of the 3PCC situation, so this would impact all calls to UA2, not just those that use 3PCC.

3.2.5. Use reliable provisional responses

If UA2 and the B2BUA support reliable provisional responses [RFC3262], UA2 can send the 183 response with SDP answer reliably (resulting in a PRACK transaction), and then the B2BUA can send an UPDATE request with the updated offer without waiting for the INVITE transaction to complete. This would seem to work, except that it requires the entities involved to support [RFC3262], which, as explained in section Section 2.2.3, is undesirable. In particular, UA2, which is the "innocent party" in 3PCC, should not be expected to provide special functionality just to make 3PCC work. Furthermore, a B2BUA performing 3PCC would not be aware of ICE and hence the need to support [RFC3262].

4. Do we really need the updated offer?

4.1. Types of devices that rely on the updated offer

Devices on the signalling path that rely on the updated offer are SIP/SDP-aware devices (e.g., policy servers) that perform admission control or resource reservation based on SDP, without modifying the SDP as the signalling messages are forwarded. Devices that modify the SDP (e.g., Session Border Controllers) generally terminate ICE, so are not an issue.

One type of behaviour that relies on the updated offer is a device that is ICE-aware and reserves resources for all the ICE candidate-pairs. Such a device would need to know which candidates have been selected, so that unwanted resources can be freed.

A second type of behaviour that relies on the updated offer is a device that is not ICE-aware and admits traffic on ports identified in the m/c lines of the SDP offer. Such a device is assumed to let a moderate amount of traffic through on other ports, so would not prevent STUN connectivity checks, but would prevent a sustained transmission of RTP. The updated offer/answer would allow such a device to admit sustained traffic on the ports that have been negotiated using ICE.

4.2. Types of environment in which ICE is deployed

ICE has seen a certain amount of deployment. ICE is not solely for use in a SIP/SDP environment, and some of those deployments are in non-SIP/SDP environments (e.g., Jingle). ICE deployment with SIP is relatively sparse in some types of environment, the fundamental reason being that NAT traversal is frequently accomplished by Session Border Controllers. This is true, for example, for most enterprise

deployments of SIP. Where such environments are migrated to IPv6, often SBCs are used at the border between IPv4 and IPv6 networks and therefore ICE is not needed for negotiating the IP version. Frequently the types of signalling path device that would require an ICE updated offer are deployed in this sort of environment, where ICE is currently not needed, not deployed and unlikely to be deployed.

On the other hand, the use of SIP across the public Internet without the use of SBCs does require ICE. In such deployments, however, it is unlikely there will be devices on the signalling path that would need an updated offer.

So basically there are environments where SIP is used and ICE is not deployed and not needed. There are also environments where SIP is used and ICE is needed, and to some extent deployed, but such environments generally do not require the ICE updated offer. Some such deployments may not be concerned with fax or with 3PCC, and therefore implementation of the updated offer might not be an issue, although it is believed that some implementations do not support the updated offer and can still operate in their target environments.

In the future, ICE is likely to be required in more environments. The present SBC-based approach in enterprise environments, for example, might not be the most appropriate for use in cloud-based deployments, where it is unrealistic to have media following the signalling path. ICE could be used in such environments, but the presence of signalling path devices that need the updated offer seems unlikely.

4.3. Relaxing the requirement

These considerations bring into question the mandatory requirement in [RFC5245] for an updated SDP offer under some circumstances. This could be relaxed such that it can be omitted in environments where it is not needed.

5. Conclusions

This document illustrates two common use cases where the introduction of ICE can lead to problems with the updated offer/answer cycle that ICE requires in certain circumstances. In the first case (fax), the problem arises at the two endpoints that are trying to accomplish ICE. In the second case (3PCC), the problem arises because of a particular B2BUA behaviour, yet the B2BUA is not involved in ICE, should not need to know anything about ICE, and should not need to implement any extensions to SIP or SDP in order for ICE to work between UAs. In both cases there are work-arounds, but these

introduce dependencies that contrive to reduce the chances of successful interoperability.

The need, in some circumstances, to conduct an updated offer/answer cycle on conclusion of ICE is common to both problems. This need arises not from ICE itself, but from the certain types of SIP/SDP-aware devices on the signalling path whose normal functioning is impacted when endpoints use ICE, unless they have been upgraded to cope with the effects of ICE.

The two use cases illustrated might not be the only cases where the ICE updated offer is problematic. As more complex multimedia situations arise, involving mid-call (and in particular early-in-the-call) offer-answer cycles for changing media, changing security, etc., the more likely it is that the additional ICE update offer-answer cycle will intrude in an unhelpful way.

According to discussions in section Section 4, it seems to be the case that the updated offer is needed, in practice, in very few environments, and therefore consideration should be given to relaxing the requirement in [RFC5245].

6. IANA considerations

This document requires no IANA actions.

7. Security considerations

This document does not introduce any new security considerations.

8. 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.
- [RFC3262] Rosenberg, J. and H. Schulzrinne, "Reliability of Provisional Responses in Session Initiation Protocol (SIP)", RFC 3262, June 2002.
- [RFC3264] Rosenberg, J. and H. Schulzrinne, "An Offer/Answer Model with Session Description Protocol (SDP)", RFC 3264, June 2002.

- [RFC3311] Rosenberg, J., "The Session Initiation Protocol (SIP) UPDATE Method", RFC 3311, October 2002.
- [RFC3312] Camarillo, G., Marshall, W., and J. Rosenberg, "Integration of Resource Management and Session Initiation Protocol (SIP)", RFC 3312, October 2002.
- [RFC3725] Rosenberg, J., Peterson, J., Schulzrinne, H., and G. Camarillo, "Best Current Practices for Third Party Call Control (3pcc) in the Session Initiation Protocol (SIP)", BCP 85, RFC 3725, April 2004.
- [RFC4566] Handley, M., Jacobson, V., and C. Perkins, "SDP: Session Description Protocol", RFC 4566, July 2006.
- [RFC5245] Rosenberg, J., "Interactive Connectivity Establishment (ICE): A Protocol for Network Address Translator (NAT) Traversal for Offer/Answer Protocols", RFC 5245, April 2010.

Authors' Addresses

John Elwell
Siemens Enterprise Communications

Phone: +44 1908 817801
Email: john.elwell@siemens-enterprise.com

Andy Hutton
Siemens Enterprise Communications

Phone: +44 1908 817920
Email: andrew.hutton@siemens-enterprise.com

Thomas Stach
Siemens Enterprise Communications

Phone: +43 5 7008 4377
Email: thomas.stach@siemens-enterprise.com

