Multi-User Jingle

Sjoerd Simons

Collabora Ltd, sponsored by NINet

LCA 2009

Multi-User Jingle

XMPP protocol extension for initiating and managing multiparty voice and video conferences within an XMPP MUC using Jingle

Jingle

XMPP protocol extension for initiating and managing peer-to-peer media sessions between two XMPP entities

XMPP

- Extensible Messaging and Presence Protocol
- The protocol formaly known as Jabber
- XML streams
- XMPP Core: RFC 3920
- XMPP IM: RFC 3921
- Extensions using XMPP Extension Proposals (XEP)

XMPP JID

Almost but not entirely unlike an e-mailaddress:

- sjoerd.simons@collabora.co.uk
- romeo@example.net
- juliet@capulet.com/balcony

Messages

```
<message to='romeo@example.net'
  from='juliet@example.com/balcony'
    type='chat' xml:lang='en'>
   <body>Wherefore art thou, Romeo?</body>
</message>
```

Presence

```
<presence from='romeo@montague.lit/orchard'>
  <show>dnd</show>
  <status>Wooing Juliet</status>
  <c xmlns='http://jabber.org/protocol/caps'
    hash='sha-1'
    node='http://telepathy.freedesktop.org/wiki/Muji'
    ver='iuNYhdSOy4nYDVPDhoTqCFubSz8='/>
</presence>
```

Info/Query

```
<iq xmlns='jabber:client'
    to='juliet@example.com/balcony'
    from='romeo@example.net/orchard'
    id='2846139002'
      type='get'>
  <query xmlns='http://jabber.org/protocol/disco#info'
    node='http://telepathy.freedesktop.org/wiki/Muji#iuNYhdSOy4nYDVPDhoTqCFubSz8=' />
</iq>
```

```
<iq xmlns='jabber:client'
  to='romeo@example.net/orchard'
  from='juliet@example.com/balcony'
  id='2846139002'
  type='result'>
<query xmlns='http://jabber.org/protocol/disco#info'>
  <identity category='client'
            type='pc'
            name='Telepathy␣Gabble␣0.7.17' />
  <feature var='urn:xmpp:jingle:0' />
  <feature var='urn:xmpp:jingle:apps:rtp:0' />
  <feature var='urn:xmpp:jingle:transports:raw-udp:0' />
</query>
</iq>
```

Jingle

XMPP protocol extension for initiating and managing peer-to-peer media sessions between two XMPP entities

- XEP-0166: Jingle
- XEP-0167: Jingle RTP Sessions
- XEP-0176: Jingle ICE-UDP Transport Method
- XEP-0177: Jingle Raw UDP Transport Method

Jingle session initiation

```
<iq from='romeo@montague.lit/orchard'
    id='jingle1'
    to='juliet@capulet.lit/balcony'
    type='set'>
  <jingle xmlns='urn:xmpp:jingle:0'
          action='session-initiate'
          initiator='romeo@montague.lit/orchard'
          sid='a73sjjvkla37jfea'>
    <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'/>
      </description>
      <transport xmlns='urn:xmpp:jingle:transports:ice-udp:0'/>
    </content>
    <content creator='initiator' name='video'>
      <description xmlns='urn:xmpp:jingle:apps:rtp:1' media='video'>
        <payload-type id='98' name='THEORA' clockrate='90000'>
          <parameter name='height' value='600'/>
          <parameter name='width' value='800'/>
          <parameter name='delivery-method' value='inline'/>
          <parameter name='configuration' value='somebase16string'/>
        </payload-type>
      </description>
      <transport xmlns='urn:xmpp:jingle:transports:ice-udp:0'/>
    </content>
  </jingle>
</iq>
```

Jingle session initiation

```
<content creator='initiator' name='video'>
  <description
    xmlns='urn:xmpp:jingle:apps:rtp:1' media='video'>
    <payload-type id='98' name='THEORA'
      clockrate='90000'>
      <parameter name='height' value='600'/>
      <parameter name='width' value='800'/>
      <parameter name='delivery-method'
        value='inline'/>
      <parameter name='configuration'
        value='somebase16string'/>
    </payload-type>
  </description>
  <transport
    xmlns='urn:xmpp:jingle:transports:ice-udp:0'/>
</content>
```

Transport information

```
<iq from='romeo@montague.lit/orchard'
    id='info1'
    to='juliet@capulet.lit/balcony'
    type='set'>
  <jingle xmlns='urn:xmpp:jingle:0'
          action='transport-info'
          initiator='romeo@montague.lit/orchard'
          sid='a73sjjvkla37jfea'>
    <content creator='initiator' name='voice'>
      <transport xmlns='urn:xmpp:jingle:transports:ice-udp:0'
                 pwd='asd88fgpdd777uzjYhagZg'
                 ufrag='8hhy'>
        <candidate component='1'
                   foundation='1'
                   generation='0'
                   ip='10.0.1.1'
                   network='1'
                   port='8998'
                   priority='2130706431'
                   protocol='udp'
                   type='host'/>
      </transport>
    </content>
  </jingle>
</iq>
```

Accepting a Jingle session

```
<iq from='juliet@capulet.lit/balcony'
    id='accept1'
    to='romeo@montague.lit/orchard'
    type='set'>
  <jingle xmlns='urn:xmpp:jingle:0'
          action='session-accept'
          initiator='romeo@montague.lit/orchard'
          responder='juliet@capulet.lit/balcony'
          sid='a73sjjvkla37jfea'>
    <content creator='initiator' name='voice'>
      <description xmlns='urn:xmpp:jingle:apps:rtp:1' media='voice'>
        <payload-type id='96' name='speex' clockrate='16000'/>
        <payload-type id='97' name='speex' clockrate='8000'/>
      </description>
      <transport xmlns='urn:xmpp:jingle:transports:ice-udp:0'>
      </transport>
    </content>
    <content creator='initiator' name='video'>
      <description xmlns='urn:xmpp:jingle:apps:rtp:1' media='video'>
        <payload-type id='98' name='theora' clockrate='90000'>
          <parameter name='height' value='600'/>
          <parameter name='width' value='800'/>
          <parameter name='delivery-method' value='inline'/>
          <parameter name='configuration' value='somebase16string'/>
          <parameter name='sampling' value='YCbCr-4:2:2'/>
        </payload-type>
      </description>
      <transport xmlns='urn:xmpp:jingle:transports:ice-udp:0'/>
    </content>
  </jingle>
</iq>
```

Multi User Chat

- MUC: XEP-0045
- Widely deployed
- Messages in the muc are ordered!

Joining a muc

```
<presence
    from='hag66@shakespeare.lit/pda'
    to='darkcave@chat.shakespeare.lit/thirdwitch'/>
```

The MUC responds with presences of all current MUC members, with our own presence as the last.

Multi-User Jingle and MUC

- The MUC server shouldn't need to provide any infrastructure for Multi-User Jingle
- Shouldn't interfere with clients not supporting Multi-User Jingle
- Shouldn't depend on the one leader

Multi-User Jingle

- Each client announces the streams it's interested in and it's payload mapping in its presence
- Each clients does its own payload mapping calculation when joining based on the information from existing clients
- Each clients sets up a Jingle session to all existing clients
- Joining the Multi-User Jingle session split in in two stages to prevent race conditions.

Joining step 1

Announce that you're preparing to start or join the session:

```
<presence xmlns="jabber:client"
          to="test0@night.luon.net/772205543123238993511248 6"
          from="muji@conference.night.luon.net/test0">
  <c xmlns="http://jabber.org/protocol/caps"
     node="http://telepathy.freedesktop.org/wiki/Mingle"
     hash="sha-1"
     ver="9KIXkd+F+C/WPRXVRn77ai5k4qI=" />
  <muji xmlns="http://telepathy.freedesktop.org/muji">
    <content name="video"
             preparing="true" />
    <content name="audio"
             preparing="true" />
  </muji>
  <x xmlns="http://jabber.org/protocol/muc#user">
    <item affiliation="owner"
          jid="test0@night.luon.net/772205543123238993511248 6"
          role="moderator" />
  </x>
</presence>
```

Joining step 2

After all clients that announced they were preparing to join have done so, update your presence with your payload mapping:

```
<presence xmlns="jabber:client"
          to="test0@night.luon.net/772205543123238993511248 6"
          from="muji@conference.night.luon.net/test0">
  <c xmlns="http://jabber.org/protocol/caps"
     node="http://telepathy.freedesktop.org/wiki/Muji"
     hash="sha-1"
     ver="9KlXkd+F+C/WPRXVRn77ai5k4qI=" />
  <muji xmlns="http://telepathy.freedesktop.org/muji">
    <content name="audio">
      <description xmlns="urn:xmpp:jingle:apps:rtp:0"
                   media="audio">
        <payload-type id='96' name='speex' clockrate='16000'/>
        <payload-type id='97' name='speex' clockrate='8000'/>
      </description>
    </content>
    <content name="video">
      <description xmlns="urn:xmpp:jingle:apps:rtp:0"
                   media="video">
        <payload-type id="96" name="THEORA" clockrate="90000">
          <parameter name="delivery-method" value="inline" />
          <parameter name="configuration" value="somebase16value" />
        </payload-type>
      </description>
    </content>
  </muji>
  <x xmlns="http://jabber.org/protocol/muc#user">
    <item affiliation="owner" jid="test0@night.luon.net/772205543123238993511248 6"
          role="moderator" />
  </x>
</presence>
```

Demo!

- Python client to make it easy to play with the protocol
- Farsight2 for the RTP streaming
- libnice (through farsight) for the ICE transport

Multi-User Jingle

Advantages:

- Only needs basic infrastructure
- Each client can decide how to present the conference

Disadvantages:

- Does not scale to large groups

Scaling Multi-User jingle

Scaling without extra infrastructure is limited by:

- Amount of streams a client can decode
- Amount of upstream bandwidth
- Amount of downstream bandwidth

RTP relay

Clients send their streams to the relay, which relays it to (all) others

Advantage:

- Removes the limitation on the clients upstream bandwidth

Disadvantages:

- Adds some latency especially when the relay is badly choosen

Mixer

Clients their streams to the mixer, which combines (mixes them) in some way and send them to others.
Advantages:

- Removes the limitation on the amount of streams can be decoded by the client
- Removes the limitation on the clients downstream bandwidth
- Removes the limitation on the clients upstream bandwidth

Mixer

Disadvantages:

- The mixer needs to reencode and mix the data
  - CPU intensive (transcoding)
  - Lowers the quality
- Clients don't have control over the presentation anymore.

Combined relay and mixer

Let clients choose whether they want to receive a mixed stream or not.
Advantages:

- Removes the limitation on clients upstream bandwidth.
- Allows clients to decide about the trade-off between downstream bandwidth and having higher quality/control over the presentation

Codecs

- Each client only encodecs the media once. Which means the conference ends up using the lowest common denominator
- Not such a big issue for audio
- Scalable video codecs might be a solution
  - Have a low resolution, low quality base layer.. On which extra layers are added to improve quality
  - Decide per client which layers to send
  - H264/SVC exists, but is heavily patented
  - No free scalable video available :(

Links

- Project wiki: `http://telepathy.freedesktop.org/wiki/MultiUserJingle`
- NINet site: `http://www.nlnet.nl/project/mujingle/`
- Farsight2: `http://farsight.freedesktop.org`
- libnice: `http://nice.freedesktop.org`