



XMPP

XEP-0176: Jingle ICE-UDP Transport Method

Joe Beda

<mailto:jbada@google.com>

<xmpp:jbada@google.com>

Scott Ludwig

<mailto:scottlu@google.com>

<xmpp:scottlu@google.com>

Peter Saint-Andre

<mailto:xsf@stpeter.im>

<xmpp:peter@jabber.org>

<http://stpeter.im/>

Joe Hildebrand

<mailto:jhildebr@cisco.com>

<xmpp:hildjj@jabber.org>

Sean Egan

<mailto:seanegan@google.com>

<xmpp:seanegan@google.com>

Robert McQueen

<mailto:robert.mcqueen@collabora.co.uk>

<xmpp:robert.mcqueen@collabora.co.uk>

2009-06-10

Version 1.0

Status	Type	Short Name
Draft	Standards Track	jingle-ice-udp

This specification defines a Jingle transport method that results in sending media data using raw datagram associations via the User Datagram Protocol (UDP). This transport method is negotiated via the Interactive Connectivity Establishment (ICE) methodology, which provides robust NAT traversal for media traffic.

Legal

Copyright

This XMPP Extension Protocol is copyright © 1999 – 2018 by the [XMPP Standards Foundation](#) (XSF).

Permissions

Permission is hereby granted, free of charge, to any person obtaining a copy of this specification (the "Specification"), to make use of the Specification without restriction, including without limitation the rights to implement the Specification in a software program, deploy the Specification in a network service, and copy, modify, merge, publish, translate, distribute, sublicense, or sell copies of the Specification, and to permit persons to whom the Specification is furnished to do so, subject to the condition that the foregoing copyright notice and this permission notice shall be included in all copies or substantial portions of the Specification. Unless separate permission is granted, modified works that are redistributed shall not contain misleading information regarding the authors, title, number, or publisher of the Specification, and shall not claim endorsement of the modified works by the authors, any organization or project to which the authors belong, or the XMPP Standards Foundation.

Warranty

NOTE WELL: This Specification is provided on an "AS IS" BASIS, WITHOUT WARRANTIES OR CONDITIONS OF ANY KIND, express or implied, including, without limitation, any warranties or conditions of TITLE, NON-INFRINGEMENT, MERCHANTABILITY, or FITNESS FOR A PARTICULAR PURPOSE.

Liability

In no event and under no legal theory, whether in tort (including negligence), contract, or otherwise, unless required by applicable law (such as deliberate and grossly negligent acts) or agreed to in writing, shall the XMPP Standards Foundation or any author of this Specification be liable for damages, including any direct, indirect, special, incidental, or consequential damages of any character arising from, out of, or in connection with the Specification or the implementation, deployment, or other use of the Specification (including but not limited to damages for loss of goodwill, work stoppage, computer failure or malfunction, or any and all other commercial damages or losses), even if the XMPP Standards Foundation or such author has been advised of the possibility of such damages.

Conformance

This XMPP Extension Protocol has been contributed in full conformance with the XSF's Intellectual Property Rights Policy (a copy of which can be found at <https://xmpp.org/about/xsf/ipr-policy>) or obtained by writing to XMPP Standards Foundation, P.O. Box 787, Parker, CO 80134 USA).

Contents

1	Introduction	1
2	Glossary	1
3	Requirements	2
4	Jingle Conformance	2
5	Protocol Description	2
5.1	Overall Flow	2
5.2	Session Initiation	3
5.3	Syntax	4
5.4	Response	7
5.5	Candidate Negotiation	8
5.6	Connectivity Checks	9
5.7	Acceptance of Successful Candidate	11
5.8	Negotiating a New Candidate	12
5.9	ICE Restarts	13
6	Fallback to Raw UDP	14
7	Determining Support	19
7.1	ICE Support	19
7.2	SDP Offer / Answer Support	19
8	Implementation Notes	20
9	Deployment Notes	20
10	Security Considerations	21
10.1	Sharing IP Addresses	21
10.2	Encryption of Media	21
11	IANA Considerations	22
12	XMPP Registrar Considerations	22
12.1	Protocol Namespaces	22
12.2	Protocol Versioning	22
12.3	Service Discovery Features	22
12.4	Jingle Transport Methods	23
13	XML Schema	23
14	Acknowledgements	25

1 Introduction

[Jingle \(XEP-0166\)](#) ¹ defines a framework for negotiating and managing out-of-band data sessions over XMPP. In order to provide a flexible framework, the base Jingle specification defines neither data transport methods nor application formats, leaving that up to separate specifications.

The current document defines a transport method for establishing and managing data exchanges between XMPP entities over the User Datagram Protocol (see [RFC 768](#) ²), using the ICE methodology developed within the IETF and specified in [RFC 5245](#) ³ (hereafter referred to as ICE-CORE). Use of this "ice-udp" method results in a datagram transport suitable for media applications where some packet loss is tolerable (e.g., audio and video).

Note: ICE-CORE has been approved for publication as an RFC but has not yet been published as an RFC. While every effort has been made to keep this document synchronized with ICE-CORE, the interested reader is referred to ICE-CORE for a detailed description of the ICE methodology.

The process for ICE negotiation is largely the same in Jingle as it is in ICE. There are several differences:

- Instead of using the Session Initiation Protocol (SIP) as the signalling channel, Jingle uses XMPP as the signalling channel.
- Syntax from the Session Description Protocol (see [RFC 4566](#) ⁴) is mapped to an XML syntax suitable for sending over the XMPP signalling channel.
- In Jingle, lists of "preferred" candidates are typically sent in the Jingle session-initiate and session-accept messages, in a way that is consistent with the SDP offer / answer model described in [RFC 3264](#) ⁵ and the process described in ICE-CORE.
- Candidates can also be sent in separate transport-info messages either before sending the session-accept message (to expedite negotiation) or after media begins to flow (to find modify existing candidates, find superior candidates, or adjust to changing network conditions).

2 Glossary

The reader is referred to ICE-CORE for a description of various terms used in the context of ICE. Those terms are not reproduced here.

¹XEP-0166: Jingle <<https://xmpp.org/extensions/xep-0166.html>>.

²RFC 768: User Datagram Protocol <<http://tools.ietf.org/html/rfc0768>>.

³RFC 5245: Interactive Connectivity Establishment (ICE) <<http://tools.ietf.org/html/rfc5245>>.

⁴RFC 4566: SDP: Session Description Protocol <<http://tools.ietf.org/html/rfc4566>>.

⁵RFC 3264: An Offer/Answer Model with the Session Description Protocol (SDP) <<http://tools.ietf.org/html/rfc3264>>.

3 Requirements

The Jingle transport method defined herein is designed to meet the following requirements:

1. Make it possible to establish and manage out-of-band connections between two XMPP entities, even if they are behind Network Address Translators (NATs) or firewalls.
2. Enable use of UDP as the transport protocol.
3. Make it relatively easy to implement support in standard Jabber/XMPP clients.
4. Where communication with non-XMPP entities is needed, push as much complexity as possible onto server-side gateways between the XMPP network and the non-XMPP network.

4 Jingle Conformance


In accordance with Section 10 of XEP-0166, this document specifies the following information related to the Jingle ice-udp transport method:

1. The transport negotiation process is defined in the [Protocol Description](#) section of this document.
2. The semantics of the <transport/> element are defined in the ICE Negotiation section of this document.
3. Successful negotiation of the ice-udp method results in use of a datagram transport that is suitable for applications where some packet loss is tolerable, such as audio and video.
4. If multiple components are to be communicated by the application type that uses the transport, the transport shall support those components and assign identifiers for them as described in the specification that defines the application type.

5 Protocol Description

5.1 Overall Flow

The overall protocol flow for negotiation of the Jingle ICE-UDP Transport Method is as follows (note: many of these events happen simultaneously, not in sequence).



Note: The examples in this document follow the scenario described in Section 17 of ICE-CORE, except that we substitute the Shakespearean characters "Romeo" and "Juliet" for the generic entities "L" and "R".

5.2 Session Initiation

In order for the initiator in a Jingle exchange to start the negotiation, it sends a Jingle "session-initiate" stanza that includes at least one content type, as described in XEP-0166. If the initiator wishes to negotiate the ice-udp transport method for an application format, it MUST include a <transport/> child element qualified by the 'urn:xmpp:jingle:transport:ice-udp:1' namespace (see [Namespace Versioning](#) regarding the possibility of incrementing the version number). This element SHOULD in turn contain one <candidate/> element for each of the initiator's higher-priority transport candidates as determined in accordance with the ICE methodology, but MAY instead be empty (with each candidate to be sent as the payload of a transport-info message).

Listing 1: Initiation

```
<iq from='romeo@montague.lit/orchard'
  id='ixt174g9'
  to='juliet@capulet.lit/balcony'
  type='set'>
  <jingle xmlns='urn:xmpp:jingle:1'
    action='session-initiate'
```

```

        initiator='romeo@montague.lit/orchard'
        sid='a73sjjvkl37jfea'>
<content creator='initiator' name='this-is-the-audio-content'>
  <description xmlns='urn:xmpp:jingle:apps:rtp:1' media='audio'>
    <payload-type id='96' name='speex' clockrate='16000' />
    <payload-type id='97' name='speex' clockrate='8000' />
    <payload-type id='18' name='G729' />
    <payload-type id='0' name='PCMU' />
    <payload-type id='103' name='L16' clockrate='16000' channels='
      2' />
    <payload-type id='98' name='x-ISAC' clockrate='8000' />
  </description>
  <transport xmlns='urn:xmpp:jingle:transports:ice-udp:1'
    pwd='asd88fgpdd777uzjYhagZg'
    ufrag='8hhy'>
    <candidate component='1'
      foundation='1'
      generation='0'
      id='el0747fg11'
      ip='10.0.1.1'
      network='1'
      port='8998'
      priority='2130706431'
      protocol='udp'
      type='host' />
    <candidate component='1'
      foundation='2'
      generation='0'
      id='y3s2b30v3r'
      ip='192.0.2.3'
      network='1'
      port='45664'
      priority='1694498815'
      protocol='udp'
      rel-addr='10.0.1.1'
      rel-port='8998'
      type='srflx' />
  </transport>
</content>
</jingle>
</iq>

```

5.3 Syntax

The <transport/> element's 'pwd' and 'ufrag' attributes MUST be included whenever sending one or more candidates to the other party, e.g. in a session-initiate, session-accept, transport-info, content-add, or transport-replace message. The values for these attributes are separately generated for both the initiator and the responder, in accordance with ICE-CORE

and as shown in the examples. The attributes of the <transport/> element are as follows.

Name	Description	SDP Syntax	Example
pwd	A Password as defined in ICE-CORE.	a=ice-pwd line	asd88fgpdd777uzjYhagZg
ufrag	A User Fragment as defined in ICE-CORE.	a=ice-ufrag line	8hhy

The attributes of the <candidate/> element are as follows.

Name	Description	SDP Syntax	Example
component	A Component ID as defined in ICE-CORE.	Component ID value in a=candidate line	1
foundation	A Foundation as defined in ICE-CORE.	Foundation value in a=candidate line	1
generation	An index, starting at 0, that enables the parties to keep track of updates to the candidate throughout the life of the session. For details, see the ICE Restarts section of this document.	extended name/value pair in a=candidate line	0
id	A unique identifier for the candidate.	N/A	el0747fg11
ip	The Internet Protocol (IP) address for the candidate transport mechanism; this can be either an IPv4 address or an IPv6 address.	IP Address value in a=candidate line	192.0.2.3
network	An index, starting at 0, referencing which network this candidate is on for a given peer (used for diagnostic purposes if the calling hardware has more than one Network Interface Card).	N/A	0
port	The port at the candidate IP address.	Port value in a=candidate line	45664

Name	Description	SDP Syntax	Example
priority	A Priority as defined in ICE-CORE. In accordance with the rules specified in Section 4.1.1 of ICE-CORE, the priority values shown in the examples within this document have been calculated as follows. The "type preference" for host candidates is stipulated to be "126" and for server reflexive candidates "100". The "local preference" for network 0 is stipulated to be "4096", for network 1 "2048", and for network 2 "1024".	Priority value in a=candidate line	2130706431
protocol	The protocol to be used. The only value defined by this specification is "udp". Future specifications might define other values such as "tcp".	Transport protocol field in a=candidate line	udp
rel-addr	A related address as defined in ICE-CORE.	raddr value in a=candidate line	10.0.1.1
rel-port	A related port as defined in ICE-CORE.	rport value in a=candidate line	8998
type	A Candidate Type as defined in ICE-CORE. The allowable values are "host" for host candidates, "prflx" for peer reflexive candidates, "relay" for relayed candidates, and "srflx" for server reflexive candidates.	Typ field in a=candidate line	srflx

5.4 Response

As described in XEP-0166, to acknowledge receipt of the session initiation request, the responder immediately returns an IQ-result.

Listing 2: Responder acknowledges receipt of session-initiate request

```
<iq from='juliet@capulet.lit/balcony'
  id='ixt174g9'
  to='romeo@montague.lit/orchard'
  type='result'/>
```

Depending on the application type, a user agent controlled by a human user might need to wait for the user to affirm a desire to proceed with the session before continuing. When the user agent has received such affirmation (or if the user agent can automatically proceed for any reason, e.g. because no human intervention is expected or because a human user has configured the user agent to automatically accept sessions with a given entity), it returns a Jingle session-accept message. This message MUST contain a <transport/> element qualified by the 'urn:xmpp:jingle:transports:ice-udp:1' namespace, which SHOULD in turn contain one <candidate/> element for each ICE-UDP candidate generated by or known to the responder, but MAY instead be empty (with each candidate to be sent as the payload of a transport-info message).

Note: See the [Security Considerations](#) section of this document regarding the exposure of IP addresses by the responder's client.

Listing 3: Responder accepts the session request

```
<iq from='juliet@capulet.lit/balcony'
  id='rw782g55'
  to='romeo@montague.lit/orchard'
  type='set'>
  <jingle xmlns='urn:xmpp:jingle:1'
    action='session-accept'
    initiator='romeo@montague.lit/orchard'
    responder='juliet@capulet.lit/balcony'
    sid='a73sjjvkl37jfea'>
    <content creator='initiator' name='this-is-the-audio-content'>
      <description xmlns='urn:xmpp:jingle:apps:rtp:1' media='audio'>
        <payload-type id='97' name='speex' clockrate='8000' />
        <payload-type id='18' name='G729' />
      </description>
      <transport xmlns='urn:xmpp:jingle:transports:ice-udp:1'
        pwd='YH75Fviy6338Vbrhr1p8Yh'
        ufrag='9uB6'>
        <candidate component='1'
          foundation='1'
          generation='0'
          id='or2ii2syr1'
```

```
        ip='192.0.2.1'  
        network='0'  
        port='3478'  
        priority='2130706431'  
        protocol='udp'  
        type='host' />  
    </transport>  
  </content>  
</jingle>  
</iq>
```

5.5 Candidate Negotiation

The initiator and responder negotiate connectivity over ICE by exchanging XML-formatted transport candidates for the channel. This negotiation proceeds immediately in order to maximize the possibility that connectivity can be established (and therefore media can be exchanged) as quickly as possible. In order to expedite session establishment, the initiator SHOULD include transport candidates in its session-initiate message but MAY also send additional transport candidates as soon as it learns of them, even before receiving acknowledgement of the session-initiate message (i.e., the initiator MUST consider the session to be live as soon as it sends the session-initiate message).⁶

The first step in negotiating connectivity is for each party to send transport candidates to the other party.⁸ These candidates SHOULD be gathered by following the procedure specified in Section 4.1.1 of ICE-CORE (typically by communicating with a standalone STUN server in order to discover the client's public IP address and port) and prioritized by following the procedure specified in Section 4.1.2 of ICE-CORE.

Each candidate shall be sent as a <candidate/> child of a <transport/> element qualified by the 'urn:xmpp:jingle:transports:ice-udp:1' namespace. The <transport/> element is sent via a Jingle message of type session-initiate, session-accept, or transport-info.

Either party MAY include multiple <candidate/> elements in one <transport/> element, especially in the session-initiate and session-accept messages sent at the beginning of the session negotiation. Including multiple candidates in the session-initiate and session-accept messages can help to ensure interoperability with entities that implement the SDP offer/answer model described in RFC 3264; in particular, an entity SHOULD include multiple candidates in its session-initiate or session-accept message if the other party advertises support for the "urn:ietf:rfc:3264" service discovery feature as described in the [SDP Offer / Answer Support](#) section of this document. However, including one candidate per subsequent transport-info message typically results in a faster negotiation because the candidates most likely to succeed are sent first (in the session-info and session-accept messages) and it is not necessary to gather all candidates before beginning to send any candidates; furthermore, because certain

⁶Given in-order delivery as mandated by [XMPP Core](#)⁷, the responder will receive such transport-info messages after receiving the session-initiate message; if not, it is appropriate for the responder to return <unknown-session/> errors since according to its state machine the session does not exist.

⁸The fact that both parties send candidates means that Jingle requires each party to be a full implementation of ICE, not a lite implementation as specified in ICE-CORE.

candidates can be more "expensive" in terms of bandwidth or processing power, either party might not want to advertise the existence of such candidates unless it is necessary to do so after other candidates have failed.

If the party that receives a candidate in a Jingle message can successfully process a given candidate or set of candidates, it returns an IQ-result (if not, for example because the candidate data is improperly formatted, it returns an IQ-error). At this point, the receiving entity is only indicating receipt of the candidate or set of candidates, not telling the other party that the candidate will be used.

The initiator can keep sending candidates (without stopping to receive an acknowledgement of receipt from the responder for each candidate) until it has exhausted its supply of possible or desirable transport candidates. The responder can also keep sending potential candidates, which the initiator will acknowledge.

5.6 Connectivity Checks

As the initiator and responder receive candidates, they probe the candidates for connectivity. In performing these connectivity checks, each party SHOULD follow the procedure specified in Section 7 of ICE-CORE. The following business rules apply:

1. Each party sends a STUN Binding Request (see [RFC 5389](#)⁹) from each local candidate it generated to each remote candidate it received.
2. In accordance with ICE-CORE, the STUN Binding Requests MUST include the PRIORITY attribute (computed according to Section 7.1.1.1. of ICE-CORE).
3. For the purposes of the Jingle ICE-UDP Transport Method, both parties are full ICE implementations and therefore the controlling role MUST be assumed by the initiator and the controlled role MUST be assumed by the responder.
4. The STUN Binding Requests generated by the initiator MAY include the USE-CANDIDATE attribute to indicate that the initiator wishes to cease checks for this component.
5. The STUN Binding Requests generated by the initiator MUST include the ICE-CONTROLLING attribute.
6. The STUN Binding Requests generated by the responder MUST include the ICE-CONTROLLED attribute.
7. The parties MUST use STUN short term credentials to authenticate requests and perform message integrity checks. As in ICE-CORE, the username in the STUN Binding Request is of the form "ufrag-of-peer:ufrag-of-sender" and the password is the value of the 'pwd'


⁹RFC 5389: Session Traversal Utilities for NAT (STUN) <<http://tools.ietf.org/html/rfc5389>>.

attribute provided by the peer.¹⁰

When it receives a STUN Binding Request, each party MUST return a STUN Binding Response, which indicates either an error case or the success case. As described in Section 7.1.2.2 of ICE-CORE, a connectivity check succeeds if *all* of the following are true:

1. The STUN transaction generated a success response.
2. The source IP address and port of the response equals the destination IP address and port to which the Binding Request was sent.
3. The destination IP address and port of the response match the source IP address and port from which the Binding Request was sent.

For the candidates exchanged in the previous section, the connectivity checks would be as follows (this diagram mirrors the example in ICE-CORE).



¹⁰Thus when Romeo sends a STUN Binding Request to Juliet the credentials will be STUN username "9uB6:8hhy" (ufrag provided by Juliet concatenated with ufrag provided by Romeo) and password "YH75Fviy6338Vbrhrp8Yh" (pwd provided by Juliet) whereas when Juliet sends a STUN Binding Request to Romeo the credentials will be STUN username "8hhy:9uB6" (ufrag provided by Romeo concatenated with ufrag provided by Juliet) and password "asd88fgpdd777uzjYhagZg" (pwd provided by Romeo).

```

| <===== |
| |
| <==Media Now Can Flow== |
| |
| | STUN Binding Request |
| | from 192.0.2.1:3478 |
| | to 192.0.2.3:45664 |
| | <===== |
| | STUN Binding Request |
| | from 192.0.2.1:3478 |
| | to 10.0.1.1:8998 |
| | <===== |
| | STUN Binding Response |
| | from 10.0.1.1:8998 |
| | to 192.0.2.1:3478 |
| | map 192.0.2.1:3478 |
| | =====> |
| | STUN Binding Response |
| | from 192.0.2.3:45664 |
| | to 192.0.2.1:3478 |
| | map 192.0.2.1:3478 |
| | =====> |
| | ==Media Now Can Flow==> |
| |

```

Note: Here the initiator (controlling agent) is using "aggressive nomination" as described in Section 8.1.1.2 of ICE-CORE and therefore includes the USE-CANDIDATE attribute in the STUN Binding Requests it sends.

5.7 Acceptance of Successful Candidate

If, based on STUN connectivity checks, the parties determine that they will be able to exchange media between a given pair of local candidates and remote candidates (i.e., the pair is "nominated" and ICE processing is "completed"), they can then begin using that candidate pair to exchange media.

Once the parties have connectivity and therefore the initiator has completed ICE as explained in ICE-CORE, the initiator MAY communicate the in-use candidate pair in the signalling channel by sending a transport-info message that contains a <remote-candidate/> element (this maps to the SDP "remote-candidates" attribute as described in Section B.6 of ICE-CORE, i.e., remote candidates are "the actual candidates at R that were selected by the offerer", of which there will be only one at this stage of the ICE-UDP negotiation).

Listing 4: Initiator communicates in-use candidate

```
<iq from='romeo@montague.lit/orchard'
```

```

    id='pd81b49s'
    to='juliet@capulet.lit/balcony'
    type='set'>
<jingle xmlns='urn:xmpp:jingle:1'
  action='transport-info'
  initiator='romeo@montague.lit/orchard'
  sid='a73sjjvkl37jfea'>
  <content creator='initiator' name='this-is-the-audio-content'>
    <transport xmlns='urn:xmpp:jingle:transports:ice-udp:1'
      pwd='asd88fgpdd777uzjYhagZg'
      ufrag='8hhy'>
      <remote-candidate component='1'
        ip='10.0.1.2'
        port='9001' />
    </transport>
  </content>
</jingle>
</iq>

```

(In accordance with Jingle core, the responder will also acknowledge the transport-info message.)

In the unlikely event that one of the parties determines that it cannot establish connectivity even after sending and checking lower-priority candidates, it SHOULD terminate the session as described in XEP-0166.

5.8 Negotiating a New Candidate

Even after media has begun to flow, either party MAY continue to send additional candidates to the other party (e.g., because the user agent has become aware of a new media proxy or network interface card). Such candidates are shared by sending a transport-info message.

Listing 5: Initiator sends a subsequent candidate

```

<iq from='romeo@montague.lit/orchard'
  id='uh3g1f48'
  to='juliet@capulet.lit/balcony'
  type='set'>
  <jingle xmlns='urn:xmpp:jingle:1'
    action='transport-info'
    initiator='romeo@montague.lit/orchard'
    sid='a73sjjvkl37jfea'>
    <content creator='initiator' name='this-is-the-audio-content'>
      <transport xmlns='urn:xmpp:jingle:transports:ice-udp:1'
        pwd='asd88fgpdd777uzjYhagZg'
        ufrag='8hhy'>
        <candidate component='1'
          foundation='1'

```

```

        generation='0'
        id='m3110wc4nd'
        ip='2001:db8::9:1'
        network='0'
        port='9001'
        priority='21149780477'
        protocol='udp'
        type='host' />
    </transport>
</content>
</jingle>
</iq>

```

The receiving party MUST acknowledge receipt of the candidate.

Listing 6: Recipient acknowledges receipt

```

<iq from='juliet@capulet.lit/balcony'
    id='uh3g1f48'
    to='romeo@montague.lit/orchard'
    type='result' />

```

The parties would check the newly-offered candidate for connectivity, as described previously. If the parties determine that media can flow over the candidate, they MAY then use the new candidate in subsequent communications.

5.9 ICE Restarts

At any time, either party MAY restart the process of ICE negotiation by sending a candidate with a 'generation' value that is greater than the previous generation of candidates; when it does so, it MUST generate new values for the 'pwd' and 'ufrag' attributes, consistent with the definition of an ICE restart in Section 9.1.1.1 of ICE-CORE. As explained in ICE-CORE, typically the ICE negotiation would be restarted to change the media target (e.g., an IP address change for one of the parties) and certain third-party-call-control scenarios.

Listing 7: Initiator restarts ICE negotiation

```

<iq from='romeo@montague.lit/orchard'
    id='kl23fs71'
    to='juliet@capulet.lit/balcony'
    type='set'>
  <jingle xmlns='urn:xmpp:jingle:1'
        action='transport-info'
        initiator='romeo@montague.lit/orchard'
        sid='a73sjjvkl37jfea'>
    <content creator='initiator' name='this-is-the-audio-content'>
      <transport xmlns='urn:xmpp:jingle:transports:ice-udp:1'

```



```

        pwd='bv71hdn38hgb39hf6x1k33'
        ufrag='g7qs'>
    <candidate component='1'
        foundation='1'
        generation='1'
        id='y3s2b30v3r'
        ip='192.0.2.3'
        network='1'
        port='45665'
        priority='1694498815'
        protocol='udp'
        type='srflx' />
</transport>
</content>
</jingle>
</iq>

```

The recipient then acknowledges receipt.

Listing 8: Recipient acknowledges transport-info

```

<iq from='juliet@capulet.lit/balcony'
    id='kl23fs71'
    to='romeo@montague.lit/orchard'
    type='result' />

```

The parties would then exchange new candidates to renegotiate connectivity and would check the new candidates for connectivity, as described previously. If the parties determine that media can flow over one of the new candidates, they can then use the successful candidate in subsequent communications. However, while ICE is being renegotiated the parties can continue to send media with the existing candidate-in-use.

6 Fallback to Raw UDP


It can happen that the responder does not support ICE, in which case it can be necessary to fall back to use of the [Jingle Raw UDP Transport Method \(XEP-0177\)](#)¹¹. One typical scenario is communication between an ICE-aware Jingle endpoint and a non-ICE-aware SIP endpoint through a Jingle-to-SIP gateway, as follows:

1. The Jingle endpoint sends a session-initiate request to the SIP endpoint, specifying a transport method of ICE-UDP.
2. Based on capabilities information, the gateway knows that the SIP endpoint does not support ICE, so it enables the endpoints to use its media relay. It does this by:

¹¹XEP-0177: Jingle Raw UDP Transport Method <<https://xmpp.org/extensions/xep-0177.html>>.

- Sending a transport-replace message to the Jingle endpoint on behalf of the SIP endpoint, specifying a transport method of Raw UDP and a candidate whose IP address and port are hosted at the gateway.
- Sending SIP INVITE to the SIP endpoint on behalf of the Jingle endpoint, specifying an IP address and port at the gateway.

The session flow is as follows.



The protocol flow is as follows, showing only the stanzas sent between Romeo and the gateway (acting on Juliet's behalf).

Listing 9: Initiator sends session-initiate

```
<iq from='romeo@montague.lit/orchard'
```

```

id='p01hf63x'
to='juliet@capulet.lit/balcony'
type='set'>
<jingle xmlns='urn:xmpp:jingle:1'
  action='session-initiate'
  initiator='romeo@montague.lit/orchard'
  sid='a73sjjvkl37jfea'>
  <content creator='initiator' name='voice'>
    <description xmlns='urn:xmpp:jingle:apps:rtp:1' media='audio'>
      <payload-type id='96' name='speex' clockrate='16000' />
      <payload-type id='97' name='speex' clockrate='8000' />
      <payload-type id='18' name='G729' />
      <payload-type id='103' name='L16' clockrate='16000' channels='
        2' />
      <payload-type id='98' name='x-ISAC' clockrate='8000' />
    </description>
    <transport xmlns='urn:xmpp:jingle:transports:ice-udp:1'
      pwd='asd88fgpdd777uzjYhagZg'
      ufrag='8hhy'>
      <candidate component='1'
        foundation='1'
        generation='0'
        id='el0747fg11'
        ip='10.0.1.1'
        network='1'
        port='8998'
        priority='2130706431'
        protocol='udp'
        type='host' />
      <candidate component='1'
        foundation='2'
        generation='0'
        id='y3s2b30v3r'
        ip='192.0.2.3'
        network='1'
        port='45664'
        priority='1694498815'
        protocol='udp'
        rel-addr='10.0.1.1'
        rel-port='8998'
        type='srflx' />
    </transport>
  </content>
</jingle>
</iq>

```

Listing 10: Responder acknowledges session-initiate

```

<iq from='juliet@capulet.lit/balcony'
  id='p01hf63x'

```

```
to='romeo@montague.lit/orchard'
type='result'/>
```

Immediately the gateway sends a transport-replace message to Romeo, specifying a transport of Raw UDP with a candidate whose IP address and port identify a media relay at the gateway.

Listing 11: Gateway sends transport-replace on behalf of responder

```
<iq from='juliet@capulet.lit/balcony'
  id='hy2gd714'
  to='romeo@montague.lit/orchard'
  type='set'>
  <jingle xmlns='urn:xmpp:jingle:1'
    action='transport-replace'
    initiator='romeo@montague.lit/orchard'
    sid='a73sjjvkl37jfea'>
    <content creator='initiator' name='voice1'>
      <transport xmlns='urn:xmpp:jingle:transports:raw-udp:1'>
        <candidate generation='0'
          id='a9j3mnbu1'
          ip='10.1.1.104'
          port='13540' />
      </transport>
    </content>
  </jingle>
</iq>
```

Romeo then acknowledges the transport-replace message and immediately also sends a transport-accept.

Listing 12: Initiator acknowledges transport-replace

```
<iq from='romeo@montague.lit/orchard'
  id='hy2gd714'
  to='juliet@capulet.lit/balcony'
  type='result'/>
```

Listing 13: Initiator accepts new transport

```
<iq from='romeo@montague.lit/orchard'
  id='rb391gs5'
  to='juliet@capulet.lit/balcony'
  type='set'>
  <jingle xmlns='urn:xmpp:jingle:1'
    action='transport-accept'
    initiator='romeo@montague.lit/orchard'
    sid='a73sjjvkl37jfea'>
    <content creator='responder' name='voice2'>
      <transport xmlns='urn:xmpp:jingle:transports:raw-udp:1'>
```

```

        <candidate generation='0'
            id='a9j3mnbtu1'
            ip='10.1.1.104'
            port='13540' />
    </transport>
</content>
</jingle>
</iq>

```

The gateway then acknowledges the acceptance on behalf of Juliet.

Listing 14: Gateway acknowledges transport-accept

```

<iq from='juliet@capulet.lit/balcony'
    id='rb391gs5'
    to='romeo@montague.lit/orchard'
    type='result' />

```

The responder then sends a session-accept through the gateway.

Listing 15: Responder sends session-accept

```

<iq from='juliet@capulet.lit/balcony'
    id='ijf61d43'
    to='romeo@montague.lit/orchard'
    type='set'>
  <jingle xmlns='urn:xmpp:jingle:1'
    action='session-accept'
    initiator='romeo@montague.lit/orchard'
    responder='juliet@capulet.lit/balcony'
    sid='a73sjjvcla37jfea'>
    <content creator='initiator' name='voice'>
      <description xmlns='urn:xmpp:jingle:apps:rtp:1' media='audio'>
        <payload-type id='18' name='G729' />
      </description>
      <transport xmlns='urn:xmpp:jingle:transports:raw-udp:1' />
    </content>
  </jingle>
</iq>

```

Listing 16: Initiator acknowledges session-accept

```

<iq from='romeo@montague.lit/orchard'
    id='ijf61d43'
    to='juliet@capulet.lit/balcony'
    type='result' />

```

The endpoints now begin to exchange session media, and can continue the session as long as desired.

7 Determining Support

7.1 ICE Support

To advertise its support for the Jingle ICE-UDP Transport Method, when replying to [Service Discovery \(XEP-0030\)](#)¹² information requests an entity MUST return URNs for any version of this protocol that the entity supports -- e.g., "urn:xmpp:jingle:transports:ice-udp:1" for this version and "urn:xmpp:jingle:transports:ice-udp:0" for the previous version (see [Namespace Versioning](#) regarding the possibility of incrementing the version number).

Listing 17: Service discovery information request

```
<iq from='romeo@montague.lit/orchard'
  id='cv5x41g9'
  to='juliet@capulet.lit/balcony'
  type='get'>
  <query xmlns='http://jabber.org/protocol/disco#info' />
</iq>
```

Listing 18: Service discovery information response

```
<iq from='juliet@capulet.lit/balcony'
  id='cv5x41g9'
  to='romeo@montague.lit/orchard'
  type='result'>
  <query xmlns='http://jabber.org/protocol/disco#info'>
    <feature var='urn:xmpp:jingle:1' />
    <feature var='urn:xmpp:jingle:transports:ice-udp:0' />
    <feature var='urn:xmpp:jingle:transports:ice-udp:1' />
    <feature var='urn:xmpp:jingle:apps:rtp:1' />
    <feature var='urn:xmpp:jingle:apps:rtp:audio' />
    <feature var='urn:xmpp:jingle:apps:rtp:video' />
  </query>
</iq>
```

In order for an application to determine whether an entity supports this protocol, where possible it SHOULD use the dynamic, presence-based profile of service discovery defined in [Entity Capabilities \(XEP-0115\)](#)¹³. However, if an application has not received entity capabilities information from an entity, it SHOULD use explicit service discovery instead.

7.2 SDP Offer / Answer Support

If an entity supports the SDP offer / answer model described in RFC 3264 and therefore prefers to receive multiple candidates in a single transport-info message, it MUST advertise support

¹²XEP-0030: Service Discovery <<https://xmpp.org/extensions/xep-0030.html>>.

¹³XEP-0115: Entity Capabilities <<https://xmpp.org/extensions/xep-0115.html>>.

for the "urn:ietf:rfc:3264" service discovery feature. Typically this feature will be advertised only by gateways between Jingle and SIP.

Listing 19: Service discovery information request

```
<iq from='romeo@montague.lit/orchard'
  id='ce81f5d6'
  to='sip.shakespeare.lit'
  type='get'>
  <query xmlns='http://jabber.org/protocol/disco#info' />
</iq>
```

Listing 20: Service discovery information response

```
<iq from='sip.shakespeare.lit'
  id='ce81f5d6'
  to='romeo@montague.lit/orchard'
  type='result'>
  <query xmlns='http://jabber.org/protocol/disco#info'>
    <feature var='urn:ietf:rfc:3264' />
    <feature var='urn:xmpp:jingle:1' />
    <feature var='urn:xmpp:jingle:transports:ice-udp:0' />
    <feature var='urn:xmpp:jingle:transports:ice-udp:1' />
    <feature var='urn:xmpp:jingle:apps:rtp:1' />
    <feature var='urn:xmpp:jingle:apps:rtp:audio' />
    <feature var='urn:xmpp:jingle:apps:rtp:video' />
  </query>
</iq>
```

8 Implementation Notes

In order to speed the negotiation process so that media can flow as quickly as possible, the initiator SHOULD gather and prioritize candidates in advance, or as soon as the principal begins the process of initiating a session.

9 Deployment Notes

This specification applies exclusively to Jingle clients and places no additional requirements on XMPP servers. However, service administrators might wish to deploy a STUN server in order to ease the client-to-client negotiation process and a TURN server for media relaying (see [TURN](#)¹⁴). Deployment of support for [External Service Discovery \(XEP-0215\)](#)¹⁵ might also

¹⁴Traversal Using Relays around NAT (TURN): Relay Extensions to Session Traversal Utilities for NAT (STUN) <<http://tools.ietf.org/html/draft-ietf-behave-turn>>. Work in progress.

¹⁵XEP-0215: External Service Discovery <<https://xmpp.org/extensions/xep-0215.html>>.

be helpful.

10 Security Considerations

10.1 Sharing IP Addresses

By definition, the exchange of transport candidates results in exposure of the sender's IP addresses, which comprise a form of personally identifying information. A Jingle client **MUST** enable a user to control which entities will be allowed to receive such information. If a human user explicitly accepts a session request, then the client **SHOULD** consider that action to imply approval of IP address sharing. However, waiting for a human user to explicitly accept the session request can result in delays during session setup, since it is more efficient to immediately begin sharing transport candidates. Therefore, it is **RECOMMENDED** for the client to immediately send transport candidates to a contact (without waiting for explicit user approval of the session request) in the following cases:

1. The user has permanently and formally authorized the contact to view the user's presence information via a presence subscription as reflected in an XMPP roster item (see [XMPP IM](#) ¹⁶).
2. The user has temporarily and dynamically shared presence with the contact via "directed presence" as described in RFC 3921.
3. The user has explicitly added the contact to a "whitelist" of entities who are allowed to access the user's personally-identifying information.

10.2 Encryption of Media

A Jingle implementation **SHOULD** support security preconditions that are enforced before application media is allowed to flow over a UDP association, such as those described in [Jingle XTLS](#) ¹⁷.

Application types that use the Jingle ICE-UDP transport method **MAY** also define their own application-specific encryption methods, such as the Secure Real-time Transport Protocol (SRTP) for RTP exchanges as described in [Jingle RTP Sessions \(XEP-0167\)](#) ¹⁸.

¹⁶RFC 6121: Extensible Messaging and Presence Protocol (XMPP): Instant Messaging and Presence <<http://tools.ietf.org/html/rfc6121>>.

¹⁷Extensible Messaging and Presence Protocol (XMPP) End-to-End Encryption Using Transport Layer Security ("XTLS") <<http://tools.ietf.org/html/draft-meyer-xmpp-e2e-encryption>>.

¹⁸XEP-0167: Jingle RTP Sessions <<https://xmpp.org/extensions/xep-0167.html>>.

11 IANA Considerations

This document requires no interaction with the [Internet Assigned Numbers Authority \(IANA\)](#)¹⁹.

12 XMPP Registrar Considerations

12.1 Protocol Namespaces

This specification defines the following XML namespace:

- urn:xmpp:jingle:transports:ice-udp:1

The [XMPP Registrar](#)²⁰ includes the foregoing namespace in its registry at [<https://xmpp.org/registrar/namespaces.html>](https://xmpp.org/registrar/namespaces.html), as governed by [XMPP Registrar Function \(XEP-0053\)](#)²¹.

12.2 Protocol Versioning

If the protocol defined in this specification undergoes a revision that is not fully backwards-compatible with an older version, the XMPP Registrar shall increment the protocol version number found at the end of the XML namespaces defined herein, as described in Section 4 of XEP-0053.

12.3 Service Discovery Features

If an entity supports the SDP offer / answer model described in RFC 3264 and therefore prefers to receive one transport-info message with multiple candidates, it MUST advertise support for the "urn:ietf:rfc:3264" feature.

The registry submission is as follows.

```
<var>
  <name>urn:ietf:rfc:3264</name>
  <desc>
    Signals support for the SDP offer / answer model
```

¹⁹The Internet Assigned Numbers Authority (IANA) is the central coordinator for the assignment of unique parameter values for Internet protocols, such as port numbers and URI schemes. For further information, see <http://www.iana.org/>.

²⁰The XMPP Registrar maintains a list of reserved protocol namespaces as well as registries of parameters used in the context of XMPP extension protocols approved by the XMPP Standards Foundation. For further information, see <https://xmpp.org/registrar/>.

²¹XEP-0053: XMPP Registrar Function [<https://xmpp.org/extensions/xep-0053.html>](https://xmpp.org/extensions/xep-0053.html).

```

    described in RFC 3264
  </desc>
  <doc>XEP-0176</doc>
</var>

```

12.4 Jingle Transport Methods

The XMPP Registrar includes "ice-udp" in its registry of Jingle transport methods at <https://xmpp.org/registrar/jingle-transports.html>. The registry submission is as follows:

```

<transport>
  <name>ice-udp</name>
  <desc>
    A method for negotiation of out-of-band UDP connections
    with built-in NAT and firewall traversal, equivalent to
    the IETF's Interactive Connectivity Establishment (ICE)
    methodology when resulting in the use of UDP as the
    transport protocol.
  </desc>
  <type>datagram</type>
  <doc>XEP-0176</doc>
</transport>

```

13 XML Schema

```

<?xml version='1.0' encoding='UTF-8'?>
<xs:schema
  xmlns:xs='http://www.w3.org/2001/XMLSchema'
  targetNamespace='urn:xmpp:jingle:transports:ice-udp:1'
  xmlns='urn:xmpp:jingle:transports:ice-udp:1'
  elementFormDefault='qualified'>
  <xs:annotation>
    <xs:documentation>
      The protocol documented by this schema is defined in
      XEP-0176: http://www.xmpp.org/extensions/xep-0176.html
    </xs:documentation>
  </xs:annotation>
  <xs:element name='transport'>
    <xs:complexType>
      <xs:choice minOccurs='0'>
        <xs:sequence>

```

```
<xs:element name='candidate'
            type='candidateElementType'
            minOccurs='1'
            maxOccurs='unbounded' />
</xs:sequence>
<xs:sequence>
  <xs:element name='remote-candidate'
              type='remoteCandidateElementType'
              minOccurs='1'
              maxOccurs='1' />
</xs:sequence>
</xs:choice>
<xs:attribute name='pwd' type='xs:string' use='optional' />
<xs:attribute name='ufrag' type='xs:string' use='optional' />
</xs:complexType>
</xs:element>

<xs:complexType name='candidateElementType'>
  <xs:simpleContent>
    <xs:extension base='empty'>
      <xs:attribute name='component' type='xs:unsignedByte' use='
        required' />
      <xs:attribute name='foundation' type='xs:unsignedByte' use='
        required' />
      <xs:attribute name='generation' type='xs:unsignedByte' use='
        required' />
      <xs:attribute name='id' type='xs:NCName' use='required' />
      <xs:attribute name='ip' type='xs:string' use='required' />
      <xs:attribute name='network' type='xs:unsignedByte' use='
        required' />
      <xs:attribute name='port' type='xs:unsignedShort' use='
        required' />
      <xs:attribute name='priority' type='xs:positiveInteger' use='
        required' />
      <xs:attribute name='protocol' type='xs:NCName' use='required' />
      <
        <xs:attribute name='rel-addr' type='xs:string' use='optional' />
      >
      <xs:attribute name='rel-port' type='xs:unsignedShort' use='
        optional' />
    <xs:attribute name='type' use='required'>
      <xs:simpleType>
        <xs:restriction base='xs:NCName'>
          <xs:enumeration value='host' />
          <xs:enumeration value='prflx' />
          <xs:enumeration value='relay' />
          <xs:enumeration value='srflx' />
        </xs:restriction>
      </xs:simpleType>
    </xs:attribute>
  </xs:simpleContent>
</xs:complexType>
```

```
        </xs:attribute>
    </xs:extension>
</xs:simpleContent>
</xs:complexType>

<xs:complexType name='remoteCandidateElementType'>
  <xs:simpleContent>
    <xs:extension base='empty'>
      <xs:attribute name='component' type='xs:unsignedByte' use='
        required' />
      <xs:attribute name='ip' type='xs:string' use='required' />
      <xs:attribute name='port' type='xs:unsignedShort' use='
        required' />
    </xs:extension>
  </xs:simpleContent>
</xs:complexType>

<xs:simpleType name='empty'>
  <xs:restriction base='xs:string'>
    <xs:enumeration value='' />
  </xs:restriction>
</xs:simpleType>
</xs:schema>
```

14 Acknowledgements

Thanks to Diana Cionoiu, Olivier Crête, Tim Julien, Steffen Larsen, Robert McQueen, Mike Ruprecht, Justin Uberti, Unnikrishnan Vikrama Panicker, and Paul Witty for their feedback.