

Network Working Group  
Internet-Draft  
Intended status: Standards Track  
Expires: May 19, 2012

E. Iovov, Ed.  
Jitsi  
E. Marocco, Ed.  
Telecom Italia  
J. Lennox  
Vidyo, Inc.  
November 16, 2011

A Real-Time Transport Protocol (RTP) Header Extension for Mixer-to-  
Client Audio Level Indication  
draft-ietf-avtext-mixer-to-client-audio-level-06

## Abstract

This document describes a mechanism for RTP-level mixers in audio conferences to deliver information about the audio level of individual participants. Such audio level indicators are transported in the same RTP packets as the audio data they pertain to.

## 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 May 19, 2012.

## 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

---

Internet-Draft    Mixer-to-Client Audio Level Indication    November 2011

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

<a href="#">1.</a>	<a href="#">Introduction</a>	<a href="#">3</a>
<a href="#">2.</a>	<a href="#">Terminology</a>	<a href="#">4</a>
<a href="#">3.</a>	<a href="#">Protocol Operation</a>	<a href="#">4</a>
<a href="#">4.</a>	<a href="#">Audio Levels</a>	<a href="#">5</a>
<a href="#">5.</a>	<a href="#">Signaling Information</a>	<a href="#">7</a>
<a href="#">6.</a>	<a href="#">Security Considerations</a>	<a href="#">10</a>
<a href="#">7.</a>	<a href="#">IANA Considerations</a>	<a href="#">11</a>
<a href="#">8.</a>	<a href="#">Acknowledgments</a>	<a href="#">11</a>
<a href="#">9.</a>	<a href="#">Changes From Earlier Versions</a>	<a href="#">11</a>
<a href="#">9.1.</a>	<a href="#">Changes From Draft -05</a>	<a href="#">11</a>
<a href="#">9.2.</a>	<a href="#">Changes From Draft -04</a>	<a href="#">12</a>
<a href="#">9.3.</a>	<a href="#">Changes From Draft -03</a>	<a href="#">12</a>
<a href="#">9.4.</a>	<a href="#">Changes From Draft -02</a>	<a href="#">12</a>
<a href="#">9.5.</a>	<a href="#">Changes From Draft -01</a>	<a href="#">12</a>
<a href="#">9.6.</a>	<a href="#">Changes From Draft -00</a>	<a href="#">13</a>
<a href="#">10.</a>	<a href="#">References</a>	<a href="#">13</a>
<a href="#">10.1.</a>	<a href="#">Normative References</a>	<a href="#">13</a>
<a href="#">10.2.</a>	<a href="#">Informative References</a>	<a href="#">13</a>
<a href="#">Appendix A.</a>	<a href="#">Reference Implementation</a>	<a href="#">14</a>
<a href="#">A.1.</a>	<a href="#">AudioLevelCalculator.java</a>	<a href="#">14</a>
	<a href="#">Authors' Addresses</a>	<a href="#">17</a>

---

Internet-Draft    Mixer-to-Client Audio Level Indication    November 2011

## 1. Introduction

The Framework for Conferencing with the Session Initiation Protocol (SIP) defined in [RFC 4353](#) [[RFC4353](#)] presents an overall architecture for multi-party conferencing. Among others, the framework borrows from RTP [[RFC3550](#)] and extends the concept of a mixer entity "responsible for combining the media streams that make up a conference, and generating one or more output streams that are delivered to recipients". Every participant would hence receive, in a flat single stream, media originating from all the others.

Using such centralized mixer-based architectures simplifies support for conference calls on the client side since they would hardly differ from one-to-one conversations. However, the method also introduces a few limitations. The flat nature of the streams that a mixer would output and send to participants makes it difficult for users to identify the original source of what they are hearing.

Mechanisms that allow the mixer to send to participants cues on current speakers (e.g. the CSRC fields in RTP [[RFC3550](#)]) only work for speaking/silent binary indications. There are, however, a number of use cases where one would require more detailed information. Possible examples include the presence of background chat/noise/music/typing, someone breathing noisily in their microphone, or other cases where identifying the source of the disturbance would make it easy to remove it (e.g. by sending a private IM to the concerned party asking them to mute their microphone). A more advanced scenario could involve an intense discussion between multiple participants that the user does not personally know. Audio level information would help better recognize the speakers by associating with them complex (but still human readable) characteristics like loudness and speed for example.

One way of presenting such information in a user friendly manner would be for a conferencing client to attach audio level indicators to the corresponding participant related components in the user

interface. One possible example is displayed in Figure 1 where levels can help users determine that Alice is currently the active speaker, Carol is mute, and Bob and Dave are sending some background noise.

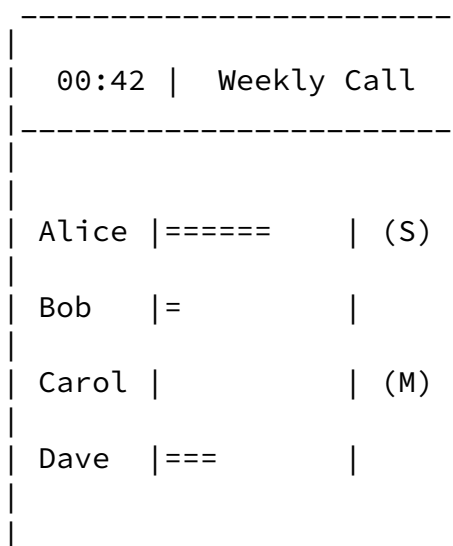


Figure 1: Displaying detailed speaker information to the user by including audio level for every participant.

Implementing a user interface like the above requires analysis of the media sent from other participants. In a conventional audio conference this is only possible for the mixer since all other conference participants are generally receiving a single, flat audio stream and have therefore no immediate way of determining individual audio levels.

This document specifies an RTP extension header that allows such mixers to deliver audio level information to conference participants

by including it directly in the RTP packets transporting the corresponding audio data.

## 2. Terminology

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in [RFC 2119](#) [[RFC2119](#)].

## 3. Protocol Operation

According to [RFC 3550](#) [[RFC3550](#)] a mixer is expected to include in outgoing RTP packets a list of identifiers (CSRC IDs) indicating the sources that contributed to the resulting stream. The presence of such CSRC IDs allows RTP clients to determine, in a binary way, the active speaker(s) in any given moment. RTCP also provides a basic mechanism to map the CSRC IDs to user identities through the CNAME

field. More advanced mechanisms can exist depending on the signaling protocol used to establish and control a conference. In the case of the Session Initiation Protocol [[RFC3261](#)] for example, the Event Package for Conference State [[RFC4575](#)] defines a <src-id> tag which binds CSRC IDs to media streams and SIP URIs.

This document describes an RTP header extension that allows mixers to indicate the audio-level of every contributing conference participant (CSRC) in addition to simply indicating their on/off status. This new header extension uses "General Mechanism for RTP Header Extensions" described in [[RFC5285](#)].

Each instance of this header contains a list of one-octet audio levels expressed in -dBov, with values from 0 to 127 representing 0 to -127 dBov(see Figure 2 and Figure 3). [Appendix A](#) provides a reference implementation indicating one way of obtaining such values from raw audio samples.

Every audio level value pertains to the CSRC identifier located at the corresponding position in the CSRC list. In other words, the first value would indicate the audio level of the conference participant represented by the first CSRC identifier in that packet

and so forth. The number and order of these values MUST therefore match the number and order of the CSRC IDs present in the same packet.

When encoding audio level information, a mixer SHOULD include in a packet information that corresponds to the audio data being transported in that same packet. It is important that these values follow the actual stream as closely as possible. Therefore a mixer SHOULD also calculate the values after the original contributing stream has undergone possible processing such as level normalization, and noise reduction for example.


It can sometimes happen that a conference involves more than a single mixer. In such cases each of the mixers MAY choose to relay the CSRC list and audio-level information they receive from peer mixers (as long as the total CSRC count remains below 16). Given that the maximum audio level is not precisely defined by this specification, it is likely that in such situations average audio levels would be perceptibly different for the participants located behind the different mixers.

#### 4. Audio Levels

The audio level header extension carries the level of the audio in the RTP payload of the packet it is associated with. This


information is carried in an RTP header extension element as defined by the "General Mechanism for RTP Header Extensions" [[RFC5285](#)].

The payload of the audio level header extension element can be encoded using the one-byte or the two-byte header defined in [[RFC5285](#)]. Figure 2 and Figure 3 show sample audio level encodings with each of them.



Sample audio level encoding using the one-byte header format

Figure 2



Sample audio level encoding using the two-byte header format

Figure 3

In the case of the one-byte header format, the 4-bit len field is the number minus one of data bytes (i.e. audio level values) transported in this header extension element following the one-byte header. Therefore, the value zero in this field indicates that one byte of data follows. In the case of the two-byte header format the 8-bit len field contains the exact number of audio levels carried in the extension. [RFC 3550](#) [RFC3550] only allows RTP packets to carry a maximum of 15 CSRC IDs. Given that audio levels directly refer to CSRC IDs, implementations MUST NOT include more than 15 audio level values. The maximum value allowed in the len field is therefore 14 for one-byte header format and 15 for two-byte header format.

Audio levels in this document are defined in the same manner as is

audio noise level in the RTP Payload Comfort Noise specification [RFC3389]. In the comfort noise specification, the overall magnitude of the noise level in comfort noise is encoded into the first byte of the payload, with spectral information about the noise in subsequent bytes. This specification's audio level parameter is defined so as to be identical to the comfort noise payload's noise-level byte.

The magnitude of the audio level itself is packed into the seven

least significant bits of the single byte of the header extension, shown in Figure 2 and Figure 3. The least significant bit of the audio level magnitude is packed into the least significant bit of the byte. The most significant bit of the byte is unused and always set to 0.

The audio level is expressed in -dBov, with values from 0 to 127 representing 0 to -127 dBov. dBov is the level, in decibels, relative to the overload point of the system, i.e. the highest-intensity signal encodable by the payload format. (Note: Representation relative to the overload point of a system is particularly useful for digital implementations, since one does not need to know the relative calibration of the analog circuitry.) For example, in the case of u-law (audio/pcmu) audio [[ITU.G.711](#)], the 0 dBov reference would be a square wave with values +/- 8031. (This translates to 6.18 dBm0, relative to u-law's dBm0 definition in Table 6 of G.711.)

The audio level for digital silence, for example for a muted audio source, MUST be represented as 127 (-127 dBov), regardless of the dynamic range of the encoded audio format.

The audio level header extension only carries the level of the audio in the RTP payload of the packet it is associated with, with no long-term averaging or smoothing applied. That level is measured as a root mean square of all the samples in the measured range.

To simplify implementation of the encoding procedures described here, this specification provides a sample Java implementation (Appendix A) of an audio level calculator that helps obtain such values from raw linear PCM audio samples.

## [5.](#) Signaling Information

The URI for declaring the audio level header extension in an SDP extmap attribute and mapping it to a local extension header identifier is "urn:ietf:params:rtp-hdext:csrc-audio-level". There is no additional setup information needed for this extension (i.e. no extensionattributes).

An example attribute line in the SDP, for a conference might be:



a=extmap:7 urn:ietf:params:rtp-hdext:csrc-audio-level

The above mapping will most often be provided per media stream (in the media-level section(s) of SDP, i.e., after an "m=" line) or globally if there is more than one stream containing audio level indicators in a session.

Presence of the above attribute in the SDP description of a media stream indicates that RTP packets in that stream, which contain the level extension defined in this document, will be carrying them with an ID of 7.

Conferencing clients that support audio level indicators and have no mixing capabilities would not be able to provide content for this audio level extension and would hence have to always include the direction parameter in the "extmap" attribute with a value of "recvonly". Conference focus entities with mixing capabilities can omit the direction or set it to "sendrecv" in SDP offers. Such entities would need to set it to "sendonly" in SDP answers to offers with a "recvonly" parameter and to "sendrecv" when answering other "sendrecv" offers.

This specification only defines use of the audio level extensions in audio streams. They MUST NOT be advertised with other media types such as video or text for example.

The following Figure 4 and Figure 5 show two example offer/answer exchanges between a conferencing client and a focus, and between two conference focus entities.

SDP Offer:

```
v=0
o=alice 2890844526 2890844526 IN IP6 host.example.com
s=-
c=IN IP6 host.example.com
t=0 0
m=audio 49170 RTP/AVP 0 4
a=rtpmap:0 PCMU/8000
a=rtpmap:4 G723/8000
a=extmap:1/recvonly urn:ietf:params:rtp-hdext:csrc-audio-level
```

SDP Answer:

```
v=0
i=A Seminar on the session description protocol
o=conf-focus 2890844730 2890844730 IN IP6 focus.example.net
s=-
c=IN IP6 focus.example.net
t=0 0
m=audio 52544 RTP/AVP 0
a=rtpmap:0 PCMU/8000
a=extmap:1/sendonly urn:ietf:params:rtp-hdext:csrc-audio-level
```

A client-initiated example SDP offer/answer exchange negotiating an audio stream with one-way flow of audio level information.

Figure 4

---

Internet-Draft    Mixer-to-Client Audio Level Indication    November 2011

SDP Offer:

```
v=0
i=Un seminaire sur le protocole de description des sessions
o=fr-focus 2890844730 2890844730 IN IP6 focus.fr.example.net
s=-
c=IN IP6 focus.fr.example.net
t=0 0
m=audio 49170 RTP/AVP 0
a=rtpmap:0 PCMU/8000
a=extmap:1/sendrecv urn:ietf:params:rtp-hdext:csrc-audio-level
```

SDP Answer:

```
v=0
i=A Seminar on the session description protocol
o=us-focus 2890844526 2890844526 IN IP6 focus.us.example.net
s=-
c=IN IP6 focus.us.example.net
t=0 0
m=audio 52544 RTP/AVP 0
a=rtpmap:0 PCMU/8000
a=extmap:1/sendrecv urn:ietf:params:rtp-hdext:csrc-audio-level
```

An example SDP offer/answer exchange between two conference focus entities with mixing capabilities negotiating an audio stream with bidirectional flow of audio level information.

Figure 5

## 6. Security Considerations

1. This document defines a means of attributing audio level to a particular participant in a conference. An attacker may try to modify the content of RTP packets in a way that would make audio activity from one participant appear as coming from another.
2. Furthermore, the fact that audio level values would not be protected even in an SRTP session might be of concern in some

cases where the activity of a particular participant in a conference is confidential. Also, as discussed in [\[I-D.ietf-avtcore-srtp-vbr-audio\]](#), an attacker might be able to infer information about the conversation, possibly with phoneme-level resolution.

3. Both of the above are concerns that stem from the design of the RTP protocol itself and they would probably also apply when using CSRC identifiers the way they were specified in [RFC 3550](#) [[RFC3550](#)]. It is therefore important that according to the needs

of a particular scenario, implementors and deployers consider use of header extension encryption [\[I-D.ietf-avtcore-srtp-encrypted-header-ext\]](#) or a lower level security and authentication mechanism such as IPsec [[RFC4301](#)] for example.

## [7.](#) IANA Considerations

This document defines a new extension URI in the RTP Compact Header Extensions sub-registry of the Real-Time Transport Protocol (RTP) Parameters registry, according to the following data:

Extension URI:	urn:ietf:params:rtp-hdext:csrc-audio-level
Description:	Mixer-to-client audio level indicators
Contact:	emcho@jitsi.org
Reference:	RFC XXXX

Note to the RFC-Editor: please replace "RFC XXXX" by the number of this RFC.

## [8.](#) Acknowledgments

Lyubomir Marinov contributed level measurement and rendering code.

Keith Drage, Roni Even, Miguel A. Garcia, John Elwell, Kevin P. Fleming, Ingemar Johansson, Michael Ramalho, Magnus Westerlund and several others provided helpful feedback over the avt and avtext mailing lists.

Jitsi's participation in this specification is funded by the NLnet

Foundation.

## 9. Changes From Earlier Versions

Note to the RFC-Editor: please remove this section prior to publication as an RFC.

### 9.1. Changes From Draft -05

- o Added proper licensing to the code samples. (Brought up by Sean Turner).
- o Added an informative reference to [RFC 4301](#) (IPsec). (Brought up by Stephen Farrell)

Ivov, et al.

Expires May 19, 2012

[Page 11]

---

Internet-Draft Mixer-to-Client Audio Level Indication November 2011

- o Briefly explained the meaning of the (S) and (M) symbols in the introductory UI example. Added "SDP Offer" and "SDP Answer" subtitles in the SDP examples. Fixed typos. Removed conditional language in the IANA considerations section (Brought up by Dan Romascanu, Lionel Morand and others).
- o Clarified the meaning of "overload point of the system". (Brought up by Robert Sparks).
- o Corrected sample source code to return the level value relative to the overload point, as specified, rather than a level relative to a reference sound pressure level; also to use round-to-nearest rather than truncate to calculate the integer return value. (Brought up by Michael Ramalho.)
- o Updated reference to [[I-D.ietf-avtcore-srtp-vbr-audio](#)].

### 9.2. Changes From Draft -04

- o Fixed problems with missing "s=" attributes and odd RTP port numbers in the SDP examples.

### 9.3. Changes From Draft -03

- o Addressed editorial comments made on the mailing list.

### 9.4. Changes From Draft -02

- o Removed the no-data use case that allowed sending levels in RTP packets. Choosing the right RTP payload type for this use case would have incurred complexity without bringing any real value.
- o Merged the "Header Format" and the "Audio level encoding" sections into a single "Audio Levels" section.
- o Changed encoding related text so that it would cover both the one-byte and the two-byte header formats.
- o Clarified use of root mean square for dBov calculation
- o Added a reference to [[I-D.ietf-avtcore-srtp-vbr-audio](#)] to better explain some "Security Considerations" .
- o Other minor editorial changes.

#### [9.5.](#) Changes From Draft -01

- o Removed code related the AudioLevelRenderer from "APPENDIX A. Reference Implementation" as it was considered an implementation matter by the working group.
- o Modified the AudioLevelCalculator in "APPENDIX A. Reference Implementation" to take overload as a parameter.
- o Clarified non-use of audio levels in video streams
- o Closed the P.56 open issue. It was agreed on IETF 80 that P.56 is mostly about speech levels and the levels transported by the extension defined here should also be able to serve as an

indication for noise.

- o The Open Issues section has been removed as all issues that were in there are now resolved or clarified.
- o Editorial changes for consistency with [[I-D.ietf-avtext-client-to-mixer-audio-level](#)].

#### [9.6.](#) Changes From Draft -00

- o Added code for sound pressure calculation and measurement in "APPENDIX A. Reference Implementation".
- o Changed affiliation for Emil Ivov.
- o Removed "Appendix: Design choices".

## [10.](#) References

### [10.1.](#) Normative References

- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", [BCP 14](#), [RFC 2119](#), March 1997.
- [RFC3550] Schulzrinne, H., Casner, S., Frederick, R., and V. Jacobson, "RTP: A Transport Protocol for Real-Time Applications", STD 64, [RFC 3550](#), July 2003.
- [RFC5285] Singer, D. and H. Desineni, "A General Mechanism for RTP Header Extensions", [RFC 5285](#), July 2008.

## [10.2](#). Informative References

- [I-D.ietf-avtcore-srtp-encrypted-header-ext]  
Lennox, J., "Encryption of Header Extensions in the Secure Real-Time Transport Protocol (SRTP)",  
[draft-ietf-avtcore-srtp-encrypted-header-ext-01](#) (work in progress), October 2011.
- [I-D.ietf-avtcore-srtp-vbr-audio]  
Perkins, C. and J. Valin, "Guidelines for the use of Variable Bit Rate Audio with Secure RTP",  
[draft-ietf-avtcore-srtp-vbr-audio-03](#) (work in progress), July 2011.
- [I-D.ietf-avtext-client-to-mixer-audio-level]  
Lennox, J., Ivov, E., and E. Marocco, "A Real-Time Transport Protocol (RTP) Header Extension for Client-to-Mixer Audio Level Indication",  
[draft-ietf-avtext-client-to-mixer-audio-level-05](#) (work in progress), September 2011.

- [ITU.G.711]  
International Telecommunications Union, "Pulse Code Modulation (PCM) of Voice Frequencies", ITU-T Recommendation G.711, November 1988.
- [ITU.P56.1993]  
International Telecommunications Union, "Objective Measurement of Active Speech Level", ITU-T Recommendation P.56, March 1988.
- [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.

- [RFC3389] Zopf, R., "Real-time Transport Protocol (RTP) Payload for Comfort Noise (CN)", [RFC 3389](#), September 2002.
- [RFC4301] Kent, S. and K. Seo, "Security Architecture for the Internet Protocol", [RFC 4301](#), December 2005.
- [RFC4353] Rosenberg, J., "A Framework for Conferencing with the Session Initiation Protocol (SIP)", [RFC 4353](#), February 2006.
- [RFC4575] Rosenberg, J., Schulzrinne, H., and O. Levin, "A Session Initiation Protocol (SIP) Event Package for Conference State", [RFC 4575](#), August 2006.

## [Appendix A](#). Reference Implementation

This appendix contains Java code for a reference implementation of the level calculation and rendering methods. The code is not normative and by no means the only possible implementation. Its purpose is to help implementors add audio level support to mixers and clients.

The Java code contains an `AudioLevelCalculator` class that calculates the sound pressure level of a signal with specific samples. It can be used in mixers to generate values suitable for the level extension headers.

The implementation is provided in Java but does not rely on any of the language specific and can be easily ported to another.

### [A.1](#). `AudioLevelCalculator.java`

/\*

Copyright (c) 2011 IETF Trust and the persons identified as authors of the code. All rights reserved.



Redistribution and use in source and binary forms, with or without modification, is permitted pursuant to, and subject to the license terms contained in, the Simplified BSD License set forth in [Section 4.c](#) of the IETF Trust's Legal Provisions Relating to IETF Documents (<http://trustee.ietf.org/license-info>).

```
*/  
  
/**  
 * Calculates the audio level of specific samples of a signal relative  
 * to overload.  
 */  
public class AudioLevelCalculator  
{  
  
    /**  
     * Calculates the audio level of a signal with specific  
     * <tt>samples</tt>.  
     *  
     * @param samples the samples of the signal to calculate the audio  
     * level of. The samples are specified as an <tt>int</tt>  
     * array starting at <tt>offset</tt>, extending <tt>length</tt>  
     * number of elements and each <tt>int</tt> element in the specified  
     * range representing a sample of the signal to calculate the audio  
     * level of. Though a sample is provided in the form of an  
     * <tt>int</tt> value, the sample size in bits is determined by the  
     * caller via <tt>overload</tt>.  
     *  
     * @param offset the offset in <tt>samples</tt> at which the samples  
     * start  
     *  
     * @param length the length of the signal specified in  
     * <tt>samples</tt> starting at <tt>offset</tt>  
     *  
     * @param overload the overload (point) of <tt>signal</tt>.  
     * For example, <tt>overload</tt> can be {@link Byte#MAX_VALUE}  
     * for 8-bit signed samples or {@link Short#MAX_VALUE} for  
     * 16-bit signed samples.  
     *  
     * @return the audio level of the specified signal  
     */  
    public static int calculateAudioLevel(  
        int[] samples, int offset, int length,  
        int overload)  
    {
```

```
/*
 * Calculate the root mean square of the signal.
 */
double rms = 0;

for (; offset < length; offset++)
{
    double sample = samples[offset];

    sample /= overload;
    rms += sample * sample;
}
rms = (length == 0) ? 0 : Math.sqrt(rms / length);

/*
 * The audio level is a logarithmic measure of the
 * rms level of an audio sample relative to a reference
 * value and is measured in decibels.
 */
double db;

/*
 * The minimum audio level permitted.
 */
final double MIN_AUDIO_LEVEL = -127;
/*
 * The maximum audio level permitted.
 */
final double MAX_AUDIO_LEVEL = 0;

if (rms > 0)
{
    /*
     * The "zero" reference level is the overload level, which
     * corresponds to 1.0 in this calculation, because the
     * samples are normalized in calculating the RMS.
     */
    db = 20 * Math.log10(rms);

    /*
     * Ensure that the calculated level is within the minimum
     * and maximum range permitted.
     */
    if (db < MIN_AUDIO_LEVEL)
        db = MIN_AUDIO_LEVEL;
    else if (db > MAX_AUDIO_LEVEL)
        db = MAX_AUDIO_LEVEL;
}
```

```
}
```

Internet-Draft    Mixer-to-Client Audio Level Indication    November 2011

```
        else
        {
            db = MIN_AUDIO_LEVEL;
        }

        return (int)Math.round(db);
    }
}
```

AudioLevelCalculator.java

#### Authors' Addresses

Emil Ivov (editor)  
Jitsi  
Strasbourg 67000  
France

Email: [emcho@jitsi.org](mailto:emcho@jitsi.org)

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

Email: [enrico.marocco@telecomitalia.it](mailto:enrico.marocco@telecomitalia.it)

Jonathan Lennox  
Vidyo, Inc.  
433 Hackensack Avenue  
Seventh Floor  
Hackensack, NJ 07601  
US

Email: [jonathan@vidyo.com](mailto:jonathan@vidyo.com)

