Audio/Video Transport Core Maintenance
Internet-Draft
Intended status: Standards Track
Expires: December 7, 2012

A. Williams
Audinate
K. Gross
AVA Networks
R. van Brandenburg
H. Stokking
TNO

June 5, 2012

RTP Clock Source Signalling
draft-williams-avtcore-clksrc-01

Abstract

NTP timestamps are used by several RTP protocols for synchronisation and statistical measurement. This memo specificies SDP signalling identifying NTP timestamp clock sources and SDP signalling identifying the media clock sources in a multimedia session.

Requirements Language

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 [1].

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 December 7, 2012.

Copyright Notice

Copyright (c) 2012 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

RTP protocols use NTP format timestamps to facilitate media stream synchronisation and for providing estimates of round trip time (RTT) and other statistical parameters.

Information about media clock timing exchanged in NTP format timestamps may come from a clock which is synchronised to a global time reference, but this cannot be assumed nor is there a standardised mechanism available to indicate that timestamps are derived from a common reference clock. Therefore, RTP implementations typically assume that NTP timestamps are taken using unsynchronised clocks and must compensate for absolute time differences and rate differences. Without a shared reference clock, RTP can time align flows from the same source at a given receiver using relative timing, however tight synchronisation between two or more different receivers (possibly with different network paths) or between two or more senders is not possible.

High performance AV systems often use a reference media clock distributed to all devices in the system. The reference media clock is often distinct from the the reference clock used to provide timestamps. A reference media clock may be provided along with an audio or video signal interface, or via a dedicated clock signal (e.g. genlock [12] or audio word clock [13]). If sending and receiving media clocks are known to be synchronised to a common reference clock, performance can improved by minimising buffering and avoiding rate conversion.

This specification defines SDP signalling of timestamp clock sources and media reference clock sources.

2. Applications

Timestamp clock source and reference media clock signalling benefit applications requiring synchronised media capture or playout and low latency operation.

Examples include, but are not limited to:

Social TV RTCP for inter-destination media synchronization [6] defines social TV as the combination of media content consumption by two or more users at different devices and locations and real-time communication between those users. An example of Social TV, is where two or more users are watching the same television broadcast at different devices and/or locations, while communicating with each other using text, audio and/or video. A

skew in the media playout of the two or more users can have adverse effects on their experience. A well-known use case here is one friend experiencing a goal in a football match well before or after other friends.

Video Walls A video wall consists of multiple computer monitors, video projectors, or television sets tiled together contiguously or overlapped in order to form one large screen. Each of the screens reproduces a portion of the larger picture. In some implementations, each screen or projector may be individually connected to the network and receive its portion of the overall image from a network-connected video server or video scaler. Screens are refreshed at 50 or 60 hertz or potentially faster. If the refresh is not synchronized, the effect of multiple screens acting as one is broken.

Networked Audio Networked loudspeakers, amplifiers and analogue I/O devices transmitting or receiving audio signals via RTP can be connected to various parts of a building or campus network. Such situations can for example be found in large conference rooms, legislative chambers, classrooms (especially those supporting distance learning) and other large-scale environments such as stadiums. Since humans are more susceptible to differences in audio delay, this use case needs even more accuracy than the video wall use case. Depending on the exact application, the need for accuracy can then be in the range of microseconds [14].

Sensor Arrays Sensor arrays contain many synchronised measurement elements producing signals which are then combined to form an overall measurement. Accurate capture of the phase relationships between the various signals arriving at each element of the array is critically important for proper operation. Examples include towed or fixed sonar arrays, seismic arrays and phased arrays used in radar applications, for instance.

3. Definitions

The definitions of streams, sources and levels of information in SDP descriptions follow the definitions found in Source-Specific Media Attributes in the Session Description Protocol (SDP) [2].

multimedia session A set of multimedia senders and receivers as well as the data streams flowing from senders to receivers. The Session Description Protocol (SDP) [3] describes multimedia sessions.

media stream An RTP session potentially containing more than one RTP source. SDP media descriptions beginning with an "m"-line define the parameters of a media stream.

media source A media source is single stream of RTP packets, identified by an RTP SSRC.

session level Session level information applies to an entire multimedia session. In an SDP description, session-level information appears before the first "m"-line.

media level Media level information applies to a single media stream (RTP session). In an SDP description, media-level information appears after each "m"-line.

source level Source level information applies to a single stream of RTP packets, identified by an RTP SSRC Source-Specific Media Attributes in the Session Description Protocol (SDP) [2] defines how source-level information is included into an SDP session description.

traceable time A clock is considered to provide traceable time if it can be proven to be synchronised to a global time reference. GPS [7] is commonly used to provide a traceable time reference. Some network time synchronisation protocols (e.g. PTP [8], NTP) can explicitly indicate that the master clock is providing a traceable time reference over the network.

4. Timestamp Reference Clock Source Signalling

The NTP timestamps used by RTP are taken by reading a local real-time clock at the sender or receiver. This local clock may be synchronised to another clock (time source) by some means or it may be unsynchronised. A variety of methods are available to synchronise local clocks to a reference time source, including network time protocols (e.g. NTP [9]) and radio clocks like GPS [7].

The following sections describe and define SDP signalling, indicating whether and how the local timestamping clock in an RTP sender/receiver is synchronised to a reference clock.

4.1. Clock synchronization

Two or more local clocks that are sufficiently synchronised will produce timestamps for a given RTP event can be used as if they cam from the same clock. Providing they are sufficiently synchronised, timestamps produced in one RTP sender/receiver can be directly

compared to a local clock in another RTP sender/receiver. The timestamps produced by synchronized local clocks in two or more RTP senders/receivers can be directly compared.

The accuracy of synchronization required is application dependent. See Applications (Section 2) section for a discussion of applications and their corresponding requirements. To serve as a reference clock, clocks must minimally be syntonized (exactly frequency matched) to one another.

Sufficient synchronization can typically be achieving by using a network time protocol (e.g. NTP, 802.1AS, IEEE 1588-2008) to synchronize all devices to a single master clock.

Another apporach is to use clocks providing a global time reference (e.g. GPS, Gallileo). This concept may be used in conjunction with network time protocols as some protocols (e.g. PTP, NTP) allow master clocks to indicate explicitly that they are "traceable" back to a global time reference.

4.2. Identifying NTP Reference Clocks

A single NTP server is identified by hostname (or IP address) and an optional port number. If the port number is not indicated, it is assumed to be the standard NTP port (123).

Two or more NTP servers may be listed at the same level in the session description to indicate that they are interchangeable. RTP senders/receivers can use any of the listed NTP servers to govern a local clock that is equivalent to a local clock slaved to a different server.

4.3. Identifying PTP Reference Clocks

The IEEE 1588 Precision Time Protocol (PTP) family of clock synchronisation protocols provides a shared reference clock in an network - typically a LAN. IEEE 1588 provides sub-microsecond synchronisation between devices on a LAN and typically locks within seconds at startup. With support from Ethernet switches, IEEE 1588 protocols can achieve nanosecond timing accuracy in LANs. Network interface chips and cards supporting hardware time-stamping of timing critical protocol messages are also available.

Three flavours of IEEE 1588 are in use today:

- o IEEE 1588-2002 [10]: the original "Standard for a Precision Clock Synchronization Protocol for Networked Measurement and Control Systems". This is also known as IEEE1588v1 or PTPv1.

- IEEE 1588-2008 [8]: the second version of the "Standard for a Precision Clock Synchronization Protocol for Networked Measurement and Control Systems". This is a revised version of the original IEEE1588-2002 standard and is also known as IEEE1588v2 or PTPv2. IEEE 1588-2008 is not protocol compatible with IEEE 1588-2002.

- IEEE 802.1AS [11]: "Timing and Synchronization for Time Sensitive Applications in Bridged Local Area Networks". This is a Layer-2 only profile of IEEE 1588-2008 for use in Audio/Video Bridged LANs.

Each IEEE 1588 clock is identified by a globally unique EUI-64 called a "ClockIdentity". A slave clock using one of the IEEE 1588 family of network time protocols acquires the ClockIdentity/EUI-64 of the grandmaster clock that is the ultimate source of timing information for the network. A master clock which is itself slaved to another master clock passes the grandmaster ClockIdentity through to its slaves.

Several instances of the IEEE 1588 protocol may operate independently on a single network, forming distinct PTP network protocol domains, each of which may have a different grandmaster clock. As the IEEE 1588 standards have developed, the definition of PTP domains has changed. IEEE 1588-2002 identifies protocol subdomains by a textual name, but IEEE 1588-2008 identifies protocol domains using a numeric domain number. 802.1AS is a Layer-2 profile of IEEE 1588-2008 supporting a single numeric clock domain (0).

When PTP subdomains are signalled via SDP, senders and receivers SHOULD check that both grandmaster ClockIdentity and PTP subdomain match when determining clock equivalence.

The PTP protocols employ a distributed election protocol called the "Best Master Clock Algorithm" (BMCA) to determine the active clock master. The clock master choices available to BMCA can be restricted or favourably biased by setting stratum values, preferred master clock bits, or other parameters to influence the election process. In some systems it may be desirable to limit the number of possible PTP clock masters to avoid re-signalling timestamp clock sources when the clock master changes.

4.4. Identifying Global Reference Clocks

Global reference clocks provide a source of traceable time, typically via a hardware radio receiver interface. Examples include GPS and Galileo. Apart from the name of the reference clock system, no further identification is required.

4.5. Other Reference Clocks

At the time of writing, it is common for RTP senders/receivers not to synchronise their local timestamp clocks to a shared master. An unsynchronised clock such as a quartz oscillator is identified as a "local" reference clock.

In some systems, all RTP senders/receivers may use a timestamp clock synchronised to a reference clock that is not provided by one of the methods listed above. Examples may include the reference time information provided by digital television or cellular services. These sources are identified as "private" reference clocks. All RTP senders/receivers in a session using a private reference clock are assumed to have a mechanism outside this specification confirming that their local timestamp clocks are equivalent.

4.6. Traceable Reference Clocks

A timestamp clock source may be labelled "traceable" if it is known to be sourced from a global time reference such as TAI or UTC. Providing adjustments are made for differing time bases, timestamps taken using clocks synchronised to a traceable time source can be directly compared even if the clocks are synchronised to different sources or via different mechanisms.

Since all NTP and PTP servers providing traceable time can be directly compared, it is not necessary to identify traceable time servers by protocol address or other identifiers.

4.7. Synchronisation Confidence

Network time protocol services periodically exchange timestamped messages between servers and clients. Assuming RTP sender/receiver clocks are based on commonly available quartz crystal hardware which is subject to drif, tight synchronisation requires frequent exchange of synchronisation messages.

Unfortunately, in some implementations, it is not possible to control the frequency of synchronisation messages nor is it possible to discover when the last sychronisation message occurred. In order to provide a measure of confidence that the timestamp clock is sufficiently synchronised, an optional timestamp may be included in the SDP clock source signalling. In addition, the frequency of synchronisation message may also be signalled.

The optional timestamp and synchronisation frequency parameters provide an indication of synchronisation quality to the receiver of those parameters. If the synchronisation confidence timestamp is far

from the timestamp clock at the receiver of the parameters, it can be assumed that synchronisation has not occured recently or the timestamp reference clock source cannot be contacted. In this case, the receiver can take action to prevent unsynchronised playout or may fall back to assuming that the timestamp clocks are not synchronised.

Synchronisation frequency is expressed as a signed (two's-compliment) 8-bit field which is the base-2 logarithm of the frequency in Hz. The synchronisation frequencies represented by this field range from 2^-128 Hz to 2^+127 Hz. The field value of 0 corresponds to an update frequency of 1 Hz.

4.8. SDP Signalling of Timestamp Clock Source

Specification of the timestamp reference clock source may be at any or all levels (session, media or source) of an SDP description (see level definitions (Section 3) earlier in this document for more information).

Timestamp clock source signalling included at session-level provides default parameters for all RTP sessions and sources in the session description. More specific signalling included at the media level overrides default session level signalling. Further, source-level signalling overrides timestamp clock source signalling at the enclosing media level and session level.

If timestamp clock source signalling is included anywhere in an SDP description, it must be properly defined for all levels in the description. This may simply be achieved by providing default signalling at the session level.

Timestamp reference clock parameters may be repeated at a given level (i.e. for a session or source) to provide information about additional servers or clock sources. If the attribute is repeated at a given level, all clocks described at that level are assumed to be equivalent. Traceable clock sources MUST NOT be mixed with non-traceable clock sources at any given level. Unless synchronisation confidence information is available for each of the reference clocks listed at a given level, it SHOULD only be included with the first reference clock source attribute at that level.

Note that clock source parameters may change from time to time, for example, as a result of a PTP clock master election. The SIP [4] protocol supports re-signalling of updated SDP information, however other protocols may require additional notification mechanisms.

Figure 1 shows the ABNF [5] grammar for the SDP reference clock source information.

```
timestamp-refclk = "a=ts-refclk:" clksrc [ SP sync-confidence ] CRLF

clksrc = ntp / ptp / gps / gal / local / private

ntp             =  "ntp=" ntp-server-addr
ntp-server-addr =  host [ ":" port ]
ntp-server-addr =/ "traceable" )

ptp             =  "ptp=" ptp-version ":" ptp-gmid [":" ptp-domain]
ptp-version     =  "IEEE1588-2002"
ptp-version     =/ "IEEE1588-2008"
ptp-version     =/ "IEEE802.1AS-2011"
ptp-gmid        =  EUI64
ptp-gmid        =/ "traceable"
ptp-domain      =  ptp-domain-name / ptp-domain-nmbr
ptp-domain-name =  "domain-name=" 16ptp-domain-char
ptp-domain-char =  %x21-7E / %x00
                   ; allowed characters: 0x21-0x7E (IEEE 1588-2002)
ptp-domain-nmbr =  "domain-nmbr=" %x00-7f
                   ; allowed number range: 0-127 (IEEE 1588-2008)

gps      =  "gps"
gal      =  "gal"
local    =  "local"
private  =  "private" [ ":" "traceable" ]


sync-confidence = sync-timestamp [SP sync-frequency]

sync-timestamp  = sync-date SP sync-time SP sync-UTCoffset

sync-date       = 4DIGIT "-" 2DIGIT "-" 2DIGIT
                  ; yyyy-mm-dd (e.g., 1982-12-02)

sync-time       = 2DIGIT ":" 2DIGIT ":" 2DIGIT "." 3DIGIT
                  ; 00:00:00.000 - 23:59:59.999

sync-UTCoffset  = ( "+" / "-" ) 2DIGIT ":" 2DIGIT
                  ; +HH:MM or -HH:MM

sync-frequency  = 2HEXDIG
                  ; If N is the field value, HZ=2^(N-127)


host          = hostname / IPv4address / IPv6reference

hostname      =  *( domainlabel "." ) toplabel [ "." ]
toplabel      =  ALPHA / ALPHA *( alphanum / "-" ) alphanum
```

```
domainlabel   =  alphanum
              =/ alphanum *( alphanum / "-" ) alphanum

IPv4address   =  1*3DIGIT "." 1*3DIGIT "." 1*3DIGIT "." 1*3DIGIT
IPv6reference =  "[" IPv6address "]"
IPv6address   =  hexpart [ ":" IPv4address ]
hexpart       =  hexseq / hexseq "::" [ hexseq ] / "::" [ hexseq ]
hexseq        =  hex4 *( ":" hex4)
hex4          =  1*4HEXDIG

port = 1*DIGIT

EUI-64 = 7(2HEXDIG "-") 2HEXDIG
```

Figure 1: Timestamp Reference Clock Source Signalling

4.8.1. Examples

Figure 2 shows an example SDP description with a timestamp reference clock source defined at the session level.

```
v=0
o=jdoe 2890844526 2890842807 IN IP4 10.47.16.5
s=SDP Seminar
i=A Seminar on the session description protocol
u=http://www.example.com/seminars/sdp.pdf
e=j.doe@example.com (Jane Doe)
c=IN IP4 224.2.17.12/127
t=2873397496 2873404696
a=recvonly
a=ts-refclk:ntp=traceable
m=audio 49170 RTP/AVP 0
m=video 51372 RTP/AVP 99
a=rtpmap:99 h263-1998/90000
```

Figure 2: Timestamp reference clock definition at the session level

Figure 3 shows an example SDP description with timestamp reference clock definitions at the media level overriding the session level defaults. Note that the synchronisation confidence timestamp appears on the first attribute at the media level only.

```
v=0
o=jdoe 2890844526 2890842807 IN IP4 10.47.16.5
s=SDP Seminar
i=A Seminar on the session description protocol
u=http://www.example.com/seminars/sdp.pdf
e=j.doe@example.com (Jane Doe)
c=IN IP4 224.2.17.12/127
t=2873397496 2873404696
a=recvonly
a=ts-refclk:local
m=audio 49170 RTP/AVP 0
a=ts-refclk:ntp=203.0.113.10 2011-02-19 21:03:20.345+01:00
a=ts-refclk:ntp=198.51.100.22
m=video 51372 RTP/AVP 99
a=rtpmap:99 h263-1998/90000
a=ts-refclk:ptp=IEEE802.1AS-2011:39-A7-94-FF-FE-07-CB-D0
```

Figure 3: Timestamp reference clock definition at the media level

Figure 4 shows an example SDP description with a timestamp reference clock definition at the source level overriding the session level default.

```
v=0
o=jdoe 2890844526 2890842807 IN IP4 10.47.16.5
s=SDP Seminar
i=A Seminar on the session description protocol
u=http://www.example.com/seminars/sdp.pdf
e=j.doe@example.com (Jane Doe)
c=IN IP4 224.2.17.12/127
t=2873397496 2873404696
a=recvonly
a=ts-refclk:local
m=audio 49170 RTP/AVP 0
m=video 51372 RTP/AVP 99
a=rtpmap:99 h263-1998/90000
a=ssrc:12345 ts-refclk:ptp=IEEE802.1AS-2011:39-A7-94-FF-FE-07-CB-D0
```

Figure 4: Timestamp reference clock signalling at the source level

5. Media Clock Source Signalling

The media clock source for a stream determines the timebase used to advance the RTP timestamps included in RTP packets. The media clock may be asynchronously generated by the sender, it may be generated in fixed relationship to the reference clock or it may be generated with respect to another stream on the network (which is presumably being

received by the sender).

5.1. Asynchronously Generated Media Clock

In the simplest sender implementation, the sender generates media by sampling audio or video according to a free-running local clock. The RTP timestamps in media packets are advanced according to this media clock and packet transmission is typically timed to regular intervals on this timeline. The sender may or may not include an NTP timestamp in sender reports to allow mapping of this asynchronous media clock to a reference clock.

The asynchronously generated media clock is the assumed mode of operation when there is no signalling of media clock source. Alternatively, asynchronous media clock me be signaled.

```
a=mediaclk:sender
```

5.2. Direct-Referenced Media Clock

A media clock may be directly derived from a reference clock. For this case it is required that a reference clock be specified. The signalling indicates a media clock offset value at the epoch (time of origin) of the reference clock. A rate for the media clock may also be specified. If include, the rate specification here overrides that specified or implied by the media description. If omitted, the rate is assumed to be the exact rate used by the media format. For example, the media clock for an 8 kHz G.711 audio stream will advance exactly 8000 units for each second advance in the reference clock from which it is derived.

The rate may optionally be expressed as the ratio of two integers. This provision is useful for accomodating certain "oddball rates" associated with NTSC video.

```
a=mediaclk:offset=<offset>[ rate=<rate numerator>[/<rate
denominator>]]
```

5.3. Stream-Referenced Media Clock

The media clock for an outgoing stream may be generated based on the media clock received with an incoming stream. In this case, the signalling identifies the session and the stream source. The received media clock is converted to a real-time clock which is used to generate outgoing media clocks. In this way, the format of the reference stream does not need to match the format of the outgoing stream.

A reference stream can be either another RTP stream or AVB stream based on the IEEE 1722 standard. An RTP stream is identified by destination IP address (for a multicast stream) or source IP address (for a unicast stream), destination port number and CNAME of the source.

```
a=mediaclk:rtp=<connection address>:<port> <CNAME>
```

An IEEE 1722 stream is identified by its StreamID, an EUI-64.

```
a=mediaclk:IEEE1722=<StreamID>
```

5.4. Signalling Grammar

Specification of the media clock source may be at any or all levels (session, media or source) of an SDP description (see level definitions (Section 3) earlier in this document for more information).

Media clock source signalling included at session level provides default parameters for all RTP sessions and sources in the session description. More specific signalling included at the media level overrides default session level signalling. Further, source-level signalling overrides media clock source signalling at the enclosing media level and session level.

Media clock source signalling may be present or absent on a per-stream basis. In the absence of media clock source signals, receivers assume an asynchronous media clock generated by the sender.

Media clock source parameters may be repeated at a given level (i.e. for a session or source) to provide information about additional clock sources. If the attribute is repeated at a given level, all clocks described at that level are comparable clock sources and may be used interchangeably.

Figure 5 shows the ABNF [5] grammar for the SDP media clock source information.

```
timestamp-mediaclk = "a=mediaclk:" mediaclock

mediaclock = refclk / rtp / streamid / sender

refclk = "offset=" 1*DIGIT [ SP "rate=" 1*DIGIT [ "/" 1*DIGIT ] ]

rtp = "rtp=" nettype SP addrtype SP connection-address SP port SP cname

streamid = "IEEE1722=" EUI-64
```

```
   sender = "sender"

   cname = non-ws-string

   nettype = token
           ;typically "IN"

   addrtype = token
           ;typically "IP4" or "IP6"

   token-char = %x21 / %x23-27 / %x2A-2B / %x2D-2E / %x30-39 / %x41-5A
                / %x5E-7E

   token = 1*(token-char)

   connection-address =  multicast-address / unicast-address

   unicast-address = IP4-address / IP6-address / FQDN / extn-addr

   multicast-address = IP4-multicast / IP6-multicast / FQDN / extn-addr

   IP4-multicast = m1 3( "." decimal-uchar ) "/" ttl [ "/" integer ]
           ; IPv4 multicast addresses may be in the
           ; range 224.0.0.0 to 239.255.255.255

   m1 = ("22" ("4"/"5"/"6"/"7"/"8"/"9")) / ("23" DIGIT )

   IP6-multicast = hexpart [ "/" integer ]
           ; IPv6 address starting with FF

   FQDN = 4*(alpha-numeric / "-" / ".")
           ; fully qualified domain name as specified
           ; in RFC 1035 (and updates)

   IP4-address = b1 3("." decimal-uchar)

   b1 = decimal-uchar
           ; less than "224"

   ; The following is consistent with RFC 2373 [30], Appendix B.
   IP6-address = hexpart [ ":" IP4-address ]

   hexpart = hexseq / hexseq "::" [ hexseq ] / "::" [ hexseq ]

   hexseq = hex4 *( ":" hex4)

   hex4 = 1*4HEXDIG
```

```
; Generic for other address families
extn-addr = non-ws-string

non-ws-string = 1*(VCHAR/%x80-FF)
        ;string of visible characters

port = 1*DIGIT

EUI-64 = 7(2HEXDIG "-") 2HEXDIG
```

Figure 5: Media Clock Source Signalling

5.5. Examples

Figure 6 shows an example SDP description 8 channels of 24-bit, 48 kHz audio transmitted as a multicast stream. Media clock is derived directly from an IEEE 1588-2008 reference.

```
v=0
o=- 1311738121 1311738121 IN IP4 192.168.1.1
c=IN IP4 239.0.0.2/255
s=
t=0 0
m=audio 5004 RTP/AVP 96
a=rtpmap:96 L24 L24/48000/8
a=sendonly
a=ts-refclk:ptp=IEEE1588-2008:39-A7-94-FF-FE-07-CB-D0:0
a=mediaclk:offset=963214424
```

Figure 6: Media clock directly referenced to IEEE 1588-2008

Figure 7 shows an example SDP description 2 channels of 24-bit, 44056 kHz NTSC "pull-down" media clock derived directly from an IEEE 1588-2008 reference clock

```
v=0
o=- 1311738121 1311738121 IN IP4 192.168.1.1
c=IN IP4 239.0.0.2/255
s=
t=0 0
m=audio 5004 RTP/AVP 96
a=rtpmap:96 L24 L24/44056/2
a=sendonly
a=ts-refclk:ptp=IEEE1588-2008:39-A7-94-FF-FE-07-CB-D0:0
a=mediaclk:offset=963214424 rate=44100000/1001
```

Figure 7: "Oddball" sample rate directly refernced to IEEE 1588-2008

Figure 8 shows the same 48 kHz audio transmission from Figure 6 with media clock derived from another RTP multicast stream. The stream providing the media clock must use the same reference clock as this stream that references it.

```
v=0
o=- 1311738121 1311738121 IN IP4 192.168.1.1
c=IN IP4 224.2.228.230/32
s=
t=0 0
m=audio 5004 RTP/AVP 96
a=rtpmap:96 L24 L24/48000/2
a=sendonly
a=ts-refclk:ptp=IEEE1588-2008:39-A7-94-FF-FE-07-CB-D0:0
a=mediaclk:rtp=IN IP4 239.0.0.1 5004 00:60:2b:20:12:if
```

Figure 8: Stream media clock derived from another RTP multicast stream

Figure 9 shows the same 48 kHz audio transmission from Figure 6 with media clock derived from an IEEE 1722 AVB stream. The stream providing the media clock must be synchronized with the IEEE 1588-2008 reference clock used by this stream.

```
v=0
o=- 1311738121 1311738121 IN IP4 192.168.1.1
c=IN IP4 224.2.228.230/32
s=
t=0 0
m=audio 5004 RTP/AVP 96
a=rtpmap:96 L24 L24/48000/2
a=sendonly
a=ts-refclk:ptp=IEEE1588-2008:39-A7-94-FF-FE-07-CB-D0:0
a=mediaclk:IEEE1722=38-D6-6D-8E-D2-78-13-2F
```

Figure 9: Stream media clock derived from another RTP multicast stream

6. IANA Considerations

The SDP attribute "ts-refclk" defined by this document is registered with the IANA registry of SDP Parameters as follows:

SDP Attribute ("att-field"):

Attribute name: ts-refclk

Long form: Timestamp reference clock source

Type of name: att-field

Type of attribute: session, media and source level

Subject to charset: no

Purpose: See section 4 of this document

Reference: This document

Values: see this document and registrations below

The attribute has an extensible parameter field and therefore a registry for these parameters is required. This document creates an IANA registry called the Timestamp Reference Clock Source Parameters Registry. It contains the six parameters defined in Figure 1: "ntp", "ptp", "gps", "gal", "local", "private".

The SDP attribute "mediaclk" defined by this document is registered with the IANA registry of SDP Parameters as follows:

SDP Attribute ("att-field"):

Attribute name: mediaclk

Long form: Media clock source

Type of name: att-field

Type of attribute: session abd media level

Subject to charset: no

Purpose: See section 6 of this document

Reference: This document

Values: see this document and registrations below

The attribute has an extensible parameter field and therefore a

registry for these parameters is required. This document creates an IANA registry called the Media Clock Source Parameters Registry. It contains the three parameters defined in Figure 5: "refclk", "ssrc", "sender".

7. References

7.1. Normative References

[1] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.

[2] Lennox, J., Ott, J., and T. Schierl, "Source-Specific Media Attributes in the Session Description Protocol (SDP)", RFC 5576, June 2009.

[3] Handley, M., Jacobson, V., and C. Perkins, "SDP: Session Description Protocol", RFC 4566, July 2006.

[4] 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.

[5] Crocker, D., Ed. and P. Overell, "Augmented BNF for Syntax Specifications: ABNF", RFC 2234, November 1997.

7.2. Informative References

[6] Brandenburg, R., Stokking, H., Boronat, F., Montagud, M., and K. Gross, "RTCP for inter-destination media synchronization", draft-ietf-avtcore-idms-04 (work in progress), May 2012.

[7] Global Positioning Systems Directorate, "Navstar GPS Space Segment/Navigation User Segment Interfaces", September 2011.

[8] Institute of Electrical and Electronics Engineers, "1588-2008 - IEEE Standard for a Precision Clock Synchronization Protocol for Networked Measurement and Control Systems", IEEE Std 1588-2008, 2008, <http://standards.ieee.org/findstds/standard/1588-2008.html>.

[9] Mills, D., Martin, J., Burbank, J., and W. Kasch, "Network Time Protocol Version 4: Protocol and Algorithms Specification", RFC 5905, June 2010.

[10] Institute of Electrical and Electronics Engineers, "1588-2002 - IEEE Standard for a Precision Clock Synchronization Protocol

for Networked Measurement and Control Systems", IEEE Std 1588-2002, 2002,
<http://standards.ieee.org/findstds/standard/1588-2002.html>.

[11] "Timing and Synchronization for Time-Sensitive Applications in Bridged Local Area Networks",
<http://standards.ieee.org/findstds/standard/802.1AS-2011.html>.

URIs

[12] <http://en.wikipedia.org/wiki/Genlock>

[13] <http://en.wikipedia.org/wiki/Word_clock>

[14] <http://www.ieee802.org/1/files/public/docs2007/as-dolsen-time-accuracy-0407.pdf>

Authors' Addresses

Aidan Williams
Audinate
Level 1, 458 Wattle St
Ultimo, NSW 2007
Australia

Phone: +61 2 8090 1000
Fax: +61 2 8090 1001
Email: aidan.williams@audinate.com
URI: http://www.audinate.com/

Kevin Gross
AVA Networks
Boulder, CO
US

Email: kevin.gross@avanw.com
URI: http://www.avanw.com/

Ray van Brandenburg
TNO
Brassersplein 2
Delft 2612CT
the Netherlands

Phone: +31-88-866-7000
Email: ray.vanbrandenburg@tno.nl

Hans Stokking
TNO
Brassersplein 2
Delft 2612CT
the Netherlands

Phone:
Email: stokking@tno.nl